


Multivariable Calculus



Jerry Shurman
Reed College

Contents

Preface	ix
1 Results from One-Variable Calculus	1
1.1 The Real Number System	1
1.2 Foundational and Basic Theorems	3
1.3 Taylor's Theorem	4
2 Euclidean Space	17
2.1 Algebra: Vectors	17
2.2 Geometry: Length and Angle	24
2.3 Analysis: Continuous Mappings	33
2.4 Topology: Compact Sets and Continuity	42
2.5 Review of the One-Variable Derivative	49
2.6 Summary	51
3 Linear Mappings and Their Matrices	53
3.1 Linear Mappings	54
3.2 Operations on Matrices	65
3.3 The Inverse of a Linear Mapping	70
3.4 Inhomogeneous Linear Equations	78
3.5 The Determinant: Characterizing Properties and Their Consequences	79
3.6 The Determinant: Uniqueness and Existence	85
3.7 An Explicit Formula for the Inverse	96
3.8 Geometry of the Determinant: Volume	98
3.9 Geometry of the Determinant: Orientation	108
3.10 The Cross Product, Lines, and Planes in \mathbf{R}^3	110
3.11 Summary	118

4	The Derivative	119
4.1	The Derivative Redefined	119
4.2	Basic Results	127
4.3	Calculating the Derivative	135
4.4	Higher Order Derivatives	145
4.5	Extreme Values	150
4.6	Directional Derivatives and the Gradient	159
4.7	Summary	168
5	Inverse and Implicit Functions	169
5.1	Preliminaries	170
5.2	The Inverse Function Theorem	174
5.3	The Implicit Function Theorem	181
5.4	Lagrange Multipliers: Geometric Motivation and Specific Examples	195
5.5	Lagrange Multipliers: Analytic Proof and General Examples ..	205
5.6	Summary	213
6	Integration	215
6.1	Machinery: Boxes, Partitions, and Sums	215
6.2	Definition of the Integral	224
6.3	Continuity and Integrability	230
6.4	Integration of Functions of One Variable	237
6.5	Integration Over Nonboxes	243
6.6	Fubini's Theorem	252
6.7	Change of Variable	265
6.8	Topological Preliminaries for the Change of Variable Theorem	279
6.9	Proof of the Change of Variable Theorem	285
6.10	Summary	295
7	Approximation by Smooth Functions	297
7.1	Spaces of Functions	298
7.2	Pulse Functions	302
7.3	Convolution	305
7.4	Smooth Approximate Identity and Convolution	311
7.5	Known-Integrable Functions	317
8	Integration of Differential Forms	323
8.1	Integration of Functions Over Surfaces	324
8.2	Flow and Flux Integrals	332
8.3	Differential Forms Syntactically and Operationally	337
8.4	Examples: 1-forms	340
8.5	Examples: 2-forms on \mathbf{R}^3	344
8.6	Algebra of Forms: Basic Properties	351
8.7	Algebra of Forms: Multiplication	352

8.8 Algebra of Forms: Differentiation 355

8.9 Algebra of Forms: the Pullback 361

8.10 Change of Variable for Differential Forms 370

8.11 Cubes and Chains 371

8.12 Geometry of Chains: the Boundary Operator 374

8.13 The Generalized Fundamental Theorem of Integral Calculus .. 379

8.14 The Classical Theorems 384

8.15 Divergence and Curl in Polar Coordinates 389

8.16 Summary 397

Index 399

Preface

This is the text for a two-semester multivariable calculus course. The setting is n -dimensional Euclidean space, with the material on differentiation culminating in the Inverse Function Theorem and its consequences, and the material on integration culminating in the Generalized Fundamental Theorem of Integral Calculus (often called Stokes's Theorem) and some of its consequences in turn. The prerequisite is a proof-based course in one-variable calculus. Some familiarity with the complex number system and complex mappings is occasionally assumed as well, but the reader can get by without it.

The book's aim is to use multivariable calculus to teach mathematics as a blend of reasoning, computing, and problem-solving to a cross section of serious liberal arts students, doing justice to the structure, the details, and the scope of the ideas. To this end, I have tried to write in a style that communicates intent early in the discussion of each topic rather than proceeding coyly from opaque definitions. Also, I have tried occasionally to speak to the pedagogy of mathematics and its effect on the process of learning the subject. Most importantly, I have tried to spread the weight of exposition among diagrams, formulas, and words. The premise is that the reader is ready to do mathematics resourcefully by marshaling the skills of

- geometric intuition (the visual cortex being quickly instinctive),
- algebraic manipulation (symbol-patterns being precise and robust),
- and incisive use of natural language (slogans that encapsulate central ideas enabling a large-scale grasp of the subject).

Thinking in these ways renders mathematics coherent, inevitable, and fluent.

In my own student days, I learned this material from books by Apostol, Buck, Rudin, and Spivak, books that thrilled me. My debt to those sources pervades these pages. There are many other fine books on the subject as well, such as the more recent one by Hubbard and Hubbard.

By way of a warm-up, chapter 1 reviews some ideas from one-variable calculus, and then covers the one-variable Taylor's Theorem in detail.

Chapters 2 and 3 cover what might be called multivariable pre-calculus, introducing the requisite algebra, geometry, analysis, and topology of Euclidean space, and the requisite linear algebra, for the calculus to follow. A pedagogical theme of these chapters is that mathematical objects can be better understood from their characterizations than from their constructions. Vector geometry follows from the intrinsic (coordinate-free) algebraic properties of the vector inner product, with no reference to the inner product formula. The fact that passing a closed and bounded subset of Euclidean space through a continuous mapping gives another such set is clear once such sets are characterized in terms of sequences. The multiplicativity of the determinant, and the fact that the determinant indicates whether a linear mapping is invertible, are consequences of the determinant's characterizing properties. The geometry of the cross product follows from its intrinsic algebraic characterization. Furthermore, the only possible formula for the inner product, or for the determinant, or for the cross product, is dictated by the relevant properties. As far as the theory is concerned, the only role of the formula to show that an object with the desired properties exists at all. The intent here is that the student who is introduced to mathematical objects via their characterizations will see quickly how the objects work, and that how they work makes their constructions inevitable.

In the same vein, chapter 4 characterizes the multivariable derivative as a well approximating linear mapping. The chapter then solves some multivariable problems that have one-variable counterparts. Specifically, the multivariable chain rule helps with change of variable in partial differential equations, a multivariable analogue of the max/min test helps with optimization, and the multivariable derivative of a scalar-valued function helps to find tangent planes and trajectories.

Chapter 5 uses the results of the three chapters preceding it to prove the Inverse Function Theorem, then the Implicit Function Theorem as a corollary, and finally the Lagrange Multiplier Criterion as a consequence of the Implicit Function Theorem. Lagrange multipliers help with a type of multivariable optimization problem that has no one-variable analogue, optimization with constraints. For example, given two curves in space, what pair of points—one on each curve—is closest to each other? Not only does this problem have six variables (the three coordinates of each point), but furthermore they are not fully independent: the first three variables must specify a point on the first curve, and similarly for the second three. In this problem, x_1 through x_6 vary though a subset of six-dimensional space, conceptually a two-dimensional subset (one degree of freedom for each curve) that is bending around in the ambient six dimensions, and we seek points of this subset where a certain function of x_1 through x_6 is optimized. That is, optimization with constraints can be viewed as a beginning example of calculus on curved spaces.

For another example, let n be a positive integer, and let e_1, \dots, e_n be positive numbers with $e_1 + \dots + e_n = 1$. Maximize the function

$$f(x_1, \dots, x_n) = x_1^{e_1} \cdots x_n^{e_n}, \quad x_i \geq 0 \text{ for all } i,$$

subject to the constraint that

$$e_1 x_1 + \cdots + e_n x_n = 1.$$

As in the previous paragraph, since this problem involves one condition on the variables x_1 through x_n , it can be viewed as optimizing over an $(n - 1)$ -dimensional space inside n dimensions. The problem may appear unmotivated, but its solution leads quickly to a generalization of the arithmetic-geometric mean inequality $\sqrt{ab} \leq (a + b)/2$ for all nonnegative a and b ,

$$a_1^{e_1} \cdots a_n^{e_n} \leq e_1 a_1 + \cdots + e_n a_n \quad \text{for all nonnegative } a_1, \dots, a_n.$$

Moving to integral calculus, chapter 6 introduces the integral of a scalar-valued function of many variables, taken over a domain of its inputs. When the domain is a box, the definitions and the basic results are essentially the same as for one variable. However, in multivariable calculus we want to integrate over regions other than boxes, and ensuring that we can do so takes a little work. After this is done, the chapter proceeds to two main tools for multivariable integration, Fubini's Theorem and the Change of Variable Theorem. Fubini's Theorem reduces one n -dimensional integral to n one-dimensional integrals, and the Change of Variable Theorem replaces one n -dimensional integral with another that may be easier to evaluate. Using these techniques one can show, for example, that the ball of radius r in n dimensions has volume

$$\text{vol}(B_n(r)) = \frac{\pi^{n/2}}{(n/2)!} r^n, \quad n = 1, 2, 3, 4, \dots$$

The meaning of the $(n/2)!$ in the display when n is odd is explained by a function called the gamma function. The sequence begins

$$2r, \quad \pi r^2, \quad \frac{4}{3} \pi r^3, \quad \frac{1}{2} \pi^2 r^4, \quad \dots$$

Chapter 7 discusses the fact that continuous functions, or once-differentiable functions, or twice-differentiable functions, are well approximated by smooth functions, meaning functions that can be differentiated endlessly. The approximation technology is an integral called the convolution. With approximation by convolution in hand, we feel free to assume in the sequel that functions are smooth.

Chapter 8 presents the integration of differential forms. This subject poses the pedagogical dilemma that fully describing its structure requires an investment in machinery untenable for students who are seeing it for the first time, whereas describing it purely operationally is unmotivated. The approach here begins with the integration of functions over k -dimensional surfaces in n -dimensional space, a natural thing to want to do, with a natural definition

of how to do it suggesting itself. For certain such integrals, called flow and flux integrals, the integrand takes a particularly workable form consisting of sums of determinants of derivatives. It is easy to see what other integrands—including integrands suitable for n -dimensional integration in the sense of chapter 6, and including functions in the usual sense—have similar features. These integrands can be uniformly described in algebraic terms as objects called differential forms. That is, differential forms comprise the smallest coherent algebraic structure encompassing the various integrands of interest to us. The fact that differential forms are algebraic makes them easy to study without thinking directly about integration. The algebra leads to a general version of the Fundamental Theorem of Integral Calculus that is rich in geometry. The theorem subsumes the three classical vector integration theorems, Green's Theorem, Stokes's Theorem, and Gauss's Theorem, also called the Divergence Theorem.

Comments and corrections should be sent to jerry@reed.edu.

Exercises

0.0.1. (a) Consider two surfaces in space, each surface having at each of its points a tangent plane and therefore a normal line, and consider pairs of points, one on each surface. Conjecture a geometric condition, phrased in terms of tangent planes and/or normal lines, about the closest pair of points.

(b) Consider a surface in space and a curve in space, the curve having at each of its points a tangent line and therefore a normal plane, and consider pairs of points, one on the surface and one on the curve. Make a conjecture about the closest pair of points.

(c) Make a conjecture about the closest pair of points on two curves.

0.0.2. (a) Assume that the factorial of a half-integer makes sense, and grant the general formula for the volume of a ball in n dimensions. Explain why it follows that $(1/2)! = \sqrt{\pi}/2$. Further assume that the half-integral factorial function satisfies the relation

$$x! = x \cdot (x - 1)! \quad \text{for } x = 3/2, 5/2, 7/2, \dots$$

Subject to these assumptions, verify that the volume of the ball of radius r in three dimensions is $\frac{4}{3}\pi r^3$ as claimed. What is the volume of the ball of radius r in five dimensions?

(b) The ball of radius r in n dimensions sits inside a circumscribing box of sides $2r$. Draw pictures of this configuration for $n = 1, 2, 3$. Determine what portion of the box is filled by the ball in the limit as the dimension n gets large. That is, find

$$\lim_{n \rightarrow \infty} \frac{\text{vol}(B_n(r))}{(2r)^n}.$$

Results from One-Variable Calculus

As a warmup, these notes begin with a quick review of some ideas from one-variable calculus. The material in the first two sections is assumed to be familiar. Section 3 discusses Taylor's Theorem at greater length, not assuming that the reader has already seen it.

1.1 The Real Number System

We assume that there is a **real number system**, a set \mathbf{R} that contains two distinct elements 0 and 1 and is endowed with the algebraic operations of addition,

$$+ : \mathbf{R} \times \mathbf{R} \longrightarrow \mathbf{R},$$

and multiplication,

$$\cdot : \mathbf{R} \times \mathbf{R} \longrightarrow \mathbf{R}.$$

The sum $+(a, b)$ is written $a + b$, and the product $\cdot(a, b)$ is written $a \cdot b$ or more briefly as ab .

Theorem 1.1.1 (Field Axioms for $(\mathbf{R}, +, \cdot)$). *The real number system, with its distinct 0 and 1 and with its addition and multiplication, is assumed to satisfy the following set of axioms.*

- (a1) *Addition is associative: $(x + y) + z = x + (y + z)$ for all $x, y, z \in \mathbf{R}$.*
- (a2) *0 is an additive identity: $x + 0 = x$ for all $x \in \mathbf{R}$.*
- (a3) *Existence of additive inverses: For each $x \in \mathbf{R}$ there exists $y \in \mathbf{R}$ such that $x + y = 0$.*
- (a4) *Addition is commutative: $x + y = y + x$ for all $x, y \in \mathbf{R}$.*
- (m1) *Multiplication is associative: $x(yz) = (xy)z$ for all $x, y, z \in \mathbf{R}$.*
- (m2) *1 is a multiplicative identity: $1x = x$ for all $x \in \mathbf{R}$.*
- (m3) *Existence of multiplicative inverses: For each nonzero $x \in \mathbf{R}$ there exists $y \in \mathbf{R}$ such that $xy = 1$.*
- (m4) *Multiplication is commutative: $xy = yx$ for all $x, y \in \mathbf{R}$.*

(d1) *Multiplication distributes over addition: $(x+y)z = xz+yz$ for all $x, y, z \in \mathbf{R}$.*

All of basic algebra follows from the field axioms. Additive and multiplicative inverses are unique, the cancellation law holds, $0 \cdot x = 0$ for all real numbers x , and so on.

Subtracting a real number from another is defined as adding the additive inverse. In symbols,

$$- : \mathbf{R} \times \mathbf{R} \longrightarrow \mathbf{R}, \quad x - y = x + (-y) \quad \text{for all } x, y \in \mathbf{R}.$$

We also assume that \mathbf{R} is an **ordered** field. This means that there is a subset \mathbf{R}^+ of \mathbf{R} (the **positive** elements) such that the following axioms hold.

Theorem 1.1.2 (Order Axioms).

(o1) *Trichotomy Axiom: For every real number x , exactly one of the following conditions holds:*

$$x \in \mathbf{R}^+, \quad -x \in \mathbf{R}^+, \quad x = 0.$$

(o2) *Closure of positive numbers under addition: For all real numbers x and y , if $x \in \mathbf{R}^+$ and $y \in \mathbf{R}^+$ then also $x + y \in \mathbf{R}^+$.*

(o3) *Closure of positive numbers under multiplication: For all real numbers x and y , if $x \in \mathbf{R}^+$ and $y \in \mathbf{R}^+$ then also $xy \in \mathbf{R}^+$.*

For all real numbers x and y , define “ $x < y$ ” to mean “ $y - x \in \mathbf{R}^+$.” The usual rules for inequalities then follow from the axioms.

Finally, we assume that the real number system is **complete**. Completeness can be phrased in various ways, all logically equivalent. The version of completeness that is currently in Ray Mayer’s notes for Mathematics 112 is as follows.

Theorem 1.1.3 (Completeness as a Binary Search Criterion). *Every binary search sequence in the real number system converges to a unique limit.*

Two other versions of completeness are phrased in terms of sequences and in terms of set-bounds:

Theorem 1.1.4 (Completeness as a Monotonic Sequence Criterion). *Every bounded monotonic sequence in \mathbf{R} converges to a unique limit.*

Theorem 1.1.5 (Completeness as a Set-Bound Criterion). *Every non-empty subset of \mathbf{R} that is bounded above has a least upper bound.*

Convergence is a concept of analysis, and therefore so is completeness. All three statements of completeness are existence statements.

Exercises

1.1.1. For any positive integer n , let $\mathbf{Z}/n\mathbf{Z}$ denote the set $\{0, 1, \dots, n-1\}$ with the usual operations of addition and multiplication carried out taking remainders. That is, add and multiply in the usual fashion but subject to the additional condition that $n = 0$. For example, in $\mathbf{Z}/5\mathbf{Z}$ we have $2 + 4 = 1$ and $2 \cdot 4 = 3$. For what values of n does $\mathbf{Z}/n\mathbf{Z}$ form a field?

1.1.2. Prove that in any ordered field, $0 < 1$. Prove that the complex number field \mathbf{C} can not be made an ordered field.

1.1.3. Use a completeness property of the real number system to show that 2 has a positive square root.

1.1.4. (a) Prove by induction that

$$\sum_{i=1}^n i^2 = \frac{n(n+1)(2n+1)}{6} \quad \text{for all } n \in \mathbf{Z}^+.$$

(b) (**Bernoulli's Inequality**) For any real number $r \geq -1$, prove that

$$(1+r)^n \geq 1+rn \quad \text{for all } n \in \mathbf{N}.$$

(c) For what positive integers n is $2^n > n^3$?

1.2 Foundational and Basic Theorems

This section reviews the foundational theorems of one-variable calculus. The first two theorems are not theorems of calculus at all, but rather they are theorems about continuous functions and the real number system. The first theorem says that under suitable conditions, an optimization problem is guaranteed to have a solution.

Theorem 1.2.1 (Extreme Value Theorem). *Let I be a nonempty closed and bounded interval in \mathbf{R} , and let $f : I \rightarrow \mathbf{R}$ be a continuous function. Then f takes a minimum value and a maximum value on I .*

The second theorem says that under suitable conditions, any value trapped between two output values of a function must itself be an output value.

Theorem 1.2.2 (Intermediate Value Theorem). *Let I be a nonempty interval in \mathbf{R} , and let $f : I \rightarrow \mathbf{R}$ be a continuous function. Let y be a real number, and suppose that*

$$f(x) < y \quad \text{for some } x \in I$$

and

$$f(x') > y \quad \text{for some } x' \in I.$$

Then

$$f(c) = y \quad \text{for some } c \in I.$$

The Mean Value Theorem relates the derivative of a function to values of the function itself with no reference to the fact that the derivative is a limit, but at the cost of introducing an unknown point.

Theorem 1.2.3 (Mean Value Theorem). *Let a and b be real numbers with $a < b$. Suppose that the function $f : [a, b] \rightarrow \mathbf{R}$ is continuous and that f is differentiable on the open subinterval (a, b) . Then*

$$\frac{f(b) - f(a)}{b - a} = f'(c) \quad \text{for some } c \in (a, b).$$

The Fundamental Theorem of Integral Calculus relates the integral of the derivative to the original function, assuming that the derivative is continuous.

Theorem 1.2.4 (Fundamental Theorem of Integral Calculus). *Let I be a nonempty interval in \mathbf{R} , and let $f : I \rightarrow \mathbf{R}$ be a continuous function. Suppose that the function $F : I \rightarrow \mathbf{R}$ has derivative f . Then for any closed and bounded subinterval $[a, b]$ of I ,*

$$\int_a^b f(x) dx = F(b) - F(a).$$

Exercises

1.2.1. Use the Intermediate Value Theorem to show that 2 has a positive square root.

1.2.2. Let $f : [0, 1] \rightarrow [0, 1]$ be continuous. Use the Intermediate Value Theorem to show that $f(x) = x$ for some $x \in [0, 1]$.

1.2.3. Let a and b be real numbers with $a < b$. Suppose that $f : [a, b] \rightarrow \mathbf{R}$ is continuous and that f is differentiable on the open subinterval (a, b) . Use the Mean Value Theorem to show that if $f' > 0$ on (a, b) then f is strictly increasing on $[a, b]$.

1.2.4. For the Extreme Value Theorem, the Intermediate Value Theorem, and the Mean Value Theorem, give examples to show that weakening the hypotheses of the theorem gives rise to examples where the conclusion of the theorem fails.

1.3 Taylor's Theorem

Let $I \subset \mathbf{R}$ be a nonempty open interval, and let $a \in I$ be any point. Let n be a nonnegative integer. Suppose that the function $f : I \rightarrow \mathbf{R}$ has n continuous derivatives,

$$f, f', f'', \dots, f^{(n)} : I \rightarrow \mathbf{R}.$$

Suppose further that we know the values of f and its derivatives at a , the $n + 1$ numbers

$$f(a), \quad f'(a), \quad f''(a), \quad \dots, \quad f^{(n)}(a).$$

(For instance, if $f : \mathbf{R} \rightarrow \mathbf{R}$ is the sine function, and $a = 0$, and n is even, then the numbers are $1, 0, -1, 0, \dots, (-1)^{n/2}$.)

Question 1 (Existence and Uniqueness): Is there a polynomial p of degree n that mimics the behavior of f at a in the sense that

$$p(a) = f(a), \quad p'(a) = f'(a), \quad p''(a) = f''(a), \quad \dots, \quad p^{(n)}(a) = f^{(n)}(a)?$$

Is there only one such polynomial?

Question 2 (Accuracy of Approximation): How well does $p(x)$ approximate $f(x)$ for $x \neq a$?

Question 1 is easy to answer. Consider a polynomial of degree n expanded about $x = a$,

$$p(x) = a_0 + a_1(x - a) + a_2(x - a)^2 + a_3(x - a)^3 + \dots + a_n(x - a)^n.$$

The goal is to choose the coefficients a_0, \dots, a_n to make p behave like the original function f at a . Note that $p(a) = a_0$. We want $p(a)$ to equal $f(a)$, so set

$$a_0 = f(a).$$

Differentiate p to obtain

$$p'(x) = a_1 + 2a_2(x - a) + 3a_3(x - a)^2 + \dots + na_n(x - a)^{n-1},$$

so that $p'(a) = a_1$. We want $p'(a)$ to equal $f'(a)$, so set

$$a_1 = f'(a).$$

Differentiate again to obtain

$$p''(x) = 2a_2 + 3 \cdot 2a_3(x - a) + \dots + n(n - 1)a_n(x - a)^{n-2},$$

so that $p''(a) = 2a_2$. We want $p''(a)$ to equal $f''(a)$, so set

$$a_2 = \frac{f''(a)}{2}.$$

Differentiate again to obtain

$$p'''(x) = 3 \cdot 2a_3 + \dots + n(n - 1)(n - 2)a_n(x - a)^{n-3},$$

so that $p'''(a) = 3 \cdot 2a_3$. We want $p'''(a)$ to equal $f'''(a)$, so set

$$a_3 = \frac{f''(a)}{3 \cdot 2}.$$

Continue in this fashion to obtain $a_4 = f^{(4)}(a)/4!$ and so on up to $a_n = f^{(n)}(a)/n!$. That is, the desired coefficients are

$$a_k = \frac{f^{(k)}(a)}{k!} \quad \text{for } k = 0, \dots, n.$$

This answers the existence part of Question 2: yes. Furthermore, since the calculation offered us no choices en route, these are the only coefficients that can work, and so the approximating polynomial is unique.

Definition 1.3.1 (*n*th degree Taylor Polynomial). Let $I \subset \mathbf{R}$ be a nonempty open interval, and let a be a point of I . Let n be a nonnegative integer. Suppose that the function $f : I \rightarrow \mathbf{R}$ has n continuous derivatives. Then the *n*th degree Taylor polynomial of f at a is

$$T_n(x) = f(a) + f'(a)(x-a) + \frac{f''(a)}{2}(x-a)^2 + \cdots + \frac{f^{(n)}(a)}{n!}(x-a)^n.$$

In more concise notation,

$$T_n(x) = \sum_{k=0}^n \frac{f^{(k)}(a)}{k!} (x-a)^k.$$

For example, if $f(x) = e^x$ and $a = 0$ then it is easy to generate the following table:

k	$f^{(k)}(x)$	$\frac{f^{(k)}(0)}{k!}$
0	e^x	1
1	e^x	1
2	e^x	$\frac{1}{2}$
3	e^x	$\frac{1}{3!}$
\vdots	\vdots	\vdots
n	e^x	$\frac{1}{n!}$

From the table we can read off the *n*th degree Taylor polynomial of f at 0,

$$\begin{aligned} T_n(x) &= 1 + x + \frac{x^2}{2} + \frac{x^3}{3!} + \cdots + \frac{x^n}{n!} \\ &= \sum_{k=0}^n \frac{x^k}{k!}. \end{aligned}$$

Recall that the second question is how well the polynomial $T_n(x)$ approximates $f(x)$ for $x \neq a$. Thus it is a question about the difference $f(x) - T_n(x)$. Giving this quantity its own name is useful.

Definition 1.3.2 (nth degree Taylor Remainder). *Let $I \subset \mathbf{R}$ be a nonempty open interval, and let a be a point of I . Let n be a nonnegative integer. Suppose that the function $f : I \rightarrow \mathbf{R}$ has n continuous derivatives. Then the n th degree Taylor remainder of f at a is*

$$R_n(x) = f(x) - T_n(x).$$

So the second question is to estimate the remainder $R_n(x)$ for points $x \in I$. The method about to be presented here for doing so proceeds very naturally, but it is perhaps a little surprising because although the Taylor polynomial $T_n(x)$ is expressed in terms of derivatives, as is the expression to be obtained for the remainder $R_n(x)$, we obtain the expression by using the Fundamental Theorem of Integral Calculus repeatedly.

The method requires a calculation, and so, guided by hindsight, we first carry it out so that then the ideas of the method itself will be uncluttered. For any positive integer k and any $x \in \mathbf{R}$, define a k -fold nested integral,

$$I_k(x) = \int_{x_1=a}^x \int_{x_2=a}^{x_1} \cdots \int_{x_k=a}^{x_{k-1}} dx_k \cdots dx_2 dx_1.$$

This expression may appear daunting, but it unwinds readily if we work a step at a time from the inside out, managing the subscripts accurately. First,

$$\int_a^{x_{k-1}} dx_k = x_k \Big|_{x_k=a}^{x_{k-1}} = x_{k-1} - a.$$

Move one layer out and use this result to get

$$\begin{aligned} \int_a^{x_{k-2}} \int_a^{x_{k-1}} dx_k dx_{k-1} &= \int_a^{x_{k-2}} (x_{k-1} - a) dx_{k-1} \\ &= \frac{1}{2} (x_{k-1} - a)^2 \Big|_{x_{k-1}=a}^{x_{k-2}} \\ &= \frac{1}{2} (x_{k-2} - a)^2. \end{aligned}$$

Again move out and quote the previous calculation,

$$\begin{aligned} \int_a^{x_{k-3}} \int_a^{x_{k-2}} \int_a^{x_{k-1}} dx_k dx_{k-1} dx_{k-2} &= \int_a^{x_{k-3}} \frac{1}{2} (x_{k-2} - a)^2 dx_{k-2} \\ &= \frac{1}{3!} (x_{k-2} - a)^3 \Big|_{x_{k-2}=a}^{x_{k-3}} \\ &= \frac{1}{3!} (x_{k-3} - a)^3. \end{aligned}$$

The pattern is clear, and the process is complete after k integrations,

$$I_k(x) = \frac{1}{k!}(x-a)^k, \quad k \in \mathbf{Z}^+.$$

Note that this is part of the k th term $\frac{f^{(k)}(a)}{k!}(x-a)^k$ of the Taylor polynomial, the part that makes no reference to the function f . That is, $f^{(k)}(a)I_k(x)$ is the k th term of the Taylor polynomial for $k = 1, 2, 3, \dots$

With the formula for $I_k(x)$ in hand, we return to using the Fundamental Theorem of Integral Calculus to study the remainder $R_n(x)$, the function $f(x)$ minus its n th degree Taylor polynomial $T_n(x)$. According to the Fundamental Theorem,

$$f(x) = f(a) + \int_a^x f'(x_1) dx_1,$$

That is, $f(x)$ is equal to the constant term of the Taylor polynomial plus an integral,

$$f(x) = T_0(x) + \int_a^x f'(x_1) dx_1.$$

By the Fundamental Theorem again, the integral is in turn

$$\int_a^x f'(x_1) dx_1 = \int_a^x \left(f'(a) + \int_a^{x_1} f''(x_2) dx_2 \right) dx_1.$$

The first term of the outer integral is $f'(a)I_1(x)$, giving the first-order term of the Taylor polynomial and leaving a doubly-nested integral,

$$\int_a^x f'(x_1) dx_1 = f'(a)(x-a) + \int_a^x \int_a^{x_1} f''(x_2) dx_2 dx_1.$$

In other words, the calculation so far has shown that

$$\begin{aligned} f(x) &= f(a) + f'(a)(x-a) + \int_a^x \int_a^{x_1} f''(x_2) dx_2 dx_1 \\ &= T_1(x) + \int_a^x \int_a^{x_1} f''(x_2) dx_2 dx_1. \end{aligned}$$

Once more by the Fundamental Theorem the doubly-nested integral is

$$\int_a^x \int_a^{x_1} f''(x_2) dx_2 dx_1 = \int_a^x \int_a^{x_1} \left(f''(a) + \int_a^{x_2} f'''(x_3) dx_3 \right) dx_2 dx_1,$$

and the first term of the outer integral is $f''(a)I_2(x)$, giving the second-order term of the Taylor polynomial and leaving a triply-nested integral,

$$\int_a^x \int_a^{x_1} f''(x_2) dx_2 dx_1 = \frac{f''(a)}{2}(x-a)^2 + \int_a^x \int_a^{x_1} \int_a^{x_2} f'''(x_3) dx_3 dx_2 dx_1.$$

So now the calculation so far has shown that

$$f(x) = T_2(x) + \int_a^x \int_a^{x_1} \int_a^{x_2} f'''(x_3) dx_3 dx_2 dx_1.$$

Continuing this process through n iterations shows that $f(x)$ is $T_n(x)$ plus an $(n + 1)$ -fold iterated integral,

$$f(x) = T_n(x) + \int_a^x \int_a^{x_1} \cdots \int_a^{x_n} f^{(n+1)}(x_{n+1}) dx_{n+1} \cdots dx_2 dx_1.$$

In other words, the remainder is the integral,

$$R_n(x) = \int_a^x \int_a^{x_1} \cdots \int_a^{x_n} f^{(n+1)}(x_{n+1}) dx_{n+1} \cdots dx_2 dx_1. \quad (1.1)$$

Note that we now are assuming that f has $n + 1$ continuous derivatives.

For simplicity, assume that $x > a$. Since $f^{(n+1)}$ is continuous on the closed and bounded interval $[a, x]$, the Extreme Value Theorem says that it takes a minimum value m and a maximum value M on the interval. That is,

$$m \leq f^{(n+1)}(x_{n+1}) \leq M, \quad x_{n+1} \in [a, x].$$

Integrating this inequality $n + 1$ times correspondingly bounds the remainder integral (1.1) on both sides by multiples of the integral that we have evaluated,

$$mI_{n+1}(x) \leq R_n(x) \leq MI_{n+1}(x),$$

and therefore by the precalculated formula for $I_{n+1}(x)$,

$$m \frac{(x-a)^{n+1}}{(n+1)!} \leq R_n(x) \leq M \frac{(x-a)^{n+1}}{(n+1)!}. \quad (1.2)$$

Recall that m and M are particular values of $f^{(n+1)}$. Define an auxiliary function that will therefore assume the sandwiching values in (1.2),

$$g: [a, x] \longrightarrow \mathbf{R}, \quad g(t) = f^{(n+1)}(t) \frac{(x-a)^{n+1}}{(n+1)!}.$$

That is, since there exist values t_m and t_M in $[a, x]$ such that $f^{(n+1)}(t_m) = m$ and $f^{(n+1)}(t_M) = M$, the result (1.2) of our calculation rephrases as

$$g(t_m) \leq R_n(x) \leq g(t_M).$$

This shows that the remainder is an intermediate value of g . And g is continuous, so by the Intermediate Value Theorem, there exists some point $c \in [a, x]$ such that $g(c) = R_n(x)$. In other words, $g(c)$ is the desired remainder, the function minus its Taylor polynomial. This proves

Theorem 1.3.3 (Taylor's Theorem). *Let $I \subset \mathbf{R}$ be a nonempty open interval, and let $a \in I$. Let n be a nonnegative integer. Suppose that the function $f : I \rightarrow \mathbf{R}$ has $n + 1$ continuous derivatives. Then for each $x \in I$,*

$$f(x) = T_n(x) + R_n(x)$$

where

$$R_n(x) = f^{(n+1)}(c) \frac{(x-a)^{n+1}}{(n+1)!} \quad \text{for some } c \text{ between } a \text{ and } x.$$

We have proved Taylor's Theorem only when $x > a$ (and it is trivial when $x = a$), but the proof for $x < a$ is the same other than a few inequalities switching direction. Whereas our proof of Taylor's Theorem relies primarily on the Fundamental Theorem of Integral Calculus, and a similar proof relies on repeated integration by parts (exercise 1.3.6), many proofs rely instead on the Mean Value Theorem. Our proof neatly uses different mathematical tools for the different parts of the argument:

- To find the Taylor polynomial $T_n(x)$ we differentiated repeatedly, using a substitution at each step to determine a coefficient.
- To get a precise (if unwieldy) expression for the remainder $R_n(x) = f(x) - T_n(x)$ we integrated repeatedly, using the Fundamental Theorem of Integral Calculus at each step to produce a term of the Taylor polynomial.
- To express the remainder in a more convenient form, we used the Extreme Value Theorem and then the Intermediate Value Theorem once each. This step involved no calculus.

The expression for $R_n(x)$ given in Theorem 1.3.3 is called the **Lagrange form** of the remainder. Other expressions for $R_n(x)$ exist as well. Whatever form is used for the remainder, it should be something that we can estimate by bounding its magnitude.

For example, we use Taylor's Theorem to estimate $\ln(1.1)$ by hand to within $1/500,000$. Let $f(x) = \ln(1+x)$ on $(-1, \infty)$, and let $a = 0$. Compute the following table:

k	$f^{(k)}(x)$	$\frac{f^{(k)}(0)}{k!}$
0	$\ln(1+x)$	0
1	$\frac{1}{(1+x)}$	1
2	$-\frac{1}{(1+x)^2}$	$-\frac{1}{2}$
3	$\frac{2}{(1+x)^3}$	$\frac{1}{3}$
4	$-\frac{3!}{(1+x)^4}$	$-\frac{1}{4}$
\vdots	\vdots	\vdots
n	$\frac{(-1)^{n-1}(n-1)!}{(1+x)^n}$	$\frac{(-1)^{n-1}}{n}$
$n+1$	$\frac{(-1)^n n!}{(1+x)^{n+1}}$	

Next, read off from the table that for $n \geq 1$, the n th degree Taylor polynomial is

$$T_n(x) = x - \frac{x^2}{2} + \frac{x^3}{3} - \dots + (-1)^{n-1} \frac{x^n}{n} = \sum_{k=1}^n (-1)^{k-1} \frac{x^k}{k},$$

and the remainder is

$$R_n(x) = \frac{(-1)^n x^{n+1}}{(1+c)^{n+1}(n+1)} \quad \text{for some } c \text{ between } 0 \text{ and } x.$$

This expression for the remainder may be a bit much to take in since it involves three variables: the point x at which we are approximating the logarithm, the degree n of the Taylor polynomial that is providing the approximation, and the unknown value c in the error term. But in particular we are interested in $x = 0.1$ (since we are approximating $\ln(1.1)$ using $f(x) = \ln(1+x)$), so that the Taylor polynomial specializes to

$$T_n(0.1) = (0.1) - \frac{(0.1)^2}{2} + \frac{(0.1)^3}{3} - \dots + (-1)^{n-1} \frac{(0.1)^n}{n},$$

and we want to bound the remainder in absolute value, so write

$$|R_n(0.1)| = \frac{(0.1)^{n+1}}{(1+c)^{n+1}(n+1)} \quad \text{for some } c \text{ between } 0 \text{ and } 0.1.$$

Now the symbol x is gone. Next, note that although we don't know the value of c , the *smallest* possible value of the quantity $(1+c)^{n+1}$ in the denominator of the absolute remainder is 1 because $c \geq 0$. And since this value occurs in

the denominator it lets us write the *greatest* possible value of the absolute remainder with no reference to c . That is,

$$|R_n(0.1)| \leq \frac{(0.1)^{n+1}}{(n+1)},$$

and the symbol c is gone as well. The only remaining variable is n , and the goal is to approximate $\ln(1.1)$ to within $1/500,000$. Set $n = 4$ in the previous display to get

$$|R_4(0.1)| \leq \frac{1}{500,000}.$$

That is, the fourth degree Taylor polynomial

$$\begin{aligned} T_4(0.1) &= \frac{1}{10} - \frac{1}{200} + \frac{1}{3000} - \frac{1}{40000} \\ &= 0.10000000 \dots - 0.00500000 \dots + 0.00033333 \dots - 0.00002500 \dots \\ &= 0.09530833 \dots \end{aligned}$$

agrees with $\ln(1.1)$ to within $0.00000200 \dots$, so that

$$0.09530633 \dots \leq \ln(1.1) \leq 0.09531033 \dots$$

Machine technology should confirm this.

Continuing to work with the function $f(x) = \ln(1+x)$ for $x > -1$, set $x = 1$ instead to get that for $n \geq 1$,

$$T_n(1) = 1 - \frac{1}{2} + \frac{1}{3} - \dots + (-1)^{n-1} \frac{1}{n},$$

and

$$|R_n(1)| = \left| \frac{1}{(1+c)^{n+1}(n+1)} \right| \quad \text{for some } c \text{ between } 0 \text{ and } 1.$$

Thus $|R_n(1)| \leq 1/(n+1)$, and this goes to 0 as $n \rightarrow \infty$. Therefore $\ln(2)$ is expressible as an infinite series,

$$\ln(2) = 1 - \frac{1}{2} + \frac{1}{3} - \frac{1}{4} + \dots$$

Repeating a formula from before, the n th degree Taylor polynomial of the natural logarithm at 1 is

$$T_n(x) = x - \frac{x^2}{2} + \frac{x^3}{3} - \dots + (-1)^{n-1} \frac{x^n}{n} = \sum_{k=1}^n (-1)^{k-1} \frac{x^k}{k},$$

The graphs of the natural logarithm and its first five Taylor polynomials at 1 are plotted from 0 to 2 in figure 1.1. A good check of your understanding is to see if you can determine which graph is which in the figure.




Figure 1.1. The natural logarithm and its Taylor polynomials

For another example, return to the exponential function $f(x) = e^x$ and let $a = 0$. For any x , the difference between $f(x)$ and the n th degree Taylor polynomial $T_n(x)$ satisfies

$$|R_n(x)| = \left| e^c \frac{x^{n+1}}{(n+1)!} \right| \quad \text{for some } c \text{ between } 0 \text{ and } x.$$

If $x \geq 0$ then e^c could be as large as e^x , while if $x < 0$ then e^c could be as large as e^0 . The worst possible case is therefore

$$|R_n(x)| \leq \max\{1, e^x\} \frac{|x|^{n+1}}{(n+1)!}.$$

As $n \rightarrow \infty$ (while x remains fixed, albeit arbitrary) the right side goes to 0 because the factorial growth of $(n+1)!$ dominates the polynomial growth of $|x|^{n+1}$, and so we have in the limit that e^x is expressible as a power series,

$$e^x = 1 + x + \frac{x^2}{2!} + \frac{x^3}{3!} + \cdots + \frac{x^n}{n!} + \cdots = \sum_{k=0}^{\infty} \frac{x^k}{k!}.$$

In Ray Mayer's notes for Mathematics 112, the power series here is used to define e^x , but then obtaining the properties of e^x depends on the technical fact that power series can be differentiated term by term.

We end this chapter by sketching two cautionary examples. First, work from earlier in the section shows that the Taylor series for the function $\ln(1+x)$ at $a = 0$ is

$$T(x) = x - \frac{x^2}{2} + \frac{x^3}{3} - \cdots + (-1)^{n-1} \frac{x^n}{n} + \cdots = \sum_{k=1}^{\infty} (-1)^{k-1} \frac{x^k}{k}.$$

The Ratio Test shows that this series converges absolutely when $|x| < 1$. The series also converges at $x = 1$, as observed earlier. Thus, while the domain of the function $\ln(1+x)$ is $(-1, \infty)$, the Taylor series has no chance to match the function outside of $(-1, 1]$. Furthermore, it is possible, using the Lagrange form $R_n(x)$ of the n th remainder, to argue similarly to the previous paragraph that

$$\ln(1+x) = T(x) \quad \text{for } x \in [-1/2, 1].$$

On the other hand, the Lagrange form can not easily be used to prove that the equality in the previous display also holds for $x \in (-1, -1/2)$. Figure 1.1 suggests why: the Taylor polynomials are converging more slowly to the original function the farther left we go on the graph. However, a different form of the remainder, given in exercise 1.3.6, proves fairly easily that indeed the equality holds for all $x \in (-1, 1]$.

For the last example, define $f : \mathbf{R} \rightarrow \mathbf{R}$ by

$$f(x) = \begin{cases} e^{-1/x^2} & \text{if } x \neq 0, \\ 0 & \text{if } x = 0. \end{cases}$$

It is possible to show that f is infinitely differentiable and that every derivative of f at 0 is 0. That is, $f^{(k)}(0) = 0$ for $k = 0, 1, 2, \dots$. Consequently, the Taylor series for f at 0 is

$$T(x) = 0 + 0x + 0x^2 + \cdots + 0x^n + \cdots.$$

This is the zero function, which certainly converges for all $x \in \mathbf{R}$. But the only value of x for which it converges to the original function f is $x = 0$. In other words, although this Taylor series converges everywhere, it fails catastrophically to equal the function it is attempting to match. The problem is that the function f decays exponentially, and since exponential behavior dominates polynomial behavior, any attempt to discern f by using polynomials will fail to see it. Figures 1.2 and 1.3 plot f to display its rapid decay. The first plot is for $x \in [-25, 25]$ and the second is for $x \in [-1/2, 1/2]$.

Exercises

1.3.1. What is the n th degree Taylor polynomial $T_n(x)$ for the function $f(x) = \sin x$ at 0? Prove that $\sin x$ is equal to the limit of $T_n(x)$ as $n \rightarrow \infty$, similarly to the argument in the text for e^x . Also find $T_n(x)$ for $f(x) = \cos x$ at 0, and explain why the argument for \sin shows that $\cos x$ is the limit of its Taylor polynomials as well.




Figure 1.2. Rapidly decaying function, wide view




Figure 1.3. Rapidly decaying function, zoom view

1.3.2. What is the n th degree Taylor polynomial $T_n(x)$ for the following functions at 0?

- (a) $f(x) = \arctan x$,
- (b) $f(x) = (1 + x)^\alpha$ where $\alpha \in \mathbf{R}$. (Warning: Although the answer can be written in a uniform way for all α , it behaves differently when $\alpha \in \mathbf{N}$.)

1.3.3. In figure 1.1, identify the graphs of T_1 through T_5 and the graph of \ln near $x = 0$ and near $x = 2$.

1.3.4. Without a calculator, use the first three terms of the Taylor series for $\sin(x)$ at 0 to approximate a decimal representation of $\sin(0.1)$. Also compute the decimal representation of an upper bound for the error of the approximation. Bound $\sin(0.1)$ between two decimal representations.

1.3.5. Use a second degree Taylor polynomial to approximate $\sqrt{4.2}$. Use Taylor's theorem to find a guaranteed accuracy of the approximation and thus to find upper and lower bounds for $\sqrt{4.2}$.

1.3.6. Another proof of Taylor's Theorem uses the Fundamental Theorem of Integral Calculus once and then integrates by parts repeatedly. Begin with the hypotheses of Theorem 1.3.3, and let $x \in I$. By the Fundamental Theorem,

$$f(x) = f(a) + \int_a^x f'(t) dt.$$

Let $u = f'(t)$ and $v = t - x$, so that the integral is $\int_a^x u dv$, and integrating by parts gives

$$f(x) = f(a) + f'(a)(x - a) - \int_a^x f''(t)(t - x) dt.$$

Let $u = f''(t)$ and $v = \frac{1}{2}(t - x)^2$, so that again the integral is $\int_a^x u dv$, and integrating by parts gives

$$f(x) = f(a) + f'(a)(x - a) + f''(a)\frac{(x - a)^2}{2} + \int_a^x f'''(t)\frac{(t - x)^2}{2} dt.$$

Show that after n steps the result is

$$f(x) = T_n(x) + (-1)^n \int_a^x f^{(n+1)}(t) \frac{(t - x)^n}{n!} dt.$$

Whereas the expression for $f(x) - T_n(x)$ in Theorem 1.3.3 is called the **Lagrange form** of the remainder, this exercise has derived the **integral form** of the remainder. Use the Extreme Value Theorem and the Intermediate Value Theorem to derive the Lagrange form of the remainder from the integral form.

Euclidean Space

Euclidean space is a mathematical construct that encompasses the line, the plane, and three-dimensional space as special cases. Its elements are called *vectors*. Vectors can be understood in various ways: as arrows, as quantities with magnitude and direction, as displacements, or as points. However, along with a sense of what vectors are, we also need to emphasize how they interact. The axioms in section 2.1 capture the idea that vectors can be added together and can be multiplied by scalars, with both of these operations obeying familiar laws of algebra. Section 2.2 expresses the geometric ideas of length and angle in Euclidean space in terms of vector algebra. Section 2.3 discusses continuity for functions (also called *mappings*) whose inputs and outputs are vectors rather than scalars. Section 2.4 introduces a special class of sets in Euclidean space, the *compact* sets, and shows that compact sets are preserved under continuous mappings. Finally, section 2.5 reviews the one-variable derivative in light of ideas from the two sections preceding it.

2.1 Algebra: Vectors

Let n be a positive integer. The set of all ordered n -tuples of real numbers,

$$\mathbf{R}^n = \{(x_1, \dots, x_n) : x_i \in \mathbf{R} \text{ for } i = 1, \dots, n\},$$

constitutes **n -dimensional Euclidean space**. When $n = 1$, the parentheses and subscript in the notation (x_1) are superfluous, so we simply view the elements of \mathbf{R}^1 as real numbers x and write \mathbf{R} for \mathbf{R}^1 . Elements of \mathbf{R}^2 and of \mathbf{R}^3 are written (x, y) and (x, y, z) to avoid needless subscripts. These first few Euclidean spaces, \mathbf{R} , \mathbf{R}^2 and \mathbf{R}^3 , are conveniently visualized as the line, the plane, and space itself. (See figure 2.1.)

Elements of \mathbf{R}^n are called **scalars**, of \mathbf{R}^n , **vectors**. The **origin** of \mathbf{R}^n , denoted $\mathbf{0}$, is defined to be

$$\mathbf{0} = (0, \dots, 0).$$




Figure 2.1. The first few Euclidean spaces

Sometimes the origin of \mathbf{R}^n will be denoted $\mathbf{0}_n$ to distinguish it from other “origins” that we will encounter later.

In the first few Euclidean spaces \mathbf{R} , \mathbf{R}^2 , \mathbf{R}^3 , one can visualize a vector as a point x or as an arrow. The arrow can have its tail at the origin and its head at the point x , or its tail at any point p and its head correspondingly translated to $p + x$. (See figure 2.2. Most illustrations will depict \mathbf{R} or \mathbf{R}^2 .)



Figure 2.2. Various ways to envision a vector

To a mathematician, the word “space” doesn’t connote volume but instead refers to a set endowed with some structure. Indeed, Euclidean space \mathbf{R}^n comes with two algebraic operations. The first is **vector addition**,

$$+ : \mathbf{R}^n \times \mathbf{R}^n \longrightarrow \mathbf{R}^n,$$

defined by adding the scalars at each component of the vectors,

$$(x_1, \dots, x_n) + (y_1, \dots, y_n) = (x_1 + y_1, \dots, x_n + y_n).$$

For example, $(1, 2, 3) + (4, 5, 6) = (5, 7, 9)$. Note that meaning of the “+” sign is now overloaded: on the left of the displayed equality, it denotes the new operation of vector addition, whereas on the right side it denotes the old addition of real numbers. This shouldn’t cause problems since which “+” is meant is clear from context, i.e., the meaning of “+” is clear from whether

it sits between vectors or scalars. (An expression such as “ $(1, 2, 3) + 4$,” with the plus sign between a vector and a scalar, makes no sense according to our grammar.)

The interpretation of vectors as arrows gives a geometric description of vector addition, at least in \mathbf{R}^2 . To add the vectors x and y , draw them as arrows starting at $\mathbf{0}$ and then complete the parallelogram P that has x and y as two of its sides. The diagonal of P starting at $\mathbf{0}$ is then the arrow depicting the vector $x + y$. (See figure 2.3.) The proof of this is a small argument with similar triangles, left to the reader as exercise 2.1.2.




Figure 2.3. The parallelogram law of vector addition

The second operation on Euclidean space is **scalar multiplication**,

$$\cdot : \mathbf{R} \times \mathbf{R}^n \longrightarrow \mathbf{R}^n,$$

defined by

$$a \cdot (x_1, \dots, x_n) = (ax_1, \dots, ax_n).$$

For example, $2 \cdot (3, 4, 5) = (6, 8, 10)$. We will almost always omit the symbol “ \cdot ” and write ax for $a \cdot x$. With this convention, juxtaposition is overloaded as “ $+$ ” was overloaded above, but again this shouldn’t cause problems.

Scalar multiplication of the vector x (viewed as an arrow) by a also has a geometric interpretation: it simply stretches (i.e., scales) x by a factor of a . When a is negative, ax turns x around and stretches it in the other direction by $|a|$. (See figure 2.4.)




Figure 2.4. Scalar multiplication as stretching

With these two operations and distinguished element $\mathbf{0}$, Euclidean space satisfies the following algebraic laws:

Theorem 2.1.1 (Vector Space Axioms).

- (A1) *Addition is associative: $(x + y) + z = x + (y + z)$ for all $x, y, z \in \mathbf{R}^n$.*
- (A2) *$\mathbf{0}$ is an additive identity: $x + \mathbf{0} = x$ for all $x \in \mathbf{R}^n$.*
- (A3) *Existence of additive inverses: For each $x \in \mathbf{R}^n$ there exists $y \in \mathbf{R}^n$ such that $x + y = \mathbf{0}$.*
- (A4) *Addition is commutative: $x + y = y + x$ for all $x, y \in \mathbf{R}^n$.*
- (M1) *Scalar multiplication is associative: $a(bx) = (ab)x$ for all $a, b \in \mathbf{R}$, $x \in \mathbf{R}^n$.*
- (M2) *1 is a multiplicative identity: $1x = x$ for all $x \in \mathbf{R}^n$.*
- (D1) *Scalar multiplication distributes over scalar addition: $(a + b)x = ax + bx$ for all $a, b \in \mathbf{R}$, $x \in \mathbf{R}^n$.*
- (D2) *Scalar multiplication distributes over vector addition: $a(x + y) = ax + ay$ for all $a \in \mathbf{R}$, $x, y \in \mathbf{R}^n$.*

All of these are consequences of how “+” and “.” and $\mathbf{0}$ are defined for \mathbf{R}^n in conjunction with the fact that the real numbers, in turn endowed with “+” and “.” and containing 0 and 1, satisfy the field axioms (see section 1.1). For example, to prove that \mathbf{R}^n satisfies (M1), take any scalars $a, b \in \mathbf{R}$ and any vector $x = (x_1, \dots, x_n) \in \mathbf{R}^n$. Then

$$\begin{aligned}
 a(bx) &= a(b(x_1, \dots, x_n)) && \text{by definition of } x \\
 &= a(bx_1, \dots, bx_n) && \text{by definition of scalar multiplication} \\
 &= (a(bx_1), \dots, a(bx_n)) && \text{by definition of scalar multiplication} \\
 &= ((ab)x_1, \dots, (ab)x_n) && \text{by } n \text{ applications of (m1) in } \mathbf{R} \\
 &= (ab)(x_1, \dots, x_n) && \text{by definition of scalar multiplication} \\
 &= (ab)x && \text{by definition of } x.
 \end{aligned}$$

The other vector space axioms for \mathbf{R}^n can be shown similarly, by unwinding vectors to their coordinates, quoting field axioms coordinatewise, and then bundling the results back up into vectors (see exercise 2.1.3). Nonetheless, the vector space axioms do not perfectly parallel the field axioms, and you are encouraged to spend a little time comparing the two axiom sets to get a feel for where they are similar and where they are different (see exercise 2.1.4). Note in particular that \mathbf{R}^n for $n > 1$ is not endowed with vector-by-vector multiplication. We know that there is a multiplication of vectors for \mathbf{R}^2 , the multiplication of complex numbers; and later (in section 3.10) we will see a noncommutative multiplication of vectors for \mathbf{R}^3 , but these are special cases.

One benefit of the vector space axioms for \mathbf{R}^n is that they are phrased **intrinsically**, meaning that they make no reference to the scalar coordinates of the vectors involved. Thus, once you use coordinates to establish the vector space axioms, your vector algebra can be intrinsic thereafter, making it lighter

and more conceptual. Also, in addition to being intrinsic, the vector space axioms are general. While \mathbf{R}^n is the prototypical set satisfying the vector space axioms, it is by no means the only one. In coming sections we will encounter other sets V (whose elements may be, for example, functions) endowed with their own addition, multiplication by elements of a field F , and distinguished element $\mathbf{0}$. If the vector space axioms are satisfied with V and F replacing \mathbf{R}^n and \mathbf{R} then we say that V is a **vector space over F** .

The pedagogical point here is that although the similarity between vector algebra and scalar algebra may initially make vector algebra seem uninspiring, in fact the similarity is exciting. It makes mathematics easier because familiar algebraic manipulations apply in a wide range of contexts. The same symbol-patterns have more meaning. For example, we use intrinsic vector algebra to show a result from Euclidean geometry, that the three bisectors of a triangle intersect. Consider a triangle with vertices x , y , and z , and form the average of the three vertices,

$$p = \frac{x + y + z}{3}.$$

This algebraic average will be the geometric **center** of the triangle, where the bisectors meet. (See figure 2.5.) Indeed, rewrite p as

$$p = x + \frac{2}{3} \left(\frac{y + z}{2} - x \right).$$

This shows that p is two thirds of the way from x along the line segment from x to the average of y and z , i.e., that p lies on the triangle bisector from vertex x to side yz . (Again see the figure. The idea is that $(y + z)/2$ is being interpreted as the midpoint of y and z , each of these viewed as a point, while on the other hand, the little mnemonic *head minus tail* helps us remember quickly that $(y + z)/2 - x$ can be viewed as the arrow-vector from x to $(y + z)/2$.) Since p is defined symmetrically in x , y , and z , and it lies on one bisector, it therefore lies on the other two bisectors as well. In fact, the vector algebra has shown that it lies two thirds of the way along each bisector.

The **standard basis** of \mathbf{R}^n is the set of vectors

$$\{e_1, e_2, \dots, e_n\}$$

where

$$e_1 = (1, 0, \dots, 0), \quad e_2 = (0, 1, \dots, 0), \quad \dots, \quad e_n = (0, 0, \dots, 1).$$

Any vector $x = (x_1, x_2, \dots, x_n)$ decomposes as

$$\begin{aligned} x &= (x_1, x_2, \dots, x_n) \\ &= (x_1, 0, \dots, 0) + (0, x_2, \dots, 0) + \dots + (0, 0, \dots, x_n) \\ &= x_1(1, 0, \dots, 0) + x_2(0, 1, \dots, 0) + \dots + x_n(0, 0, \dots, 1) \\ &= x_1 e_1 + x_2 e_2 + \dots + x_n e_n, \end{aligned}$$




Figure 2.5. Three bisectors of a triangle

or, more succinctly,

$$x = \sum_{i=1}^n x_i e_i. \quad (2.1)$$

Note that in equation (2.1), x and the e_i are vectors while the x_i are scalars. The equation shows that any $x \in \mathbf{R}^n$ is expressible as a **linear combination** (sum of scalar multiples) of the standard basis vectors. The expression is unique, for if also $x = \sum_{i=1}^n x'_i e_i$ for some scalars x'_1, \dots, x'_n then the equality says that $x = (x'_1, x'_2, \dots, x'_n)$, so that $x'_i = x_i$ for $i = 1, \dots, n$.

The standard basis is handy in that it is a finite set of vectors from which each of the infinitely many vectors of \mathbf{R}^n can be obtained in exactly one way as a linear combination. But it is not the only such set, nor is it always the optimal one.

Definition 2.1.2 (Basis). *A set of vectors $\{f_i\}$ is a **basis** of \mathbf{R}^n if every $x \in \mathbf{R}^n$ is uniquely expressible as a linear combination of the f_i .*

For example, the set $\{f_1, f_2\} = \{(1, 1), (1, -1)\}$ is a basis of \mathbf{R}^2 . To see this, consider an arbitrary vector $(x, y) \in \mathbf{R}^2$. This vector is expressible as a linear combination of f_1 and f_2 if and only if there are scalars a and b such that

$$(x, y) = a f_1 + b f_2.$$

Since $f_1 = (1, 1)$ and $f_2 = (1, -1)$, this vector equation is equivalent to a pair of scalar equations,

$$\begin{aligned} x &= a + b, \\ y &= a - b. \end{aligned}$$

Add these equations and divide by 2 to get $a = (x + y)/2$, and similarly $b = (x - y)/2$. In other words, we have found that

$$(x, y) = \frac{x+y}{2}(1, 1) + \frac{x-y}{2}(1, -1),$$

and the coefficients $a = (x+y)/2$ and $b = (x-y)/2$ on the right side of the equation are the only possible coefficients a and b for the equation to hold. That is, scalars a and b exist to express the vector (x, y) as a linear combination of $\{f_1, f_2\}$, and the scalars are uniquely determined by the vector. This shows that $\{f_1, f_2\}$ is a basis of \mathbf{R}^2 as claimed.

The set $\{g_1, g_2\} = \{(1, 3), (2, 6)\}$ is not a basis of \mathbf{R}^2 , because any linear combination $ag_1 + bg_2$ is $(a+2b, 3a+6b)$, with the second entry equal to three times the first. The vector $(1, 0)$ is therefore not a linear combination of g_1 and g_2 .

Nor is the set $\{h_1, h_2, h_3\} = \{(1, 0), (1, 1), (1, -1)\}$ a basis of \mathbf{R}^2 , because $h_3 = 2h_1 - h_2$, so that h_3 is a nonunique linear combination of the h_j .

See exercises 2.1.9 and 2.1.10 for practice with bases.

Exercises

2.1.1. Write down any three specific nonzero vectors u, v, w from \mathbf{R}^3 and any two specific nonzero scalars a, b from \mathbf{R} . Compute $u + v$, aw , $b(v + w)$, $(a + b)u$, $u + v + w$, abw , and the additive inverse to u .

2.1.2. Give a geometric proof that in \mathbf{R}^2 if we view the vectors x and y as arrows from $\mathbf{0}$ and form the parallelogram P with these arrows as two of its sides, then the diagonal z starting at $\mathbf{0}$ is the vector sum $x + y$ viewed as an arrow.

2.1.3. Verify that \mathbf{R}^n satisfies vector space axioms (A2), (A3), (D1).

2.1.4. Are all the field axioms used in verifying that Euclidean space satisfies the vector space axioms?

2.1.5. Show that $\mathbf{0}$ is the unique additive identity in \mathbf{R}^n . Show that each vector $x \in \mathbf{R}^n$ has a unique additive inverse, which can therefore be denoted $-x$. (And it follows that vector subtraction can now be defined,

$$- : \mathbf{R}^n \times \mathbf{R}^n \longrightarrow \mathbf{R}^n, \quad x - y = x + (-y) \quad \text{for all } x, y \in \mathbf{R}^n.)$$

Show that $0x = \mathbf{0}$ for all $x \in \mathbf{R}^n$.

2.1.6. Repeat the previous exercise, but with \mathbf{R}^n replaced by an arbitrary vector space V over a field F . (Work with the axioms.)

2.1.7. Show the uniqueness of additive identity and additive inverse using only (A1), (A2), (A3). (This is tricky; the opening pages of some books on group theory will help.)

2.1.8. Let x and y be non-collinear vectors in \mathbf{R}^3 . Give a geometric description of the set of all linear combinations of x and y .

2.1.9. Which of the following sets are bases of \mathbf{R}^3 ?

$$\begin{aligned} S_1 &= \{(1, 0, 0), (1, 1, 0), (1, 1, 1)\}, \\ S_2 &= \{(1, 0, 0), (0, 1, 0), (0, 0, 1), (1, 1, 1)\}, \\ S_3 &= \{(1, 1, 0), (0, 1, 1)\}, \\ S_4 &= \{(1, 1, 0), (0, 1, 1), (1, 0, -1)\}. \end{aligned}$$

How many elements do you think a basis for \mathbf{R}^n must have? Give (without proof) geometric descriptions of all bases for \mathbf{R}^2 , for \mathbf{R}^3 .

2.1.10. Recall the field \mathbf{C} of complex numbers. Define **complex n -space** \mathbf{C}^n analogously to \mathbf{R}^n :

$$\mathbf{C}^n = \{(z_1, \dots, z_n) : z_i \in \mathbf{C} \text{ for } i = 1, \dots, n\},$$

and endow it with addition and scalar multiplication defined by the same formulas as for \mathbf{R}^n . Feel free to take my word that under these definitions, \mathbf{C}^n is a vector space over \mathbf{R} and also a vector space over \mathbf{C} . Give a basis in each case.

2.2 Geometry: Length and Angle

The geometric notions of length and angle in \mathbf{R}^n are readily described in terms of the algebraic notion of inner product.

Definition 2.2.1 (Inner Product). *The inner product is a function from pairs of vectors to scalars,*

$$\langle \cdot, \cdot \rangle : \mathbf{R}^n \times \mathbf{R}^n \longrightarrow \mathbf{R},$$

defined by the formula

$$\langle (x_1, \dots, x_n), (y_1, \dots, y_n) \rangle = \sum_{i=1}^n x_i y_i.$$

For example,

$$\begin{aligned} \langle (1, 1, \dots, 1), (1, 2, \dots, n) \rangle &= \frac{n(n+1)}{2}, \\ \langle e_i, e_j \rangle &= \delta_{ij} \text{ (this means 1 if } i = j, 0 \text{ otherwise)}. \end{aligned}$$

Proposition 2.2.2 (Inner Product Properties).

(IP1) *The inner product is positive definite: $\langle x, x \rangle \geq 0$ for all $x \in \mathbf{R}^n$, with equality if and only if $x = \mathbf{0}$.*

(IP2) *The inner product is symmetric: $\langle x, y \rangle = \langle y, x \rangle$ for all $x, y \in \mathbf{R}^n$.*

(IP3) *The inner product is bilinear:*

$$\begin{aligned}\langle x + x', y \rangle &= \langle x, y \rangle + \langle x', y \rangle, & \langle ax, y \rangle &= a\langle x, y \rangle, \\ \langle x, y + y' \rangle &= \langle x, y \rangle + \langle x, y' \rangle, & \langle x, by \rangle &= b\langle x, y \rangle\end{aligned}$$

for all $a, b \in \mathbf{R}$, $x, x', y, y' \in \mathbf{R}^n$.

Proof. Exercise 2.2.4. □

Be aware that the Inner Product Properties do not imply the relation “ $\langle x + x', y + y' \rangle = \langle x, y \rangle + \langle x', y' \rangle$.” In general, this relation does not hold.

Like the vector space axioms, the inner product properties are phrased intrinsically, although they need to be proved using coordinates. As mentioned in the previous section, intrinsic methods are neater and more conceptual than using coordinates. More importantly, *the rest of the results of this section are proved by reference to the inner product properties, with no further reference to the inner product formula.* The notion of an inner product generalizes beyond Euclidean space—this will be demonstrated in exercise 2.3.4, for example—and thanks to the previous sentence, once the properties (IP1) through (IP3) are established for any inner product, all of the pending results in the section will follow automatically with no further work.

Definition 2.2.3 (Modulus). *The modulus (or absolute value) of a vector $x \in \mathbf{R}^n$ is defined as*

$$|x| = \sqrt{\langle x, x \rangle}.$$

Thus the modulus is defined in terms of the inner product, rather than by its own formula. The inner product formula shows that the modulus formula is

$$|(x_1, \dots, x_n)| = \sqrt{x_1^2 + \dots + x_n^2},$$

so that some particular examples are

$$\begin{aligned}|(1, 2, \dots, n)| &= \sqrt{\frac{n(n+1)(2n+1)}{6}}, \\ |e_i| &= 1.\end{aligned}$$

However, the definition of the modulus in terms of inner product combines with the inner product properties to show, with no reference to the inner product formula or the modulus formula, that the modulus satisfies (exercise 2.2.5)

Proposition 2.2.4 (Modulus Properties).

(Mod1) *The modulus is positive: $|x| \geq 0$ for all $x \in \mathbf{R}^n$, with equality if and only if $x = \mathbf{0}$.*

(Mod2) *The modulus is absolute-homogeneous: $|ax| = |a||x|$ for all $a \in \mathbf{R}$ and $x \in \mathbf{R}^n$.*

Like other symbols, the absolute value signs are now overloaded, but their meaning can be inferred from context, as in property (Mod2). When n is 1, 2, or 3, the modulus $|x|$ gives the distance from $\mathbf{0}$ to the point x , or the length of x viewed as an arrow. (See figure 2.6.)




Figure 2.6. Modulus as length

The following relation between inner product and modulus will help to show that distance in \mathbf{R}^n behaves as it should and that angle in \mathbf{R}^n makes sense. Since the relation is not obvious, its proof is a little sneaky.

Theorem 2.2.5 (Cauchy–Schwarz Inequality). *For all $x, y \in \mathbf{R}^n$,*

$$|\langle x, y \rangle| \leq |x| |y|,$$

with equality if and only if one of x, y is a scalar multiple of the other.

Note that the absolute value signs mean different things on each side of the Cauchy–Schwarz inequality. On the left side, the quantities x and y are vectors, their inner product $\langle x, y \rangle$ is a scalar, and $|\langle x, y \rangle|$ is its scalar absolute value, while on the right side, $|x|$ and $|y|$ are the scalar absolute values of vectors, and $|x| |y|$ is their product.

The Cauchy–Schwarz inequality can be written out in coordinates, temporarily abandoning the principle that we should avoid reference to formulas,

$$(x_1 y_1 + \cdots + x_n y_n)^2 \leq (x_1^2 + \cdots + x_n^2)(y_1^2 + \cdots + y_n^2).$$

And this inequality can be proved bare-handedly, as follows (the reader is encouraged only to skim the following computation). The desired inequality rewrites as

$$\left(\sum_i x_i y_i \right)^2 \leq \sum_i x_i^2 \cdot \sum_j y_j^2,$$

where the indices of summation run from 1 to n . This expands to

$$\sum_i x_i^2 y_i^2 + \sum_{\substack{i,j \\ i \neq j}} x_i y_i x_j y_j \leq \sum_{i,j} x_i^2 y_j^2,$$

and canceling the terms common to both sides reduces it to

$$\sum_{i \neq j} x_i y_i x_j y_j \leq \sum_{i \neq j} x_i^2 y_j^2,$$

or

$$\sum_{i \neq j} (x_i^2 y_j^2 - x_i y_i x_j y_j) \geq 0.$$

Rather than sum over all pairs (i, j) with $i \neq j$, sum over the pairs with $i < j$, collecting the (i, j) -term and the (j, i) -term for each such pair. This makes the previous inequality

$$\sum_{i < j} (x_i^2 y_j^2 + x_j^2 y_i^2 - 2x_i y_j x_j y_i) \geq 0.$$

Thus the desired inequality has reduced to a true inequality,

$$\sum_{i < j} (x_i y_j - x_j y_i)^2 \geq 0.$$

So the main proof is done, although there is still the question of when equality holds.

But this can't be the graceful way to argue. The computation draws on the minutiae of the formulas for the inner product and the modulus, rather than using their properties. It is uninformative, making the Cauchy–Schwarz inequality look like a low-level accident. To prove the inequality in a way that is enlightening and general, we should work intrinsically, keeping the scalars $\langle x, y \rangle$, $|x|$, and $|y|$ notated in their concise forms, and we should use properties, not formulas. The idea is that the calculation in coordinates reduces to the fact that squares are nonnegative. That is, the Cauchy–Schwarz inequality is somehow quadratically hard, and its verification amounted to completing many squares. The argument to be given here is guided by this insight to prove the inequality by citing facts about quadratic polynomials, facts established by completing one square back in high school algebra at the moment that doing so was called for. This eliminates redundancy and clutter. So the argument to follow will involve an auxiliary object, a quadratic polynomial, but in return it will become coherent.

Proof. The result is clear when $x = \mathbf{0}$, so assume $x \neq \mathbf{0}$. For any $a \in \mathbf{R}$,

$$\begin{aligned} 0 &\leq \langle ax - y, ax - y \rangle && \text{by positive definiteness} \\ &= a \langle x, ax - y \rangle - \langle y, ax - y \rangle && \text{by linearity in the first variable} \\ &= a^2 \langle x, x \rangle - a \langle x, y \rangle - a \langle y, x \rangle + \langle y, y \rangle && \text{by linearity in the second variable} \\ &= |x|^2 a^2 - 2 \langle x, y \rangle a + |y|^2 && \text{by symmetry, definition of modulus.} \end{aligned}$$

View the right side as a quadratic polynomial in the scalar variable a , where the scalar coefficients of the polynomial depend on the generic but fixed vectors x and y ,

$$f(a) = |x|^2 a^2 - 2\langle x, y \rangle a + |y|^2.$$

We have shown that $f(a)$ is always nonnegative, so f has at most one root. Thus by the quadratic formula its discriminant is nonpositive,

$$4\langle x, y \rangle^2 - 4|x|^2|y|^2 \leq 0,$$

and the inequality follows. Equality holds exactly when the quadratic polynomial $f(a) = |ax - y|^2$ has a root a , i.e., exactly when $ax - y = 0$ for some $a \in \mathbf{R}$. \square

Geometrically, the condition for equality in Cauchy–Schwarz is that the vectors x and y , viewed as arrows at the origin, are parallel, though perhaps pointing in opposite directions. A geometrically conceived proof of Cauchy–Schwarz is given in exercise 2.2.15 to complement the algebraic argument that has been given here.

The Cauchy–Schwarz inequality shows that the modulus function satisfies

Theorem 2.2.6 (Triangle Inequality). *For all $x, y \in \mathbf{R}^n$,*

$$|x + y| \leq |x| + |y|,$$

with equality if and only if one of x, y is a nonnegative scalar multiple of the other.

Proof. To show this, compute

$$\begin{aligned} |x + y|^2 &= \langle x + y, x + y \rangle \\ &= |x|^2 + 2\langle x, y \rangle + |y|^2 \quad \text{by bilinearity} \\ &\leq |x|^2 + 2|x||y| + |y|^2 \quad \text{by Cauchy–Schwarz} \\ &= (|x| + |y|)^2, \end{aligned}$$

proving the inequality. Equality holds exactly when $\langle x, y \rangle = |x||y|$, or equivalently when $|\langle x, y \rangle| = |x||y|$ and $\langle x, y \rangle \geq 0$. These hold when one of x, y is a scalar multiple of the other and the scalar is nonnegative. \square

The Triangle Inequality’s name is explained by its geometric interpretation in \mathbf{R}^2 and \mathbf{R}^3 . View x as an arrow at the origin, y as an arrow with tail at the head of x , and $x + y$ as an arrow at the origin. These three arrows form a triangle, and the assertion is that the lengths of two sides sum to at least the length of the third. (See figure 2.7.)

The **full Triangle Inequality** says that for all $x, y \in \mathbf{R}^n$,

$$||x| - |y|| \leq |x \pm y| \leq |x| + |y|.$$




Figure 2.7. Sides of a triangle

The proof is exercise 2.2.7.

A small argument, which can be formalized as induction if one is painstaking, shows that the basic Triangle Inequality extends from two vectors to any finite number of vectors. For example,

$$|x + y + z| \leq |x + y| + |z| \leq |x| + |y| + |z|.$$

The only obstacle to generalizing the basic Triangle Inequality in this fashion is notation. The argument can't use the symbol n to denote the number of vectors since n already denotes the dimension of the Euclidean space where we are working; and furthermore, the vectors can't be denoted with subscripts since for now a subscript denotes a component of an individual vector. Thus, for now we are stuck with something like

$$|x^{(1)} + \cdots + x^{(k)}| \leq |x^{(1)}| + \cdots + |x^{(k)}| \quad \text{for all } x^{(1)}, \dots, x^{(k)} \in \mathbf{R}^n,$$

or

$$\left| \sum_{i=1}^k x^{(i)} \right| \leq \sum_{i=1}^k |x^{(i)}|, \quad x^{(1)}, \dots, x^{(k)} \in \mathbf{R}^n.$$

As our work with vectors becomes more intrinsic, vector entries will demand less of our attention, and we will be able to denote vectors by subscripts. The notation-change will be implemented in the next section.

For any vector $x = (x_1, \dots, x_n) \in \mathbf{R}^n$, useful bounds on the modulus $|x|$ in terms of the scalar absolute values $|x_i|$ are

Proposition 2.2.7 (Size Bounds). *For any $j \in \{1, \dots, n\}$,*

$$|x_j| \leq |x| \leq \sum_{i=1}^n |x_i|.$$

The proof is exercise 2.2.8.

The modulus gives rise to a distance function on \mathbf{R}^n that behaves as distance should. Define

$$d: \mathbf{R}^n \times \mathbf{R}^n \longrightarrow \mathbf{R}$$

by

$$d(x, y) = |y - x|.$$

For example, $d(e_i, e_j) = \sqrt{2}(1 - \delta_{ij})$.

Theorem 2.2.8 (Distance Properties).

(D1) *Distance is positive: $d(x, y) \geq 0$ for all $x, y \in \mathbf{R}^n$, and $d(x, y) = 0$ if and only if $x = y$.*

(D2) *Distance is symmetric: $d(x, y) = d(y, x)$ for all $x, y \in \mathbf{R}^n$.*

(D3) *Triangle Inequality: $d(x, z) \leq d(x, y) + d(y, z)$ for all $x, y, z \in \mathbf{R}^n$.*

(D1) and (D2) are clearly desirable as properties of a distance function. Property (D3) says that you can't shorten your trip from x to z by making a stop at y .

Proof. Exercise 2.2.9. □

The Cauchy–Schwarz inequality also lets us define the angle between two nonzero vectors in terms of the inner product. If x and y are nonzero vectors in \mathbf{R}^n , define their angle $\theta_{x,y}$ by the condition

$$\cos \theta_{x,y} = \frac{\langle x, y \rangle}{|x||y|}, \quad 0 \leq \theta_{x,y} \leq \pi. \quad (2.2)$$

This makes sense because $-1 \leq \frac{\langle x, y \rangle}{|x||y|} \leq 1$ by Cauchy–Schwarz. For example, $\cos \theta_{(1,0), (1,1)} = 1/\sqrt{2}$, so $\theta_{(1,0), (1,1)} = \pi/4$. In particular, two nonzero vectors x, y are **orthogonal** when $\langle x, y \rangle = 0$. Naturally, we would like $\theta_{x,y}$ to correspond to the usual notion of angle, at least in \mathbf{R}^2 , and indeed it does—see exercise 2.2.10. For convenience, define any two vectors x and y to be orthogonal if $\langle x, y \rangle = 0$, thus making $\mathbf{0}$ orthogonal to all vectors.

Rephrasing geometry in terms of intrinsic vector algebra not only extends the geometric notions of length and angle uniformly to any dimension, it also makes some low-dimensional geometry easier. For example, vectors show in a natural way that the three altitudes of any triangle must meet. Let x and y denote two sides of the triangle, making the third side $x - y$ by the *head minus tail* mnemonic. Let p be the point where the altitudes to x and y meet. (See figure 2.8, which also shows the third altitude.) Thus

$$p - y \perp x \quad \text{and} \quad p - x \perp y.$$

We want to show that also p lies on the third altitude, i.e., that

$$p \perp x - y.$$

To rephrase matters in terms of inner products, we know that

$$\langle p - y, x \rangle = 0 \quad \text{and} \quad \langle p - x, y \rangle = 0,$$

and we want to show that

$$\langle p, x - y \rangle = 0.$$

Since the inner product is linear in each of its arguments, a further rephrase is that we know that

$$\langle p, x \rangle = \langle y, x \rangle \quad \text{and} \quad \langle p, y \rangle = \langle x, y \rangle,$$

and we want to show that

$$\langle p, x \rangle = \langle p, y \rangle.$$

But since the inner product is symmetric, this is immediate: $\langle p, x \rangle$ and $\langle p, y \rangle$ both equal $\langle x, y \rangle$, and hence they equal each other. The point where the three altitudes meet is called the **orthocenter** of the triangle. In general the orthocenter of a triangle is not the center, cf. the previous section.




Figure 2.8. Three altitudes of a triangle

Exercises

2.2.1. Let $x = (\frac{\sqrt{3}}{2}, -\frac{1}{2}, 0)$, $y = (\frac{1}{2}, \frac{\sqrt{3}}{2}, 1)$, $z = (1, 1, 1)$. Compute $\langle x, x \rangle$, $\langle x, y \rangle$, $\langle y, z \rangle$, $|x|$, $|y|$, $|z|$, $\theta_{x,y}$, θ_{y,e_1} , θ_{z,e_2} .

2.2.2. Show that the points $x = (2, -1, 3, 1)$, $y = (4, 2, 1, 4)$, $z = (1, 3, 6, 1)$ form the vertices of a triangle in \mathbf{R}^4 with two equal angles.

2.2.3. Explain why for all $x \in \mathbf{R}^n$, $x = \sum_{j=1}^n \langle x, e_j \rangle e_j$.

2.2.4. Prove the Inner Product Properties.

2.2.5. Use the Inner Product Properties and the definition of the modulus in terms of the inner product to prove the Modulus Properties.

2.2.6. In the text, the modulus is defined in terms of the inner product. Prove that this can be turned around by showing that for every $x, y \in \mathbf{R}^n$,

$$\langle x, y \rangle = \frac{|x + y|^2 - |x - y|^2}{4}.$$

2.2.7. Prove the full Triangle Inequality: for any $x, y \in \mathbf{R}^n$,

$$||x| - |y|| \leq |x \pm y| \leq |x| + |y|.$$

Do not do this by writing three more variants of the proof of the Triangle Inequality, but by substituting suitably into the basic Triangle Inequality, which is already proved.

2.2.8. Let $x = (x_1, \dots, x_n) \in \mathbf{R}^n$. Prove the Size Bounds: For any $j \in \{1, \dots, n\}$,

$$|x_j| \leq |x| \leq \sum_{i=1}^n |x_i|.$$

When can each “ \leq ” be an “ $=$ ”?

2.2.9. Prove the Distance Properties.

2.2.10. In \mathbf{R}^2 , depict the nonzero vectors x and y as arrows from the origin and depict $x - y$ as an arrow from the endpoint of y to the endpoint of x . Let θ denote the angle (in the usual geometric sense) between x and y . Use the Law of Cosines to show that

$$\cos \theta = \frac{\langle x, y \rangle}{|x||y|},$$

so that our notion of angle agrees with the geometric one, at least in \mathbf{R}^2 .

2.2.11. Prove that for any nonzero $x \in \mathbf{R}^n$, $\sum_{i=1}^n \cos^2 \theta_{x, e_i} = 1$.

2.2.12. Prove that two nonzero vectors x, y are orthogonal if and only if $|x + y|^2 = |x|^2 + |y|^2$.

2.2.13. Use vectors in \mathbf{R}^2 to show that the diagonals of a parallelogram are perpendicular if and only if the parallelogram is a rhombus.

2.2.14. Use vectors to show that every angle inscribed in a semicircle is right.

2.2.15. Let x and y be vectors, with x nonzero. Define the **parallel** component of y along x and the **normal** component of y to x to be

$$y_{(\parallel x)} = \frac{\langle x, y \rangle}{|x|^2} x \quad \text{and} \quad y_{(\perp x)} = y - y_{(\parallel x)}.$$

(a) Show that $y = y_{(\parallel x)} + y_{(\perp x)}$; show that $y_{(\parallel x)}$ is a scalar multiple of x ; show that $y_{(\perp x)}$ is orthogonal to x . Show that the decomposition of y as a sum of vectors parallel and perpendicular to x is unique. Draw an illustration.

(b) Show that

$$|y|^2 = |y_{(\parallel x)}|^2 + |y_{(\perp x)}|^2.$$

What theorem from classical geometry does this encompass?

(c) Explain why it follows from (b) that

$$|y_{(\|x\|)}| \leq |y|,$$

with equality if and only y is a scalar multiple of x . Use this inequality to give another proof of the Cauchy–Schwarz Inequality. This argument gives the geometric content of Cauchy–Schwarz: *The parallel component of one vector along another is at most as long as the original vector.*

(d) The proof of the Cauchy–Schwarz inequality in part (c) refers to parts (a) and (b), part (a) refers to orthogonality, orthogonality refers to an angle, and as explained in the text, the fact that angles make sense depends on the Cauchy–Schwarz inequality. This suggests that the proof in part (c) relies on circular logic. Explain why the logic is in fact not circular.

2.2.16. Given nonzero vectors x_1, x_2, \dots, x_n in \mathbf{R}^n , the **Gram–Schmidt process** is to set

$$\begin{aligned} x'_1 &= x_1 \\ x'_2 &= x_2 - (x_2)_{(\|x'_1\|)} \\ x'_3 &= x_3 - (x_3)_{(\|x'_2\|)} - (x_3)_{(\|x'_1\|)} \\ &\vdots \\ x'_n &= x_n - (x_n)_{(\|x'_{n-1}\|)} - \cdots - (x_n)_{(\|x'_1\|)}. \end{aligned}$$

(a) What is the result of applying the Gram–Schmidt process to the vectors $x_1 = (1, 0, 0)$, $x_2 = (1, 1, 0)$, and $x_3 = (1, 1, 1)$?

(b) Returning to the general case, show that x'_1, \dots, x'_n are pairwise orthogonal and that each x'_j has the form

$$x'_j = a_{j1}x_1 + a_{j2}x_2 + \cdots + a_{j,j-1}x_{j-1} + x_j.$$

Thus any linear combination of the new $\{x'_j\}$ is also a linear combination of the original $\{x_j\}$. The converse is also true and will be shown in exercise 3.3.13.

2.3 Analysis: Continuous Mappings

A **mapping** from \mathbf{R}^n to \mathbf{R}^m is some rule that assigns to each point x in \mathbf{R}^n a point in \mathbf{R}^m . Generally mappings will be denoted by letters such as f, g, h . When $m = 1$ we say “function” instead of mapping. For example, the mapping

$$f : \mathbf{R}^2 \longrightarrow \mathbf{R}^2$$

defined by

$$f(x, y) = (x^2 - y^2, 2xy)$$

takes the real and imaginary parts of a complex number $z = x + iy$ and returns the real and imaginary parts of z^2 . By the nature of multiplication of complex numbers, this means that each output point has modulus equal to the square of the modulus of the input point and has angle equal to twice the angle of the input point. Make sure that you see how this is shown in figure 2.9.




Figure 2.9. The complex square as a mapping from \mathbf{R}^2 to \mathbf{R}^2

Mappings expressed by formulas may be undefined at certain points (e.g., $f(x) = 1/|x|$ is undefined at $\mathbf{0}$), so we need to restrict their domains. For a given dimension n , a given set $A \subset \mathbf{R}^n$, and a second dimension m , let $\mathcal{M}(A, \mathbf{R}^m)$ denote the set of all mappings $f : A \rightarrow \mathbf{R}^m$. This set forms a vector space over \mathbf{R} (whose “points” are functions) under the operations

$$+ : \mathcal{M}(A, \mathbf{R}^m) \times \mathcal{M}(A, \mathbf{R}^m) \rightarrow \mathcal{M}(A, \mathbf{R}^m),$$

defined by

$$(f + g)(x) = f(x) + g(x) \quad \text{for all } x \in A,$$

and

$$\cdot : \mathbf{R} \times \mathcal{M}(A, \mathbf{R}^m) \rightarrow \mathcal{M}(A, \mathbf{R}^m),$$

defined by

$$(a \cdot f)(x) = a \cdot f(x) \quad \text{for all } x \in A.$$

As usual, “+” and “ \cdot ” are overloaded: on the left they denote operations on $\mathcal{M}(A, \mathbf{R}^m)$, while on the right they denote the operations on \mathbf{R}^m defined in section 2.1. Also as usual, the “ \cdot ” is generally omitted. The origin in $\mathcal{M}(A, \mathbf{R}^m)$ is the **zero mapping**, $\mathbf{0} : A \rightarrow \mathbf{R}^m$, defined by

$$\mathbf{0}(x) = \mathbf{0}_m \quad \text{for all } x \in A.$$

For example, to verify that $\mathcal{M}(A, \mathbf{R}^m)$ satisfies (A1), consider any functions $f, g, h \in \mathcal{M}(A, \mathbf{R}^m)$. For any $x \in A$,

$$\begin{aligned}
((f + g) + h)(x) &= (f + g)(x) + h(x) && \text{by definition of “+” in } \mathcal{M}(A, \mathbf{R}^m) \\
&= (f(x) + g(x)) + h(x) && \text{by definition of “+” in } \mathcal{M}(A, \mathbf{R}^m) \\
&= f(x) + (g(x) + h(x)) && \text{by associativity of “+” in } \mathbf{R}^m \\
&= f(x) + (g + h)(x) && \text{by definition of “+” in } \mathcal{M}(A, \mathbf{R}^m) \\
&= (f + (g + h))(x) && \text{by definition of “+” in } \mathcal{M}(A, \mathbf{R}^m).
\end{aligned}$$

Since x was arbitrary, $(f + g) + h = f + (g + h)$.

Let A be a subset of \mathbf{R}^n . A **sequence in A** is an infinite list of vectors x_1, x_2, x_3, \dots in A , often written $\{x_\nu\}$. (The symbol n is already in use, so its Greek counterpart ν —pronounced nu—is used as the index-counter.) Since a vector has n entries, each vector x_ν in the sequence takes the form $(x_{1,\nu}, \dots, x_{n,\nu})$. Convergence of vector sequences is defined as for scalar sequences.

Definition 2.3.1 (Sequence Convergence, Sequence Limit). *Let A be a subset of \mathbf{R}^n . Consider a sequence $\{x_\nu\}$ in A and a point $a \in \mathbf{R}^n$. The sequence $\{x_\nu\}$ **converges to a** (or **has limit a**), written $\{x_\nu\} \rightarrow a$, if for every $\varepsilon > 0$ there exists some ν_0 such that*

$$\text{if } \nu > \nu_0 \text{ then } |x_\nu - a| < \varepsilon.$$

*When the limit a is a point of A , the sequence $\{x_\nu\}$ **converges in A** .*

In particular, a **null sequence** in A is a sequence that converges to $\mathbf{0}_n$. Since for any vector $x \in \mathbf{R}^n$,

$$|x - \mathbf{0}_n| = |x| = ||x| - 0|,$$

it follows that a vector sequence $\{x_\nu\}$ is null if and only if the scalar sequence $\{|x_\nu|\}$ is null.

Lemma 2.3.2 (Componentwise Nature of Nullness). *The vector sequence $\{(x_{1,\nu}, \dots, x_{n,\nu})\}$ is null if and only if each of its component scalar sequences $\{x_{j,\nu}\}$ ($j \in \{1, \dots, n\}$) is null.*

Proof. By the preceding remark it suffices to show that $\{|(x_{1,\nu}, \dots, x_{n,\nu})|\}$ is null if and only if each $\{|x_{j,\nu}|\}$ is null. The Size Bounds give for any $j \in \{1, \dots, n\}$ and any ν ,

$$|x_{j,\nu}| \leq |(x_{1,\nu}, \dots, x_{n,\nu})| \leq \sum_{i=1}^n |x_{i,\nu}|.$$

If $\{|(x_{1,\nu}, \dots, x_{n,\nu})|\}$ is null, then by the first inequality so is each $\{|x_{j,\nu}|\}$. On the other hand, if each $\{|x_{j,\nu}|\}$ is null then so is $\{\sum_{i=1}^n |x_{i,\nu}|\}$, and thus by the second inequality $\{|(x_{1,\nu}, \dots, x_{n,\nu})|\}$ is null as well. \square

Any sequence $\{x_\nu\}$ converges to a if and only if $\{x_\nu - a\}$ is null, so the corollary is (exercise 2.3.5)

Proposition 2.3.3 (Componentwise Nature of Convergence). *The vector sequence $\{(x_{1,\nu}, \dots, x_{n,\nu})\}$ converges to the vector (a_1, \dots, a_n) if and only if each component scalar sequence $\{x_{j,\nu}\}$ ($j = 1, \dots, n$) converges to the scalar a_j .*

Vector versions of the Sum Rule and the Constant Multiple Rule for convergent sequences follow immediately from their scalar counterparts and from the componentwise nature of convergence.

Proposition 2.3.4 (Linearity of Convergence). *Let $\{x_\nu\}$ be a sequence in \mathbf{R}^n converging to a , and let $\{y_\nu\}$ be a sequence in \mathbf{R}^n converging to b . Let c be a scalar. Then the sequence $\{x_\nu + y_\nu\}$ converges to $a + b$, and the sequence $\{cx_\nu\}$ converges to ca .*

Continuity, like convergence, is typographically indistinguishable in \mathbf{R} and \mathbf{R}^n .

Definition 2.3.5 (Continuity). *Let A be a subset of \mathbf{R}^n , let $f : A \rightarrow \mathbf{R}^m$ be a mapping, and let a be a point of A . Then f is **continuous at a** if for every sequence $\{x_\nu\}$ in A converging to a , the sequence $\{f(x_\nu)\}$ converges to $f(a)$. The mapping f is **continuous on A** (or just **continuous** when A is clearly established) if it is continuous at each point $a \in A$.*

For example, the modulus function

$$| \cdot | : \mathbf{R}^n \rightarrow \mathbf{R}$$

is continuous on \mathbf{R}^n . To see this, consider any point $a \in \mathbf{R}^n$ and consider any sequence $\{x_\nu\}$ in \mathbf{R}^n that converges to a . We need to show that the sequence $\{|x_\nu|\}$ in \mathbf{R} converges to $|a|$. But by the full Triangle Inequality,

$$||x_\nu| - |a|| \leq |x_\nu - a|.$$

Since the right side is the ν th term of a null sequence, so is the left, giving the result.

For another example, let $v \in \mathbf{R}^n$ be any fixed vector and consider the function defined by taking the inner product of this vector with other vectors,

$$T : \mathbf{R}^n \rightarrow \mathbf{R}, \quad T(x) = \langle v, x \rangle.$$

This function is also continuous on \mathbf{R}^n . To see this, again consider any $a \in \mathbf{R}^n$ and any sequence $\{x_\nu\}$ in \mathbf{R}^n converging to a . Then the definition of T , the bilinearity of the inner product, and the Cauchy-Schwarz inequality combine to show that

$$|T(x_\nu) - T(a)| = |\langle v, x_\nu \rangle - \langle v, a \rangle| = |\langle v, x_\nu - a \rangle| \leq |v| |x_\nu - a|.$$

Since $|v|$ is a constant, the right side is the ν th term of a null sequence, hence so is the left, and the proof is complete. We will refer to this example in section 3.1. Also, note that as a special case of this example, take any $j \in \{1, \dots, n\}$, and set the fixed vector v to e_j . This shows that the j th **coordinate function map**,

$$\pi_j : \mathbf{R}^n \longrightarrow \mathbf{R}, \quad \pi_j(x_1, \dots, x_n) = x_j,$$

is continuous.

Proposition 2.3.6 (Linearity of Continuity). *Let A be a subset of \mathbf{R}^n , let $f, g : A \longrightarrow \mathbf{R}^m$ be continuous mappings, and let $c \in \mathbf{R}$. Then the sum and the scalar multiple mappings*

$$f + g, cf : A \longrightarrow \mathbf{R}^m$$

are continuous.

This follows immediately from the Linearity of Convergence and from the definition of continuity. Another consequence of the definition of continuity is

Proposition 2.3.7 (Persistence of Continuity Under Composition).

Let A be a subset of \mathbf{R}^n , and let $f : A \longrightarrow \mathbf{R}^m$ be a continuous mapping. Let B be a superset of $f(A)$ in \mathbf{R}^m , and let $g : B \longrightarrow \mathbf{R}^\ell$ be a continuous mapping. Then the composition mapping

$$g \circ f : A \longrightarrow \mathbf{R}^\ell$$

is continuous.

Let A be a subset of \mathbf{R}^n . Any mapping $f : A \longrightarrow \mathbf{R}^m$ decomposes as m functions f_1, \dots, f_m with each $f_i : A \longrightarrow \mathbf{R}$, by the formula

$$f(x) = (f_1(x), \dots, f_m(x)).$$

For example, if $f(x, y) = (x^2 - y^2, 2xy)$ then $f_1(x, y) = x^2 - y^2$ and $f_2(x, y) = 2xy$. The decomposition of f can also be written

$$f(x) = \sum_{i=1}^m f_i(x)e_i,$$

or equivalently, the functions f_i are defined by the condition

$$f_i(x) = f(x)_i \quad \text{for } i = 1, \dots, m.$$

Conversely, given m functions f_1, \dots, f_m from A to \mathbf{R} , any of the preceding three displayed formulas assembles a mapping $f : A \longrightarrow \mathbf{R}^m$. Thus, each mapping f determines and is determined by its **component functions** f_1, \dots, f_m . Conveniently, to check continuity of the vector-valued mapping f we only need to check its scalar-valued component functions.

Theorem 2.3.8 (Componentwise Nature of Continuity). *Let $A \subset \mathbf{R}^n$, let $f : A \rightarrow \mathbf{R}^m$ have component functions f_1, \dots, f_m , and let a be a point in A . Then*

$$f \text{ is continuous at } a \iff \text{each } f_i \text{ is continuous at } a.$$

This follows from the componentwise nature of convergence and is left as exercise 2.3.6.

Let A be a subset of \mathbf{R}^n , let f and g be continuous functions from A to \mathbf{R} , and let $c \in \mathbf{R}$. Then the familiar Sum Rule, Constant Multiple Rule, Product Rule, and Quotient Rule for continuous functions hold. That is, the sum $f + g$, the constant multiple cf , the product fg , and the quotient f/g (at points $a \in A$ such that $g(a) \neq 0$) are again continuous. The first two of these facts are special cases of the Linearity of Continuity principle. The proofs of the other two are typographically identical to their one-variable counterparts. With the various continuity results so far in hand, it is clear that a function such as

$$f : \mathbf{R}^3 \rightarrow \mathbf{R}, \quad f(x, y, z) = \frac{\sin(\sqrt{x^2 + y^2 + z^2})}{e^{xy+z}}$$

is continuous. The continuity of such functions, and of mappings with such functions as their components, will go without comment from now on.

However, the continuity of functions of n variables also has new, subtle features when $n > 1$. In \mathbf{R} , a sequence $\{x_\nu\}$ can approach the point a in only two essential ways: from the left and from the right. But in \mathbf{R}^n where $n \geq 2$, $\{x_\nu\}$ can approach a from infinitely many directions, or not approach along a line at all, so the convergence of $\{f(x_\nu)\}$ can be trickier. For example, consider the function $f : \mathbf{R}^2 \rightarrow \mathbf{R}$ defined by

$$f(x, y) = \begin{cases} \frac{2xy}{x^2+y^2} & \text{if } (x, y) \neq \mathbf{0}, \\ b & \text{if } (x, y) = \mathbf{0}. \end{cases}$$

Can the constant b be specified to make f continuous at $\mathbf{0}$?

It can't. Take a sequence $\{(x_\nu, y_\nu)\}$ approaching $\mathbf{0}$ along the line $y = mx$ of slope m . For any point (x_ν, y_ν) of this sequence,

$$f(x_\nu, y_\nu) = f(x_\nu, mx_\nu) = \frac{2x_\nu mx_\nu}{x_\nu^2 + m^2 x_\nu^2} = \frac{2mx_\nu^2}{(1 + m^2)x_\nu^2} = \frac{2m}{1 + m^2}.$$

Thus, as $\{(x_\nu, y_\nu)\}$ approaches $\mathbf{0}$ along the line of slope m , $f(x, y)$ holds steady at $2m/(1 + m^2)$, and so $f(\mathbf{0})$ needs to take this value for continuity. Taking sequences $\{(x_\nu, y_\nu)\}$ that approach $\mathbf{0}$ along lines of different slope shows that $f(\mathbf{0})$ needs to take different values for continuity, and hence f can not be made continuous at $\mathbf{0}$. The graph of f away from $\mathbf{0}$ is a sort of spiral staircase, and no height over $\mathbf{0}$ is compatible with all the stairs. (See figure 2.10. The figure displays only the portion of the graph for slopes between 0 and 1 in the input plane.)




Figure 2.10. A spiral staircase

The last example was actually fairly simple in that we only needed to study $f(x, y)$ as (x, y) approached $\mathbf{0}$ along straight lines. Consider the function $g : \mathbf{R}^2 \rightarrow \mathbf{R}$ defined by

$$g(x, y) = \begin{cases} \frac{x^2 y}{x^4 + y^2} & \text{if } (x, y) \neq \mathbf{0}, \\ b & \text{if } (x, y) = \mathbf{0}. \end{cases}$$

For any nonzero slope m , take a sequence $\{(x_\nu, y_\nu)\}$ approaching $\mathbf{0}$ along the line $y = mx$. Compute that for each point of this sequence,

$$g(x_\nu, y_\nu) = g(x_\nu, mx_\nu) = \frac{mx_\nu^3}{x_\nu^4 + m^2 x_\nu^2} = \frac{mx_\nu}{x_\nu^2 + m^2}.$$

This quantity tends to 0 as x_ν goes to 0. That is, as $\{(x_\nu, y_\nu)\}$ approaches $\mathbf{0}$ along the line of slope m , $g(x, y)$ approaches 0, and so $g(\mathbf{0})$ needs to take the value 0 for continuity. Since g is 0 at the nonzero points of either axis in the (x, y) -plane, this requirement extends to the cases that $\{(x_\nu, y_\nu)\}$ approaches $\mathbf{0}$ along a horizontal or vertical line. However, next consider a sequence $\{(x_\nu, y_\nu)\}$ approaching $\mathbf{0}$ along the parabola $y = x^2$. For each point of this sequence,

$$g(x_\nu, y_\nu) = g(x_\nu, x_\nu^2) = \frac{x_\nu^4}{x_\nu^4 + x_\nu^4} = \frac{1}{2}.$$

Thus, as $\{(x_\nu, y_\nu)\}$ approaches $\mathbf{0}$ along the parabola, $g(x, y)$ holds steady at $1/2$, and so $g(\mathbf{0})$ needs to be $1/2$ for continuity as well. This shows that g can not be made continuous at $\mathbf{0}$, even though approaching $\mathbf{0}$ only along lines suggests that it can.

Thus, given a function $f : \mathbf{R}^2 \rightarrow \mathbf{R}$, letting $\{(x_\nu, y_\nu)\}$ approach $\mathbf{0}$ along lines can disprove continuity at $\mathbf{0}$, but it can only suggest continuity at $\mathbf{0}$, not prove it. To prove continuity, the Size Bounds may be helpful. For example, let

$$h(x, y) = \begin{cases} \frac{x^3}{x^2 + y^2} & \text{if } (x, y) \neq \mathbf{0}, \\ b & \text{if } (x, y) = \mathbf{0}. \end{cases}$$

Can b be specified to make h continuous at $\mathbf{0}$? The estimate $|x| \leq |(x, y)|$ gives for any $(x, y) \neq \mathbf{0}$,

$$0 \leq |h(x, y)| = \frac{|x|^3}{|(x, y)|^2} \leq \frac{|(x, y)|^3}{|(x, y)|^2} = |(x, y)|,$$

so as any sequence $\{(x_\nu, y_\nu)\}$ of nonzero vectors converges to $\mathbf{0}$, the corresponding sequence of outputs $\{h(x_\nu, y_\nu)\}$ is squeezed to 0 in absolute value and hence converges to 0. Setting $b = 0$ makes h continuous at $\mathbf{0}$.

Returning to the spiral staircase example,

$$f(x, y) = \begin{cases} \frac{2xy}{x^2+y^2} & \text{if } (x, y) \neq \mathbf{0}, \\ b & \text{if } (x, y) = \mathbf{0}, \end{cases}$$

the Size Bounds show that that for any $(x, y) \neq \mathbf{0}$,

$$0 \leq |f(x, y)| = \frac{2|x||y|}{|(x, y)|^2} \leq \frac{2|(x, y)|^2}{|(x, y)|^2} = 2.$$

This tells us only that as a sequence of inputs $\{(x_\nu, y_\nu)\}$ approaches $\mathbf{0}$, the sequence of outputs $\{f(x_\nu, y_\nu)\}$ might converge to some limit between -2 and 2 . The outputs needn't converge to 0 (or converge at all), but according to this diagnostic they possibly could. Thus the Size Bounds tell us only that f *could* be discontinuous at $(0, 0)$, but they give no conclusive information.

In sum, these examples show that

- The straight line test can prove that a limit does not exist, but it can only suggest that a limit does exist.
- The Size Bounds can prove that a limit does exist, but they can only suggest that a limit does not exist.

The next proposition is a handy encoding of an intuitively plausible property of continuous mappings. The result is so natural that it often is tacitly taken for granted, but it is worth stating and proving carefully.

Proposition 2.3.9 (Persistence of Inequality). *Let A be a subset of \mathbf{R}^n and let $f : A \rightarrow \mathbf{R}^m$ be a continuous mapping. Let a be a point of A , let b be a point of \mathbf{R}^m , and suppose that $f(a) \neq b$. Then there exists some $\varepsilon > 0$ such that*

$$\text{for all } x \in A \text{ such that } |x - a| < \varepsilon, f(x) \neq b.$$

Proof. Assume that the displayed statement in the proposition fails for every $\varepsilon > 0$. Then in particular it fails for $\varepsilon = 1/\nu$ for $\nu = 1, 2, 3, \dots$. So there is a sequence $\{x_\nu\}$ in A such that

$$|x_\nu - a| < 1/\nu \quad \text{and} \quad f(x_\nu) = b, \quad \nu = 1, 2, 3, \dots$$

Since f is continuous at a , this condition shows that $f(a) = b$. But in fact $f(a) \neq b$, and so our assumption that the displayed statement in the proposition fails for every $\varepsilon > 0$ leads to a contradiction. Therefore the statement holds for some $\varepsilon > 0$, as desired. \square

Exercises

2.3.1. For $A \subset \mathbf{R}^n$, partially verify that $\mathcal{M}(A, \mathbf{R}^m)$ is a vector space over \mathbf{R} by showing that it satisfies vector space axioms (A4) and (D1).

2.3.2. Define multiplication $*$: $\mathcal{M}(A, \mathbf{R}) \times \mathcal{M}(A, \mathbf{R}) \rightarrow \mathcal{M}(A, \mathbf{R})$. Is $\mathcal{M}(A, \mathbf{R})$ a field with “+” from the section and this multiplication? Does it have a subspace that is a field?

2.3.3. For $A \subset \mathbf{R}^n$ and $m \in \mathbf{Z}^+$ define

$$\mathcal{C}(A, \mathbf{R}^m) = \{f \in \mathcal{M}(A, \mathbf{R}^m) : f \text{ is continuous on } A\}.$$

What facts about continuous functions are needed to prove that $\mathcal{C}(A, \mathbf{R}^m)$ is a vector space? Prove them.

2.3.4. Define an inner product and a modulus on $\mathcal{C}([0, 1], \mathbf{R})$ by

$$\langle f, g \rangle = \int_0^1 f(t)g(t)dt, \quad |f| = \sqrt{\langle f, f \rangle}.$$

How much of the material on inner product and modulus in \mathbf{R}^n carries over to $\mathcal{C}([0, 1], \mathbf{R})$? Express the Cauchy–Schwarz inequality as a relation between integrals.

2.3.5. Prove the componentwise nature of convergence.

2.3.6. Prove the componentwise nature of continuity.

2.3.7. Prove the persistence of continuity under composition.

2.3.8. Define $f : \mathbf{Q} \rightarrow \mathbf{R}$ by the rule

$$f(x) = \begin{cases} 1 & \text{if } x^2 < 2 \\ 0 & \text{if } x^2 > 2 \end{cases}$$

Is f continuous on \mathbf{Q} ?

2.3.9. Which of the following functions on \mathbf{R}^2 can be defined continuously at $\mathbf{0}$?

$$f(x, y) = \begin{cases} \frac{x^2 - y^2}{x^2 + y^2} & \text{if } (x, y) \neq \mathbf{0}, \\ b & \text{if } (x, y) = \mathbf{0}, \end{cases}$$

$$g(x, y) = \begin{cases} \frac{x^2 - y^3}{x^2 + y^2} & \text{if } (x, y) \neq \mathbf{0}, \\ b & \text{if } (x, y) = \mathbf{0}, \end{cases}$$

$$h(x, y) = \begin{cases} \frac{x^2 - y^2}{(x^2 + y^2)^{1/2}} & \text{if } (x, y) \neq \mathbf{0}, \\ b & \text{if } (x, y) = \mathbf{0}, \end{cases}$$

2.3.10. Let $k(x, y) = l(xy)$ where $l : \mathbf{R} \rightarrow \mathbf{R}$ is continuous. Is k continuous?

2.3.11. Let $f, g \in \mathcal{M}(\mathbf{R}^n, \mathbf{R})$ be such that $f + g$ and fg are continuous. Are f and g necessarily continuous?

2.4 Topology: Compact Sets and Continuity

The Extreme Value Theorem from one-variable calculus states: *Let I be a nonempty closed and bounded interval in \mathbf{R} , and let $f : I \rightarrow \mathbf{R}$ be a continuous function. Then f takes a minimum value and a maximum value on I .* This section generalizes the theorem from scalars to vectors. That is, we want a result that if A is a set in \mathbf{R}^n with certain properties, and if $f : A \rightarrow \mathbf{R}^m$ is a continuous mapping, then the output set $f(A)$ will also have certain properties. The questions are, for what sorts of properties do such statements hold, and when they hold, how do we prove them?

The one-variable theorem involves two pieces of data, the nonempty closed and bounded interval I and the continuous function f . Each of these is described in its own terms— I takes the form $[a, b]$ where $a \leq b$, while the continuity of f is an assertion about convergence of sequences. Because the two data have differently-phrased descriptions, a proof of the Extreme Value Theorem doesn't suggest itself immediately; no ideas at hand bear obviously on all the information. Thus the work of this section is not only to define the sets to appear in the pending theorem, but also to describe them compatibly with the sequential description of continuous mappings. The theorem itself will then be easy to prove. Accordingly, most of the section will consist of describing sets in two ways.

We begin with a little machinery to quantify the intuitive notion of nearness.

Definition 2.4.1 (ε -ball). *For any point $a \in \mathbf{R}^n$ and any positive real number $\varepsilon > 0$, the ε -ball centered at a is the set*

$$B(a, \varepsilon) = \{x \in \mathbf{R}^n : |x - a| < \varepsilon\}.$$

(See figure 2.11.)




Figure 2.11. Balls in various dimensions

With ε -balls it is easy to describe the points that are approached by a set A .

Definition 2.4.2 (Limit Point). *Let A be a subset of \mathbf{R}^n , and let a be a point of \mathbf{R}^n . The point a is a **limit point** of A if every ε -ball centered at a contains some point $x \in A$ such that $x \neq a$.*

A limit point of A need not belong to A (exercise 2.4.2). On the other hand, a point in A need not be a limit point of A (exercise 2.4.2 again); such a point is called an **isolated point** of A . Equivalently, a is an isolated point of A if $a \in A$ and there exists some $\varepsilon > 0$ such that $B(a, \varepsilon) \cap A = \{a\}$. The next lemma justifies the terminology of the previous definition: limit points of A are precisely the (nontrivial) limits of sequences in A .

Lemma 2.4.3 (Sequential Characterization of Limit Points). *Let A be a subset of \mathbf{R}^n , and let a be a point of \mathbf{R}^n . Then a is the limit of a sequence $\{x_\nu\}$ in A with each $x_\nu \neq a$ if and only if a is a limit point of A .*

Proof. (\implies) If a is the limit of sequence $\{x_\nu\}$ in A with each $x_\nu \neq a$ then any ε -ball about a contains an x_ν (in fact, infinitely many), so a is a limit point of A .

(\impliedby) Conversely, if a is a limit point of A then $B(a, 1/2)$ contains some $x_1 \in A$, $x_1 \neq a$. Let $\varepsilon_2 = |x_1 - a|/2$. The ball $B(a, \varepsilon_2)$ contains some $x_2 \in A$, $x_2 \neq a$, and since $|x_2 - a| < |x_1 - a|/2$, $x_2 \neq x_1$. Set $\varepsilon_3 = |x_2 - a|/2$ and continue defining a sequence $\{x_\nu\}$ in this fashion with $|x_\nu - a| < 1/2^\nu$ for all ν . This sequence converges to a and each $x_\nu \neq a$. \square

The lemma shows that Definition 2.4.2 is more powerful than it appears—every ε -ball centered at a limit point of A contains not only one but *infinitely many* points of A .

Definition 2.4.4 (Closed Set). *A subset A of \mathbf{R}^n is closed if it contains all of its limit points.*

For example, the x_1 -axis is closed as a subset of \mathbf{R}^n since any point off the axis is surrounded by a ball that misses the axis. The interval $(0, 1)$ is not closed because it does not contain the limit points at its ends. These examples illustrate the fact that with a little practice it becomes easy to recognize quickly whether a set is closed. Loosely speaking, a set is closed when it contains all the points that it seems to want to contain.

Proposition 2.4.5 (Sequential Characterization of Closed Sets). *Let A be a subset of \mathbf{R}^n . Then A is closed if and only if every sequence in A that converges in \mathbf{R}^n in fact converges in A .*

Proof. (\implies) Suppose that A is closed, and let $\{x_\nu\}$ be a sequence in A converging in \mathbf{R}^n to a . If $x_\nu = a$ for some ν then $a \in A$ since $x_\nu \in A$, and if $x_\nu \neq a$ for all ν then a is a limit point of A by “ \implies ” of Lemma 2.4.3, and so $a \in A$ since A is closed.

(\impliedby) Conversely, suppose that every convergent sequence in A has its limit in A . Then all limit points of A are in A by “ \impliedby ” of Lemma 2.4.3, and so A is closed. \square

The proposition equates an easily recognizable condition that we can understand intuitively (a set being closed) with a technical condition that we can use in further arguments (the sequential characterization).

Closed sets do not necessarily have good properties under continuous mappings. So next we describe another property of sets, boundedness. This is again an easily recognizable condition that also has a characterization in terms of sequences. The sequential characterization will turn out to be complementary to the sequential characterization of closed sets, foreshadowing that the properties of being closed and bounded will work well together.

Definition 2.4.6 (Bounded Set). *The set A in \mathbf{R}^n is bounded if $A \subset B(\mathbf{0}, R)$ for some $R > 0$.*

Thus a bounded set is enclosed in some finite “corral” centered at the origin, possibly a very big one. For example, any ball $B(p, \varepsilon)$, not necessarily centered at the origin, is bounded, by a nice application of the Triangle Inequality. On the other hand, the Archimedean property of the real number system says that \mathbf{Z} is an unbounded subset of \mathbf{R} . The Size Bounds show that any subset of \mathbf{R}^n is bounded if and only if the j th coordinates of its points form a bounded subset of \mathbf{R} for each $j \in \{1, \dots, n\}$. The geometric content of this statement is that a set sits inside a ball centered at the origin if and only if it sits inside a box centered at the origin.

Blurring the distinction between a sequence and the set of its elements allows the definition of boundedness to apply to sequences. That is, a sequence $\{x_\nu\}$ is bounded if there is some $R > 0$ such that $|x_\nu| < R$ for all $\nu \in \mathbf{Z}^+$. The proof of the next fact in \mathbf{R}^n is symbol-for-symbol the same as in \mathbf{R} (or in \mathbf{C}), so it is only sketched.

Proposition 2.4.7 (Convergence Implies Boundedness). *If the sequence $\{x_\nu\}$ converges in \mathbf{R}^n then it is bounded.*

Proof. Let $\{x_\nu\}$ converge to a . Then there exists a starting index ν_0 such that $x_\nu \in B(a, 1)$ for all $\nu > \nu_0$. Consider any real number R such that

$$R > \max\{|x_1|, \dots, |x_{\nu_0}|, |a| + 1\}.$$

Then clearly $x_\nu \in B(\mathbf{0}, R)$ for $\nu = 1, \dots, \nu_0$, and the Triangle Inequality shows that also $x_\nu \in B(\mathbf{0}, R)$ for all $\nu > \nu_0$. Thus $\{x_\nu\} \subset B(\mathbf{0}, R)$ as a set. \square

Definition 2.4.8 (Subsequence). *A subsequence of the sequence $\{x_\nu\}$ is a sequence consisting of some (possibly all) of the original terms, in ascending order of indices.*

Since a subsequence of $\{x_\nu\}$ consists of terms x_ν only for some values of ν , it is often written $\{x_{\nu_k}\}$, where now k is the index variable. For example, given the sequence

$$\{x_1, x_2, x_3, x_4, x_5, \dots\},$$

a subsequence is

$$\{x_2, x_3, x_5, x_7, x_{11}, \dots\},$$

with $\nu_1 = 2$, $\nu_2 = 3$, $\nu_3 = 5$, and generally $\nu_k =$ the k th prime.

Lemma 2.4.9 (Persistence of Convergence). *Let $\{x_\nu\}$ converge to a . Then any subsequence $\{x_{\nu_k}\}$ also converges to a .*

Proof. The k th term of the subsequence is x_{ν_k} . Given $\varepsilon > 0$ there exists ν_0 such that $k > \nu_0 \implies |x_k - a| < \varepsilon$. Since $\nu_k \geq k$ for all k , also $k > \nu_0 \implies |x_{\nu_k} - a| < \varepsilon$. Thus all terms of the subsequence past the ν_0 th are within ε of a . \square

The sequence property that characterizes bounded sets is called the **Bolzano–Weierstrass** property. Once it is proved in \mathbf{R} , the result follows in \mathbf{R}^n by arguing one component at a time.

Theorem 2.4.10 (Bolzano–Weierstrass Property in \mathbf{R}). *Let A be a bounded subset of \mathbf{R} . Then every sequence in A has a convergent subsequence.*

Proof. Let $\{x_\nu\}$ be a sequence in A . Call a term x_ν of the sequence a max-point if it is at least as big as all later terms, i.e., $x_\nu \geq x_\mu$ for all $\mu > \nu$. (For visual intuition, draw a graph plotting x_ν as a function of ν , with line segments connecting consecutive points. A max-point is a peak of the graph at least as high as all points to its right.) If there are infinitely many max-points in $\{x_\nu\}$ then these form a decreasing sequence. If there are only finitely many max-points then $\{x_\nu\}$ has an increasing sequence starting after the last max-point—this follows almost immediately from the definition of max-point. In either case, $\{x_\nu\}$ has a monotonic subsequence which, being bounded, converges because the real number system is complete. \square

Theorem 2.4.11 (Bolzano–Weierstrass Property in \mathbf{R}^n : Sequential Characterization of Bounded Sets). *Let A be a subset of \mathbf{R}^n . Then A is bounded if and only if every sequence in A has a subsequence that converges in \mathbf{R}^n .*

Proof. Suppose that A is bounded. Consider any sequence in A , written as $\{(x_{1,\nu}, \dots, x_{n,\nu})\}$. The real sequence $\{x_{1,\nu}\}$ takes values in a bounded subset of \mathbf{R} and thus has a convergent subsequence, $\{x_{1,\nu_k}\}$. The subscripts are getting out of hand, so keep only the ν_k th terms of the original sequence and relabel it. In other words, we may as well assume that the sequence of first components, $\{x_{1,\nu}\}$, converges. The real sequence of second components, $\{x_{2,\nu}\}$, in turn has a convergent subsequence, and by Lemma 2.4.9 the corresponding subsequence of first components, $\{x_{1,\nu}\}$, converges too. Relabeling again, we may assume that $\{x_{1,\nu}\}$ and $\{x_{2,\nu}\}$ both converge. Continuing in this fashion exhibits a subsequence that converges at each component.

Conversely, suppose that A is not bounded. Then there is a sequence $\{x_\nu\}$ in A with $|x_\nu| > \nu$ for all ν . This has no bounded subsequence, and hence it has no convergent subsequence by Proposition 2.4.7. \square

Note how the sequential characterizations in Proposition 2.4.5 and in the Bolzano–Weierstrass Property complement each other. The proposition characterizes any closed set in \mathbf{R}^n by the fact that if a sequence converges in the ambient space then it converges in the set. The Bolzano–Weierstrass Property characterizes any bounded set in \mathbf{R}^n by the fact that every sequence in the set has a subsequence that converges, but not necessarily in the set.

Definition 2.4.12 (Compact Set). *A subset K of \mathbf{R}^n is compact if it is closed and bounded.*

Since the notions of closed and bounded are reasonably intuitive, we can usually recognize compact sets on sight. But it is not obvious from how compact sets look that they are particularly useful objects in relation to continuity. So the program now has two steps: first, combine Proposition 2.4.5 and the Bolzano–Weierstrass property to characterize compact sets in terms of sequences, and second, use the characterization to prove that compactness is preserved by continuous mappings.

Theorem 2.4.13 (Sequential Characterization of Compact Sets). *Let K be a subset of \mathbf{R}^n . Then K is compact if and only if every sequence in K has a subsequence that converges in K .*

Proof. (\implies) Suppose that K is compact and $\{x_\nu\}$ is a sequence in K . Then K is bounded, so by “ \implies ” of the Bolzano–Weierstrass property $\{x_\nu\}$ has a convergent subsequence. But K is also closed, so by “ \implies ” of Proposition 2.4.5, this subsequence converges in K .

(\impliedby) On the other hand, suppose that every sequence in K has a subsequence that converges in K . Then in particular, every sequence in K that converges in \mathbf{R}^n has a subsequence that converges in K . By Lemma 2.4.9 the limit of the sequence is the limit of the subsequence, so the sequence converges in K . That is, every sequence in K that converges in \mathbf{R}^n converges in K . Thus K is closed by “ \impliedby ” of Proposition 2.4.5. Another consequence of the assumption at the beginning of this paragraph is that every sequence in K has a subsequence that converges in \mathbf{R}^n . Thus K is bounded by “ \impliedby ” of the Bolzano–Weierstrass Property. \square

The next theorem is the main result of this section. Now that all of the objects involved are described in the common language of sequences, its proof is natural.

Theorem 2.4.14 (The Continuous Image of a Compact Set is Compact). *Let K be a compact subset of \mathbf{R}^n and let $f : K \rightarrow \mathbf{R}^m$ be continuous. Then $f(K)$, the image set of K under f , is a compact subset of \mathbf{R}^m .*

Proof. Let $\{y_\nu\}$ be any sequence in $f(K)$; by (\impliedby) of Theorem 2.4.13, it suffices to exhibit a subsequence converging in $f(K)$. Each y_ν has the form $f(x_\nu)$, and this defines a sequence $\{x_\nu\}$ in K . By (\implies) of Theorem 2.4.13,

since K is compact, $\{x_\nu\}$ necessarily has a subsequence $\{x_{\nu_k}\}$ converging in K , say to x . By the continuity of f at x , the sequence $\{f(x_{\nu_k})\}$ converges in $f(K)$ to $f(x)$. Since $\{f(x_{\nu_k})\}$ is a subsequence of $\{y_\nu\}$, the proof is complete. \square

Again, the sets in Theorem 2.4.14 are defined with no direct reference to sequences, but the theorem is proved entirely by using sequences. The point is that with the theorem proved, we can easily see that it applies in particular contexts without having to think any more about the sequences that were used to prove it.

A corollary of Theorem 2.4.14 generalizes the theorem that was quoted to begin the section:

Theorem 2.4.15 (Extreme Value Theorem). *Let K be a nonempty compact subset of \mathbf{R}^n and let the function $f : K \rightarrow \mathbf{R}$ be continuous. Then f takes a minimum and a maximum value on K .*

Proof. By Theorem 2.4.14, $f(K)$ is a compact subset of \mathbf{R} . As a nonempty bounded subset of \mathbf{R} , $f(K)$ has a greatest lower bound and a least upper bound by the completeness of the real number system. Each of these bounds is an isolated point or a limit point of $f(K)$, since otherwise some ε -ball about it would be disjoint from $f(K)$, giving rise to greater lower bounds or lesser upper bounds of $f(K)$. Since $f(K)$ is also closed it contains its limit points, so in particular it contains its greatest lower bound and its least upper bound. This means precisely that f takes a minimum and a maximum value on K . \square

Even when $n = 1$, Theorem 2.4.15 generalizes the Extreme Value Theorem from the beginning of the section—in the theorem here, K can be a finite union of closed and bounded intervals in \mathbf{R} rather than only one interval.

A *topological property* of sets is a property that is preserved under continuity. Theorem 2.4.14 says that compactness is a topological property. Neither the property of being closed nor of being bounded is in itself topological. That is, the continuous image of a closed set need not be closed, and the continuous image of a bounded set need not be bounded; for that matter, the continuous image of a closed set need not be bounded, and the continuous image of a bounded set need not be closed (exercise 2.4.7).

Actually, the nomenclature *continuous image* in the slogan-title of Theorem 2.4.14 and in the previous paragraph is inaccurate: the image of a mapping is a set, and the notion of a set being continuous doesn't even make sense according to our grammar. As stated correctly in the body of the theorem, *continuous image* is short for *image under a continuous mapping*. The property that students often have in mind when they call a set continuous is in fact called *connectedness*. Loosely, a set is connected if it has only one piece, so that a better approximating word from everyday language is *contiguous*. Connectedness is rather technical to define carefully, and so we omit it since it is not needed in this course. The remark after Theorem 2.4.15 points out

that connectedness is not even needed for the one-variable Extreme Value Theorem. However, it deserves passing mention that connectedness is also a topological property: again using language loosely, the continuous image of a connected set is connected. This generalizes another theorem that underlies one-variable calculus, the Intermediate Value Theorem.

The ideas of this section readily extend to broader environments. The first generalization of Euclidean space is a *metric space*, a set with a well-behaved distance function. Even more general is a *topological space*, a set with some of its subsets designated as closed. Continuous functions and compact sets can be defined meaningfully in these environments, and the theorems remain the same: the continuous image of a compact set is compact, and the continuous image of a connected set is connected.

Exercises

2.4.1. Are the following subsets of \mathbf{R}^n closed, bounded, compact?

- (a) $B(\mathbf{0}, 1)$,
- (b) $\{(x, y) \in \mathbf{R}^2 : y - x^2 = 0\}$,
- (c) $\{(x, y, z) \in \mathbf{R}^3 : x^2 + y^2 + z^2 - 1 = 0\}$,
- (d) $\{x : f(x) = \mathbf{0}_m\}$ where $f \in \mathcal{M}(\mathbf{R}^n, \mathbf{R}^m)$ is continuous (this generalizes (b) and (c)),
- (e) \mathbf{Q}^n where \mathbf{Q} denotes the rational numbers,
- (f) $\{(x_1, \dots, x_n) : x_1 + \dots + x_n > 0\}$.

2.4.2. Give a set $A \subset \mathbf{R}^n$ and limit point b of A such that $b \notin A$. Give a set $A \subset \mathbf{R}^n$ and a point $a \in A$ such that a is not a limit point of A .

2.4.3. Let A be a closed subset of \mathbf{R}^n and let $f \in \mathcal{M}(A, \mathbf{R}^m)$. Define the **graph** of f to be

$$G(f) = \{(a, f(a)) : a \in A\},$$

a subset of \mathbf{R}^{n+m} . Show that if f is continuous then its graph is closed.

2.4.4. Prove the closed set properties: (1) The empty set \emptyset and the full space \mathbf{R}^n are closed subsets of \mathbf{R}^n , (2) any intersection of closed sets is closed, (3) any finite union of closed sets is closed.

2.4.5. Prove that any ball $B(p, \varepsilon)$ is bounded in \mathbf{R}^n .

2.4.6. Show that A is a bounded subset of \mathbf{R}^n if and only if for each $j \in \{1, \dots, n\}$, the j th coordinates of its points form a bounded subset of \mathbf{R} .

2.4.7. Show by example that the continuous image of a closed set need not be closed, that the continuous image of a closed set need not be bounded, that the continuous image of a bounded set need not be closed, and that the continuous image of a bounded set need not be bounded.

2.4.8. A subset A of \mathbf{R}^n is called **discrete** if each of its points is isolated. (Recall that the term “isolated” was defined in the section.) Show or take for granted the (perhaps surprising at first) fact that every mapping whose domain is discrete must be continuous. Is discreteness a topological property? That is, need the continuous image of a discrete set be discrete?

2.4.9. A subset A of \mathbf{R}^n is called **path-connected** if for any two points $x, y \in A$, there is a continuous mapping

$$\gamma : [0, 1] \longrightarrow A$$

such that $\gamma(0) = x$ and $\gamma(1) = y$. (This γ is the path that connects x and y .) Draw a picture to illustrate the definition of a path-connected set. Prove that path-connectedness is a topological property.

2.5 Review of the One-Variable Derivative

The limit of a sequence was introduced in Definition 2.3.1. The limit of a mapping will now be defined as the common limit of all suitable sequences, if this common sequence limit exists. Recall from Definition 2.4.2 that a point a is a limit point of a set A if every ε -ball centered at a contains some point $x \in A$ such that $x \neq a$. A limit point of A may or may not be a point of A . Also recall from Lemma 2.4.3 that a point a is a limit point of a set A if and only if a is the limit of a sequence $\{x_\nu\}$ in A with each $x_\nu \neq a$.

Definition 2.5.1 (Function Limit). Let A be a subset of \mathbf{R}^n , let $f : A \longrightarrow \mathbf{R}^m$ be a mapping, and let a be a limit point of A . Let ℓ be a point of \mathbf{R}^m . Then f has limit ℓ as x approaches a , written

$$\lim_{x \rightarrow a} f(x) = \ell,$$

if for every sequence $\{x_\nu\}$ in A with each $x_\nu \neq a$ such that $\{x_\nu\}$ converges to a , the corresponding output sequence $\{f(x_\nu)\}$ converges to ℓ .

Thus the notion of $\lim_{x \rightarrow a} f(x)$ makes no reference to $f(a)$ (which may not even be defined), but only to values $f(x)$ for x near a .

The Sum Rule and the Constant Multiple Rule for sequence limits give rise to the same rules for mapping limits as well, but there is one technical issue. The Sum Rule seems obvious,

$$\lim_{x \rightarrow a} (f(x) + g(x)) = \lim_{x \rightarrow a} f(x) + \lim_{x \rightarrow a} g(x),$$

where $f : A \longrightarrow \mathbf{R}^m$ and a is a limit point of A , and $g : B \longrightarrow \mathbf{R}^m$ and a is a limit point of B . But one needs to observe that the domain of $f + g$ is $A \cap B$, and so the limit on the left can not exist unless the limit point a of A and

of B is also a limit point of the smaller set $A \cap B$. For example, the functions \sqrt{x} and $\sqrt{-x}$ have the respective domains $[0, \infty)$ and $(-\infty, 0]$, and

$$\lim_{x \rightarrow 0} \sqrt{x} = \lim_{x \rightarrow 0} \sqrt{-x} = 0,$$

but the function $\sqrt{x} + \sqrt{-x}$ has the domain $[0, \infty) \cap (-\infty, 0] = \{0\}$, and 0 is not a limit point of this set, so therefore

$$\lim_{x \rightarrow 0} (\sqrt{x} + \sqrt{-x}) \text{ does not exist.}$$

This should be clear in light of the sentence immediately after Definition 2.5.1. Returning to the general Sum Rule for mappings, other than this additional detail to check, it follows from its counterpart for sequences. The Constant Multiple Rule for mappings follows from its counterpart for sequences without any additional technical considerations, since any constant multiple of a mapping has the same domain as the original mapping.

Let $A \subset \mathbf{R}^n$ be a set and let $a \in \mathbf{R}^n$ be a point. A mapping $f : A \rightarrow \mathbf{R}^m$ is **null at a** if $\lim_{x \rightarrow a} f(x) = \mathbf{0}_m$. Thus if f is null at a then a must be a limit point of A . Formulating the Sum Rule and the Constant Multiple Rule for null mappings is left to you (exercise 2.5.1).

The notions of limit and continuity are closely related for mappings, but again with a small technical issue present. The proof of the following proposition is exercise 2.5.2.

Proposition 2.5.2 (Continuity in Terms of Function Limits). *Let A be a subset of \mathbf{R}^n , and let a be a point of A , and let $f : A \rightarrow \mathbf{R}^m$ be a mapping.*

Suppose that a is a limit point of A . Then f is continuous at a if and only if $\lim_{x \rightarrow a} f(x)$ exists and is equal to $f(a)$.

Suppose that a is not a limit point of A . Then f is continuous at a .

A careful discussion of the derivative is surprisingly technical even for functions of one variable. The one-variable derivative is defined as a limit of a difference quotient function. Usually the underlying assumptions, which can easily get lost, are that f is a function from some interval $I \subset \mathbf{R}$ to \mathbf{R} and that a is a point of I but not an endpoint. (Some authors allow differentiation at endpoints, but then the derivative can exist and be nonzero at an extremum of the function.) The difference quotient function is defined at all points except 0 of the interval J obtained by translating I by $-a$, moving a to 0,

$$g : J - \{0\} \rightarrow \mathbf{R}, \quad g(h) = \frac{f(a+h) - f(a)}{h}.$$

Thus 0 is a limit point of the domain of g (though not a point of the domain of g), so that according to Definition 2.5.1, $\lim_{h \rightarrow 0} g(h)$ might exist. When it does, the derivative of f at a is this function limit,

$$f'(a) = \lim_{h \rightarrow 0} g(h) = \lim_{h \rightarrow 0} \frac{f(a+h) - f(a)}{h}.$$

In sum, the derivative of f at a is

- a limit of a different function g , the difference quotient function whose domain is obtained by translating and puncturing the domain of f ,
- the limit being taken at the limit point 0 of the domain of g , which is not in the domain of g ,
- and the function limit being defined as the common value of output-sequence limits over all input-sequences that approach but do not reach 0, if this common value of output-sequence limits exists.

If you found the definition of the derivative in Mathematics 111 or 112 difficult to digest, it may be because all of these ideas had to be covered under intense time pressure in the midst of everything else that was happening in those courses, and because the very process of getting all the ideas into play necessarily rendered their presentation diffuse.

However, the author of these notes does not know any useful way to simplify the setup without waving his hands. One can study an alternate difference quotient function $g(x) = (f(x) - f(a))/(x - a)$ instead and thus avoid translating the domain of f to place the puncture-point at 0, but this is not not a good idea: in the definition of multivariable derivative to be introduced in chapter 4, translating the situation to the origin will clarify rather than complicate it. Also, one can define the limit of a function without reference to sequence-limits: this is the so-called “epsilon–delta” definition rather than our “epsilon– ν .” For example, the formulation of the completeness of the real number system as a set-bound criterion in Theorem 1.1.5 makes no reference to sequences, and if continuity of mappings is defined in epsilon–delta language then the Persistence of Inequality principle, which was a small nuisance to prove, becomes true by definition. However, eschewing sequences and basing all of the ideas in play here on an epsilon–delta formulation of limit makes other parts of the material harder. In particular, proving that compactness is a topological property without using the sequential characterization of compact sets requires considerable subtlety.

Exercises

2.5.1. Carefully state and prove the Sum Rule and the Constant Multiple Rule for mappings and then for null mappings.

2.5.2. Prove Proposition 2.5.2.

2.6 Summary

Along with introducing Euclidean space and its properties, this chapter is meant to provide a quick review of some ideas from one-variable calculus while generalizing them to higher dimension. This chapter has also emphasized working with vectors intrinsically rather than using coordinates. The multivariable Extreme Value Theorem will play a crucial role in our proof of the Inverse Function Theorem in chapter 5.

Linear Mappings and Their Matrices

The basic idea of differential calculus is to approximate smooth-but-curved objects in the small by straight ones. To prepare for doing so, this chapter studies the multivariable analogues of lines. With one variable, lines are easily manipulated by explicit formulas (e.g., the point-slope form is $y = mx + b$), but with many variables we want to use the language of mappings. Section 3.1 gives an algebraic description of “straight” mappings, the *linear* mappings, proceeding from an intrinsic definition to a description in coordinates. Each linear mapping is described by a box of numbers called a *matrix*, so section 3.2 derives mechanical matrix manipulations corresponding to the natural ideas of adding, scaling, and composing linear mappings. Section 3.3 discusses in matrix terms the question of whether a linear mapping has an *inverse*, i.e., whether there is a second linear mapping such that each undoes the other’s effect. Section 3.5 discusses the *determinant*, an elaborate matrix-to-scalar function that extracts from a linear mapping a single number with remarkable properties:

- (Linear Invertibility Theorem) The mapping is invertible if and only if the determinant is nonzero.
- An explicit formula for the inverse of an invertible linear mapping can be written using the determinant (section 3.7).
- The factor by which the mapping magnifies volume is the absolute value of the determinant (section 3.8).
- The mapping preserves or reverses orientation according to the sign of the determinant (section 3.9). Here orientation is an algebraic generalization of clockwise versus counterclockwise in the plane and of right-handed versus left-handed in space.

Finally, section 3.10 defines the cross product (a vector-by-vector multiplication special to three dimensions) and uses it to derive formulas for lines and planes in space.

3.1 Linear Mappings

The simplest interesting mappings from \mathbf{R}^n to \mathbf{R}^m are those whose output is proportional to their input, the *linear mappings*. Proportionality means that a linear mapping should take a sum of inputs to the corresponding sum of outputs,

$$T(x + y) = T(x) + T(y) \quad \text{for all } x, y \in \mathbf{R}^n, \quad (3.1)$$

and a linear mapping should take a scaled input to the correspondingly scaled output,

$$T(\alpha x) = \alpha T(x) \quad \text{for all } \alpha \in \mathbf{R}, x \in \mathbf{R}^n. \quad (3.2)$$

(Here we use the symbol α because a will be used heavily in other ways during this chapter.) More formally,

Definition 3.1.1 (Linear Mapping). *The mapping $T : \mathbf{R}^n \rightarrow \mathbf{R}^m$ is linear if*

$$T\left(\sum_{i=1}^k \alpha_i x_i\right) = \sum_{i=1}^k \alpha_i T(x_i)$$

for all positive integers k , all real numbers α_1 through α_k , and all vectors x_1 through x_k .

The reader may find this definition discomfiting. It does not say what form a linear mapping takes, and this raises some immediate questions. How are we to recognize linear mappings when we encounter them? Or are we supposed to think about them without knowing what they look like? For that matter, are there even any linear mappings to encounter? Another troublesome aspect of Definition 3.1.1 is semantic: despite the geometric sound of the word “linear,” the definition is in fact algebraic, describing how T behaves with respect to the algebraic operations of vector addition and scalar multiplication. (Note that on the left of the equality in the definition, the operations are set in \mathbf{R}^n , while on the right they are in \mathbf{R}^m .) So what is the connection between the definition and actual lines? Finally, how exactly do conditions (3.1) and (3.2) relate to the condition in the definition?

On the other hand, Definition 3.1.1 has the virtue of illustrating the principle that *to do mathematics effectively we should characterize our objects rather than construct them*. The characterizations are admittedly guided by hindsight, but there is nothing wrong with that. Definition 3.1.1 says how a linear mapping behaves. It says that whatever form linear mappings will turn out to take, our reflex should be to think of them as mappings through which we can pass sums and constants. The definition tells us how to use linear mappings once we know what they are. Another virtue of Definition 3.1.1 is that it is intrinsic, making no reference to coordinates.

Some of the questions raised by Definition 3.1.1 have quick answers. The connection between the definition and actual lines will quickly emerge from our pending investigations. Also, an induction argument shows that (3.1) and (3.2)

are equivalent to the characterization in the definition, despite appearing weaker (exercise 3.1.1). Thus, to verify that a mapping is linear, we only need to show that it satisfies the easier-to-check conditions (3.1) and (3.2); but to derive properties of mappings that are known to be linear, we may want to use the more powerful condition in the definition. As for finding linear mappings, the definition suggests a two-step strategy: first, derive the form that a linear mapping necessarily takes in consequence of satisfying the definition; and second, verify that the mappings of that form are indeed linear, i.e., show that the necessary form of a linear mapping is also sufficient for a mapping to be linear. We now turn to this.

The easiest case to study is linear mappings from \mathbf{R} to \mathbf{R} . Following the program, first we assume that we have such a mapping and determine its form, obtaining the mappings that are candidates to be linear. Second we show that all the candidates are indeed linear mappings. Thus suppose that some mapping $T : \mathbf{R} \rightarrow \mathbf{R}$ is linear. The map determines a scalar, $a = T(1)$. And then for any $x \in \mathbf{R}$,

$$\begin{aligned} T(x) &= T(x \cdot 1) && \text{since } x \cdot 1 = x \\ &= xT(1) && \text{by (3.1)} \\ &= xa && \text{by definition of } a \\ &= ax && \text{since multiplication in } \mathbf{R} \text{ commutes.} \end{aligned}$$

Thus, T is simply multiplication by a , where $a = T(1)$. But to reiterate, this calculation does not show that any mapping is linear, it shows only what form a mapping must necessarily have once it is already known to be linear. We don't yet know that any linear mappings exist at all.

So the next thing to do is show that conversely any mapping of the derived form is indeed linear—the necessary condition is also sufficient. Fix a real number a and define a mapping $T : \mathbf{R} \rightarrow \mathbf{R}$ by $T(x) = ax$. Then the claim is that T is linear and $T(1) = a$. Let's partially show this by verifying that T satisfies (3.2). For any $\alpha \in \mathbf{R}$ and any $x \in \mathbf{R}$,

$$\begin{aligned} T(\alpha x) &= a\alpha x && \text{by definition of } T \\ &= \alpha ax && \text{since multiplication in } \mathbf{R} \text{ commutes} \\ &= \alpha T(x) && \text{by definition of } T, \end{aligned}$$

as needed. You can check (3.1) similarly, and the calculation that $T(1) = a$ is immediate. These last two paragraphs combine to show

Proposition 3.1.2 (Description of Linear Mappings from Scalars to Scalars). *The linear mappings $T : \mathbf{R} \rightarrow \mathbf{R}$ are precisely the mappings*

$$T(x) = ax$$

where $a \in \mathbf{R}$. That is, each linear mapping $T : \mathbf{R} \rightarrow \mathbf{R}$ is multiplication by a unique $a \in \mathbf{R}$ and conversely.

This explains the term “linear”: the graphs of linear mappings from \mathbf{R} to \mathbf{R} are lines through the origin. (Mappings $f(x) = ax + b$ with $b \neq 0$ are not linear according to our definition even though their graphs are also lines. However, see exercises 3.1.14 and 3.2.5.) For example, a typical linear mapping from \mathbf{R} to \mathbf{R} is $T(x) = (1/2)x$. Figure 3.1 shows two ways of visualizing this mapping. The left half of the figure plots the domain axis and the codomain axis in one plane, orthogonally to each other, the familiar way to graph a function. The right half of the figure plots the axes separately, using the spacing of the dots to describe the mapping instead. The uniform spacing along the rightmost axis depicts the fact that $T(x) = xT(1)$ for all $x \in \mathbf{Z}$, and the spacing is half as big because the multiplying factor is $1/2$. Figures of this second sort can generalize up to three dimensions of input and three dimensions of output, whereas figures of the first sort can display at most three dimensions of input and output combined.




Figure 3.1. A linear mapping from \mathbf{R} to \mathbf{R}

Next consider a linear mapping $T : \mathbf{R}^n \rightarrow \mathbf{R}$. Recall the standard basis of \mathbf{R}^n ,

$$\{e_1, \dots, e_n\} = \{(1, 0, \dots, 0), \dots, (0, 0, \dots, 1)\}.$$

Take the n real numbers

$$a_1 = T(e_1), \quad \dots, \quad a_n = T(e_n),$$

and define the vector $a = (a_1, \dots, a_n) \in \mathbf{R}^n$. Any $x \in \mathbf{R}^n$ can be written

$$x = (x_1, \dots, x_n) = \sum_{i=1}^n x_i e_i, \quad \text{each } x_i \in \mathbf{R}.$$

(So here each x_i is a scalar entry of the vector x , whereas in Definition 3.1.1, each x_i was itself a vector. The author does not know any graceful way to avoid this notation collision, the systematic use of boldface or arrows to adorn vector names being heavyhanded, and the systematic use of the Greek letter ξ rather than its Roman counterpart x to denote scalars being alien. Since mathematics involves finitely many symbols and infinitely many ideas, the reader will in any case eventually need the skill of discerning meaning from context, a skill that may as well start receiving practice now.) Returning to

the main discussion, since $x = \sum_{i=1}^n x_i e_i$ and T is linear, Definition 3.1.1 shows that

$$T(x) = T\left(\sum_{i=1}^n x_i e_i\right) = \sum_{i=1}^n x_i T(e_i) = \sum_{i=1}^n x_i a_i = \langle x, a \rangle = \langle a, x \rangle.$$

Again, the only possibility for the linear mapping is multiplication (but this time in the sense of inner product) by an element a , where now $a = (T(e_1), \dots, T(e_n))$, but we don't yet know that such a mapping is linear. However, fix a vector $a = (a_1, \dots, a_n)$ and define the corresponding mapping $T : \mathbf{R}^n \rightarrow \mathbf{R}$ by $T(x) = \langle a, x \rangle$. Then it is straightforward to show that indeed T is linear and $T(e_j) = a_j$ for $j = 1, \dots, n$ (exercise 3.1.3). This shows

Proposition 3.1.3 (Description of Linear Mappings from Vectors to Scalars). *The linear mappings $T : \mathbf{R}^n \rightarrow \mathbf{R}$ are precisely the mappings*

$$T(x) = \langle a, x \rangle$$

where $a \in \mathbf{R}^n$. That is, each linear mapping $T : \mathbf{R}^n \rightarrow \mathbf{R}$ is multiplication by a unique $a \in \mathbf{R}^n$ and conversely.

In light of this proposition, you should be able to recognize linear mappings from \mathbf{R}^n to \mathbf{R} on sight. For example, the mapping $T : \mathbf{R}^3 \rightarrow \mathbf{R}$ given by $T(x, y, z) = \pi x + ey + \sqrt{2}z$ is linear, being multiplication by the vector $(\pi, e, \sqrt{2})$.

In the previous chapter, the second example after Definition 2.3.5 showed that every linear mapping $T : \mathbf{R}^n \rightarrow \mathbf{R}$ is continuous. You are encouraged to reread that example now before continuing. (Warning: the fixed vector a here is denoted v in the example, since in the context of the example the symbol a is reserved for another purpose.)

A depiction of a linear mapping from \mathbf{R}^2 to \mathbf{R} can again plot the domain plane and the codomain axis orthogonally to each other or separately. See figures 3.2 and 3.3 for examples of each type of plot. The first figure suggests that the graph forms a plane in \mathbf{R}^3 and that a line of inputs is taken to the output value 0. The second figure shows more clearly how the mapping compresses the plane into the line. As in the right half of figure 3.1, the idea is that $T(x, y) = xT(1, 0) + yT(0, 1)$ for all $x, y \in \mathbf{Z}$. The compression is that although $(1, 0)$ and $(0, 1)$ lie on separate input axes, $T(1, 0)$ and $T(0, 1)$ lie on the same output axis.

The most general mapping is $T : \mathbf{R}^n \rightarrow \mathbf{R}^m$. This decomposes as $T = (T_1, \dots, T_m)$ where each $T_i : \mathbf{R}^n \rightarrow \mathbf{R}$ is the i th component function of T . The next proposition reduces the linearity of such T to the linearity of its components T_i , which we already understand.

Proposition 3.1.4 (Componentwise Nature of Linearity). *The vector-valued mapping $T = (T_1, \dots, T_m) : \mathbf{R}^n \rightarrow \mathbf{R}^m$ is linear if and only if each scalar-valued component function $T_i : \mathbf{R}^n \rightarrow \mathbf{R}$ is linear.*




Figure 3.2. The graph of a linear mapping from \mathbf{R}^2 to \mathbf{R}




Figure 3.3. Second depiction of a linear mapping from \mathbf{R}^2 to \mathbf{R}

Proof. For any $x, y \in \mathbf{R}^n$,

$$T(x + y) = (T_1(x + y), \dots, T_m(x + y))$$

and

$$\begin{aligned} T(x) + T(y) &= (T_1(x), \dots, T_m(x)) + (T_1(y), \dots, T_m(y)) \\ &= (T_1(x) + T_1(y), \dots, T_m(x) + T_m(y)). \end{aligned}$$

But T satisfies (3.1) exactly when the left sides are equal, the left sides are equal exactly when the right sides are equal, and the right sides are equal exactly when each T_i satisfies (3.1). A similar argument with (3.2), left as exercise 3.1.5, completes the proof. \square

The componentwise nature of linearity combines with the fact that scalar-valued linear mappings are continuous (as observed after Proposition 3.1.3) and with the componentwise nature of continuity to show that all linear mappings are continuous. Despite being so easy to prove, this fact deserves a prominent statement.

Theorem 3.1.5 (Linear Mappings are Continuous). *Let the mapping $T : \mathbf{R}^n \rightarrow \mathbf{R}^m$ be linear. Then T is continuous.*

By the previous proposition, a mapping $T : \mathbf{R}^n \rightarrow \mathbf{R}^m$ is linear if and only if each T_i determines n real numbers a_{i1}, \dots, a_{in} . Putting all mn numbers a_{ij} into a box with m rows and n columns gives a **matrix**

$$A = \begin{bmatrix} a_{11} & a_{12} & \cdots & a_{1n} \\ a_{21} & a_{22} & \cdots & a_{2n} \\ \vdots & \vdots & & \vdots \\ a_{m1} & a_{m2} & \cdots & a_{mn} \end{bmatrix} \tag{3.3}$$

whose i th row is the vector determined by T_i , and whose (i, j) th entry (this means i th row, j th column) is thus given by

$$a_{ij} = T_i(e_j). \tag{3.4}$$

Sometimes one saves writing by abbreviating the right side of (3.3) to $[a_{ij}]_{m \times n}$, or even just $[a_{ij}]$ when m and n are firmly established.

The set of all m -by- n matrices (those with m rows and n columns) of real numbers is denoted $M_{m,n}(\mathbf{R})$. The n -by- n square matrices are denoted $M_n(\mathbf{R})$. Euclidean space \mathbf{R}^n is often identified with $M_{n,1}(\mathbf{R})$ and vectors written as columns,

$$(x_1, \dots, x_n) = \begin{bmatrix} x_1 \\ \vdots \\ x_n \end{bmatrix}.$$

This typographical convention may look odd, but it is useful. The idea is that a vector in parentheses is merely an ordered list of entries, not inherently a row or a column; but when a vector—or, more generally, a matrix—is enclosed by square brackets, the distinction between rows and columns is significant.

To make the linear mapping $T : \mathbf{R}^n \rightarrow \mathbf{R}^m$ be multiplication by its matrix $A \in M_{m,n}(\mathbf{R})$, we need to *define* multiplication of an m -by- n matrix A by an n -by-1 vector x appropriately. The result must be an m -by-1 vector whose i th entry is the inner product of A 's i th row and x . Thus,

Definition 3.1.6 (Matrix-by-Vector Multiplication). *Let $A \in M_{m,n}(\mathbf{R})$ and let $x \in \mathbf{R}^n$. Then the product $Ax \in \mathbf{R}^m$ is defined as*

$$Ax = \begin{bmatrix} a_{11} & a_{12} & \cdots & \cdots & a_{1n} \\ a_{21} & a_{22} & \cdots & \cdots & a_{2n} \\ \vdots & \vdots & & & \vdots \\ a_{m1} & a_{m2} & \cdots & \cdots & a_{mn} \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \\ \vdots \\ \vdots \\ x_n \end{bmatrix} = \begin{bmatrix} a_{11}x_1 + \cdots + a_{1n}x_n \\ a_{21}x_1 + \cdots + a_{2n}x_n \\ \vdots \\ a_{m1}x_1 + \cdots + a_{mn}x_n \end{bmatrix}.$$

In summation form this says,

$$A \left(\sum_{j=1}^n x_j e_j \right) = \sum_{i=1}^m \left(\sum_{k=1}^n a_{ik} x_k \right) e_i.$$

For example,

$$\begin{bmatrix} 1 & 2 & 3 \\ 4 & 5 & 6 \end{bmatrix} \begin{bmatrix} 7 \\ 8 \\ 9 \end{bmatrix} = \begin{bmatrix} 1 \cdot 7 + 2 \cdot 8 + 3 \cdot 9 \\ 4 \cdot 7 + 5 \cdot 8 + 6 \cdot 9 \end{bmatrix} = \begin{bmatrix} 50 \\ 122 \end{bmatrix}.$$

Definition 3.1.6 gives the following theorem, which encompasses Propositions 3.1.2 and 3.1.3 as special cases.

Theorem 3.1.7 (Description of Linear Mappings from Vectors to Vectors). *The linear mappings $T : \mathbf{R}^n \rightarrow \mathbf{R}^m$ are precisely the mappings*

$$T(x) = Ax$$

where $A \in M_{m,n}(\mathbf{R})$. That is, each linear mapping $T : \mathbf{R}^n \rightarrow \mathbf{R}^m$ is multiplication by a unique $A \in M_{m,n}(\mathbf{R})$ and conversely.

The columns of a matrix A , like the rows, have a description in terms of the corresponding mapping T . To see this, observe that the j th column is

$$\begin{bmatrix} a_{1j} \\ \vdots \\ a_{mj} \end{bmatrix} = \begin{bmatrix} T_1(e_j) \\ \vdots \\ T_m(e_j) \end{bmatrix} = T(e_j).$$

Thus, the i th row of A describes the i th component function T_i of T and the j th column of A gives the value of T on the j th standard basis vector e_j .

For an example of using this last principle, consider the mapping $r : \mathbf{R}^2 \rightarrow \mathbf{R}^2$ given by rotating the plane counterclockwise through an angle of $\pi/6$. It is geometrically evident that r is linear: rotating the parallelogram P with sides x_1 and x_2 (and thus with diagonal $x_1 + x_2$) by $\pi/6$ yields the parallelogram $r(P)$ with sides $r(x_1)$ and $r(x_2)$, so the diagonal of $r(P)$ is equal to both $r(x_1 + x_2)$ and $r(x_1) + r(x_2)$. This shows that r satisfies (3.1). The geometric verification of (3.2) is similar. (See figure 3.4.)




Figure 3.4. The rotation mapping is linear

To find the matrix of r , simply note that

$$r(e_1) = r(1, 0) = \begin{bmatrix} \sqrt{3}/2 \\ 1/2 \end{bmatrix}, \quad r(e_2) = r(0, 1) = \begin{bmatrix} -1/2 \\ \sqrt{3}/2 \end{bmatrix},$$

and these are the columns of r 's matrix,

$$A = \begin{bmatrix} \sqrt{3}/2 & -1/2 \\ 1/2 & \sqrt{3}/2 \end{bmatrix}.$$

This matrix A explicitly describes r .

$$r(x, y) = \begin{bmatrix} \sqrt{3}/2 & -1/2 \\ 1/2 & \sqrt{3}/2 \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix} = \begin{bmatrix} \frac{\sqrt{3}}{2}x - \frac{1}{2}y \\ \frac{1}{2}x + \frac{\sqrt{3}}{2}y \end{bmatrix} = \left(\frac{\sqrt{3}}{2}x - \frac{1}{2}y, \frac{1}{2}x + \frac{\sqrt{3}}{2}y \right).$$

Figures 3.5 through 3.8 show more depictions of linear mappings between spaces of various dimensions. Note that although these mappings stretch and torque their basic input grids, the grids still get taken to configurations of straight lines. Contrast this to how the nonlinear mapping of figure 2.9 bent the basic grid lines into curves.




Figure 3.5. A linear mapping from \mathbf{R} to \mathbf{R}^2




Figure 3.6. A linear mapping from \mathbf{R}^2 to \mathbf{R}^2

We end this section by returning from calculations to intrinsic methods. The following result could have been come immediately after Definition 3.1.1,




Figure 3.7. A linear mapping from \mathbf{R}^3 to \mathbf{R}^3




Figure 3.8. A linear mapping from \mathbf{R}^3 to \mathbf{R}^2

but it has been deferred to this point in order to present some of the objects involved more explicitly first, to make them familiar. However, it is most easily proved abstractly.

Let $\mathcal{L}(\mathbf{R}^n, \mathbf{R}^m)$ denote the set of all linear mappings from \mathbf{R}^n to \mathbf{R}^m . This set not only sits inside the vector space $\mathcal{M}(\mathbf{R}^n, \mathbf{R}^m)$, it is a vector space in its own right:

Proposition 3.1.8 ($\mathcal{L}(\mathbf{R}^n, \mathbf{R}^m)$ Forms a Vector Space). *Suppose that $S, T : \mathbf{R}^n \rightarrow \mathbf{R}^m$ are linear and that $a \in \mathbf{R}$. Then the mappings*

$$S + T, aS : \mathbf{R}^n \rightarrow \mathbf{R}^m$$

are also linear. Consequently, the set of linear mappings from \mathbf{R}^n to \mathbf{R}^m forms a vector space.

Proof. The mappings S and T satisfy (3.1) and (3.2). We must show that $S + T$ and aS do the same. Compute for any $x, y \in \mathbf{R}^n$,

$$\begin{aligned}
& (S + T)(x + y) \\
&= S(x + y) + T(x + y) && \text{by definition of “+” in } \mathcal{M}(\mathbf{R}^n, \mathbf{R}^m) \\
&= S(x) + S(y) + T(x) + T(y) && \text{since } S \text{ and } T \text{ satisfy (3.1)} \\
&= S(x) + T(x) + S(y) + T(y) && \text{since addition in } \mathbf{R}^m \text{ commutes} \\
&= (S + T)(x) + (S + T)(y) && \text{by definition of “+” in } \mathcal{M}(\mathbf{R}^n, \mathbf{R}^m).
\end{aligned}$$

This shows that $S + T$ satisfies (3.1). The other three statements about $S + T$ and aS satisfying (3.1) and (3.2) are similar and left as exercise 3.1.12. Once those are established, the rest of the vector space axioms in $\mathcal{L}(\mathbf{R}^n, \mathbf{R}^m)$ are readily seen to be inherited from $\mathcal{M}(\mathbf{R}^n, \mathbf{R}^m)$. \square

Also, linearity is preserved under composition. That is, if $S : \mathbf{R}^n \rightarrow \mathbf{R}^m$ and $T : \mathbf{R}^p \rightarrow \mathbf{R}^n$ are linear then so is $S \circ T : \mathbf{R}^p \rightarrow \mathbf{R}^m$ (exercise 3.1.13).

Exercises

3.1.1. Prove that $T : \mathbf{R}^n \rightarrow \mathbf{R}^m$ is linear if and only if it satisfies (3.1) and (3.2). (It may help to rewrite (3.1) with the symbols x_1 and x_2 in place of x and y . Then prove one direction by showing that (3.1) and (3.2) are implied by the defining condition for linearity, and prove the other direction by using induction to show that (3.1) and (3.2) imply the defining condition. Note that as pointed out in the text, one direction of this argument has a bit more substance than the other.)

3.1.2. Suppose that $T : \mathbf{R}^n \rightarrow \mathbf{R}^m$ is linear. Show that $T(\mathbf{0}_n) = \mathbf{0}_m$. (An intrinsic argument is nicer.)

3.1.3. Fix a vector $a \in \mathbf{R}^n$. Show that the mapping $T : \mathbf{R}^n \rightarrow \mathbf{R}$ given by $T(x) = \langle a, x \rangle$ is linear, and that $T(e_j) = a_j$ for $j = 1, \dots, n$.

3.1.4. Find the linear mapping $T : \mathbf{R}^3 \rightarrow \mathbf{R}$ such that $T(0, 1, 1) = 1$, $T(1, 0, 1) = 2$, and $T(1, 1, 0) = 3$.

3.1.5. Complete the proof of the componentwise nature of linearity.

3.1.6. Carry out the matrix-by-vector multiplications

$$\begin{bmatrix} 1 & 0 & 0 \\ 1 & 1 & 0 \\ 1 & 1 & 1 \end{bmatrix} \begin{bmatrix} 1 \\ 2 \\ 3 \end{bmatrix}, \quad \begin{bmatrix} a & b \\ c & d \\ e & f \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix}, \quad [x_1 \ \dots \ x_n] \begin{bmatrix} y_1 \\ \vdots \\ y_n \end{bmatrix}, \quad \begin{bmatrix} 1 & -1 & 0 \\ 0 & 1 & -1 \\ -1 & 0 & 1 \end{bmatrix} \begin{bmatrix} 1 \\ 1 \\ 1 \end{bmatrix}.$$

3.1.7. Prove that the identity mapping $id : \mathbf{R}^n \rightarrow \mathbf{R}^n$ is linear. What is its matrix? Explain.

3.1.8. Let θ denote a fixed but generic angle. Argue geometrically that the mapping $R : \mathbf{R}^2 \rightarrow \mathbf{R}^2$ given by counterclockwise rotation by θ is linear, and then find its matrix.

3.1.9. Show that the mapping $Q : \mathbf{R}^2 \rightarrow \mathbf{R}^2$ given by reflection through the x -axis is linear. Find its matrix.

3.1.10. Show that the mapping $P : \mathbf{R}^2 \rightarrow \mathbf{R}^2$ given by orthogonal projection onto the diagonal line $x = y$ is linear. Find its matrix. (See exercise 2.2.15.)

3.1.11. Draw the graph of a generic linear mapping from \mathbf{R}^2 to \mathbf{R}^3 .

3.1.12. Continue the proof of Proposition 3.1.8 by proving the other three statements about $S + T$ and aS satisfying (3.1) and (3.2).

3.1.13. If $S \in \mathcal{L}(\mathbf{R}^n, \mathbf{R}^m)$ and $T \in \mathcal{L}(\mathbf{R}^p, \mathbf{R}^n)$, show that $S \circ T : \mathbf{R}^p \rightarrow \mathbf{R}^m$ lies in $\mathcal{L}(\mathbf{R}^p, \mathbf{R}^m)$.

3.1.14. A mapping $f : \mathbf{R}^n \rightarrow \mathbf{R}^m$ is called **affine** if it has the form $f(x) = T(x) + b$ where $T \in \mathcal{L}(\mathbf{R}^n, \mathbf{R}^m)$ and $b \in \mathbf{R}^m$. State precisely and prove: the composition of affine mappings is affine.

3.1.15. Let $T : \mathbf{R}^n \rightarrow \mathbf{R}^m$ be a linear mapping. Note that since T is continuous and since the absolute value function on \mathbf{R}^m is continuous, the composite function

$$|T| : \mathbf{R}^n \rightarrow \mathbf{R}$$

is continuous.

(a) Let $B = \{x \in \mathbf{R}^n : |x| = 1\}$. Explain why B is a compact subset of \mathbf{R}^n . Explain why it follows that $|T|$ takes a maximum value K on B .

(b) Show that $|T(x)| \leq K|x|$ for all $x \in \mathbf{R}^n$. This result is the **Linear Magnification Boundedness Lemma**. We will use it in chapter 4.

3.1.16. Let $T : \mathbf{R}^n \rightarrow \mathbf{R}^m$ be a linear mapping.

(a) Explain why the set $D = \{x \in \mathbf{R}^n : |x| = 1\}$ is compact.

(b) Use part (a) of this exercise and part (b) of the preceding exercise to explain why therefore the set $\{|T(x)| : x \in D\}$ has a maximum. This maximum is called the **norm** of T and is denoted $\|T\|$.

(c) Explain why $\|T\|$ is the smallest value K that satisfies the displayed condition in part (a) of the preceding exercise. (On the other hand, $\|T\|$ has come to us abstractly while the value of K in the previous exercise is easy to compute.)

(d) Show that for any $S, T \in \mathcal{L}(\mathbf{R}^n, \mathbf{R}^m)$ and any $a \in \mathbf{R}$,

$$\|S + T\| \leq \|S\| + \|T\| \quad \text{and} \quad \|aT\| = |a| \|T\|.$$

Define a distance function

$$d : \mathcal{L}(\mathbf{R}^n, \mathbf{R}^m) \times \mathcal{L}(\mathbf{R}^n, \mathbf{R}^m) \rightarrow \mathbf{R}, \quad d(S, T) = \|T - S\|.$$

Show that this function satisfies the distance properties of Theorem 2.2.8.

(e) Show that for any $S \in \mathcal{L}(\mathbf{R}^n, \mathbf{R}^m)$ and any $T \in \mathcal{L}(\mathbf{R}^p, \mathbf{R}^n)$,

$$\|ST\| \leq \|S\| \|T\|.$$

3.2 Operations on Matrices

Having described abstract objects, the linear mappings $T \in \mathcal{L}(\mathbf{R}^n, \mathbf{R}^m)$, with explicit ones, the matrices $A \in M_{m,n}(\mathbf{R})$ with (i, j) th entry $a_{ij} = T_i(e_j)$, we naturally want to study linear mappings via their matrices. The first step is to develop rules for matrix manipulation corresponding to operations on mappings. Thus if $S, T : \mathbf{R}^n \rightarrow \mathbf{R}^m$ are linear mappings with matrices $A = [a_{ij}]$ and $B = [b_{ij}]$, and if a is a real number, then the matrices for the linear mappings

$$S + T, \quad aS$$

naturally should be denoted

$$A + B, \quad aA.$$

So “+” and “.” (or juxtaposition) are about to acquire new meanings yet again,

$$+ : M_{m,n}(\mathbf{R}) \times M_{m,n}(\mathbf{R}) \rightarrow M_{m,n}(\mathbf{R})$$

and

$$\cdot : \mathbf{R} \times M_{m,n}(\mathbf{R}) \rightarrow M_{m,n}(\mathbf{R}).$$

This is straightforward enough. To define the sum, fix i between 1 and m , and j between 1 and n . Then $S_i(e_j) = a_{ij}$ and $T_i(e_j) = b_{ij}$; the (i, j) th entry of $A + B$ should be

$$\begin{aligned} (S + T)_i(e_j) &= ((S + T)(e_j))_i && \text{by definition of component function} \\ &= (S(e_j) + T(e_j))_i && \text{by definition of addition in } \mathcal{L}(\mathbf{R}^n, \mathbf{R}^m) \\ &= S_i(e_j) + T_i(e_j) && \text{by definition of addition in } \mathbf{R}^m \\ &= S_i(e_j) + T_i(e_j) && \text{by definition of component function} \\ &= a_{ij} + b_{ij}. \end{aligned}$$

Thus, the natural definition for matrix addition is

Definition 3.2.1 (Matrix Addition).

If $A = [a_{ij}]_{m \times n}$ and $B = [b_{ij}]_{m \times n}$ then $A + B = [a_{ij} + b_{ij}]_{m \times n}$.

For example,

$$\begin{bmatrix} 1 & 2 \\ 3 & 4 \end{bmatrix} + \begin{bmatrix} -1 & 0 \\ 2 & 1 \end{bmatrix} = \begin{bmatrix} 0 & 2 \\ 5 & 5 \end{bmatrix}.$$

A similar argument shows that the appropriate definition to make for scalar multiplication of matrices is

Definition 3.2.2 (Scalar-by-Matrix Multiplication).

If $\alpha \in \mathbf{R}$ and $A = [a_{ij}]_{m \times n}$ then $\alpha A = [\alpha a_{ij}]_{m \times n}$.

For example,

$$2 \begin{bmatrix} 1 & 2 \\ 3 & 4 \end{bmatrix} = \begin{bmatrix} 2 & 4 \\ 6 & 8 \end{bmatrix}.$$

The **zero matrix** $\mathbf{0}_{m,n} \in M_{m,n}(\mathbf{R})$, corresponding to the zero mapping in $\mathcal{L}(\mathbf{R}^n, \mathbf{R}^m)$, is the obvious one, with all entries 0. The operations in $M_{m,n}(\mathbf{R})$ precisely mirror those in $\mathcal{L}(\mathbf{R}^n, \mathbf{R}^m)$, so

Proposition 3.2.3 ($M_{m,n}(\mathbf{R})$ Forms a Vector Space). *The set $M_{m,n}(\mathbf{R})$ of m -by- n matrices forms a vector space over \mathbf{R} .*

The remaining important operation on linear mappings is composition. As shown in exercise 3.1.13, if $S : \mathbf{R}^n \rightarrow \mathbf{R}^m$ and $T : \mathbf{R}^p \rightarrow \mathbf{R}^n$ are linear then $S \circ T : \mathbf{R}^p \rightarrow \mathbf{R}^m$ is linear as well. Suppose that S has matrix $A \in M_{m,n}(\mathbf{R})$ and T has matrix $B \in M_{n,p}(\mathbf{R})$. Then $S \circ T$ has a matrix in $M_{m,p}(\mathbf{R})$ that is naturally defined as the matrix-by-matrix product AB , the order of multiplication being chosen for consistency with the composition. Under this specification, the (i, j) th entry of AB must be

$$\begin{aligned} (AB)_{ij} &= (S \circ T)_i(e_j) && \text{since } AB \text{ is the matrix of } S \circ T \\ &= (S_i \circ T)(e_j) && \text{since } S_i \circ T \text{ is the } i\text{th component function} \\ & && \text{of } S \circ T \\ &= S_i(T(e_j)) && \text{by definition of composition} \\ &= S_i\left(\sum_{k=1}^n T_k(e_j)e_k\right) && \text{decomposing } T(e_j) \text{ into its components} \\ &= S_i\left(\sum_{k=1}^n b_{kj}e_k\right) && \text{since } B \text{ is the matrix of } T \\ &= \sum_{k=1}^n b_{kj}S_i(e_k) && \text{since } S_i \text{ is linear} \\ &= \sum_{k=1}^n b_{kj}a_{ik} && \text{since } A \text{ is the matrix of } S \\ &= \sum_{k=1}^n a_{ik}b_{kj} && \text{since multiplication in } \mathbf{R} \text{ commutes.} \end{aligned}$$

So matrix-by-matrix multiplication is specified by

Definition 3.2.4 (Matrix Multiplication).

$$\text{If } A = [a_{ij}]_{m \times n} \text{ and } B = [b_{ij}]_{n \times p} \text{ then } AB = \left[\sum_{k=1}^n a_{ik}b_{kj} \right]_{m \times p}.$$

In words, to multiply A by B , A must have as many columns as B has rows, in which case the (i, j) th entry of AB is the inner product of the i th row of A and the j th column of B .

Note that when B has only one column, this reduces to the matrix-by-vector multiplication from the previous section.

For example, consider the matrices

$$A = \begin{bmatrix} 1 & 2 & 3 \\ 4 & 5 & 6 \end{bmatrix}, \quad B = \begin{bmatrix} 1 & -2 \\ 2 & -3 \\ 3 & -4 \end{bmatrix}, \quad C = \begin{bmatrix} 4 & 5 \\ 6 & 7 \end{bmatrix},$$

$$D = \begin{bmatrix} 1 & 1 & 1 \\ 0 & 1 & 1 \\ 0 & 0 & 1 \end{bmatrix}, \quad E = [a \ b \ c], \quad F = \begin{bmatrix} x \\ y \\ z \end{bmatrix}.$$

Some products among these (verify!) are

$$AB = \begin{bmatrix} 14 & -20 \\ 32 & -47 \end{bmatrix}, \quad BC = \begin{bmatrix} -8 & -9 \\ -10 & -11 \\ -12 & -13 \end{bmatrix}, \quad AD = \begin{bmatrix} 1 & 3 & 6 \\ 4 & 9 & 15 \end{bmatrix},$$

$$DB = \begin{bmatrix} 6 & -9 \\ 5 & -7 \\ 3 & -4 \end{bmatrix}, \quad AF = \begin{bmatrix} x + 2y + 3z \\ 4x + 5y + 6z \end{bmatrix}, \quad FE = \begin{bmatrix} ax & bx & cx \\ ay & by & cy \\ az & bz & cz \end{bmatrix},$$

$$EF = ax + by + cz.$$

Matrix multiplication is not commutative. Indeed, when the product AB is defined, the product BA may not be, or it may be but have different dimensions from AB ; cf. EF and FE above. Even when A and B are both n -by- n , so that AB and BA are likewise n -by- n , the products need not agree. For example,

$$\begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix} \begin{bmatrix} 0 & 0 \\ 1 & 0 \end{bmatrix} = \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix}, \quad \begin{bmatrix} 0 & 0 \\ 1 & 0 \end{bmatrix} \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix} = \begin{bmatrix} 0 & 0 \\ 0 & 1 \end{bmatrix}.$$

Of particular interest is the matrix associated to the identity mapping,

$$id : \mathbf{R}^n \longrightarrow \mathbf{R}^n, \quad id(x) = x.$$

Naturally, this matrix is denoted the **identity matrix**; it is written I_n . Since $id_i(e_j) = \delta_{ij}$,

$$I_n = [\delta_{ij}]_{n \times n} = \begin{bmatrix} 1 & 0 & \dots & 0 \\ 0 & 1 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & 1 \end{bmatrix}.$$

Although matrix multiplication fails to commute, it does have the following properties.

Proposition 3.2.5 (Properties of Matrix Multiplication). *Matrix multiplication is associative,*

$$A(BC) = (AB)C \quad \text{for } A \in M_{m,n}(\mathbf{R}), B \in M_{n,p}(\mathbf{R}), C \in M_{p,q}(\mathbf{R}).$$

Matrix multiplication distributes over matrix addition,

$$\begin{aligned} A(B + C) &= AB + AC \quad \text{for } A \in M_{m,n}(\mathbf{R}), B, C \in M_{n,p}(\mathbf{R}), \\ (A + B)C &= AC + BC \quad \text{for } A, B \in M_{m,n}(\mathbf{R}), C \in M_{n,p}(\mathbf{R}). \end{aligned}$$

Scalar multiplication passes through matrix multiplication,

$$\alpha(AB) = (\alpha A)B = A(\alpha B) \quad \text{for } \alpha \in \mathbf{R}, A \in M_{m,n}(\mathbf{R}), B \in M_{n,p}(\mathbf{R}).$$

The identity matrix is a multiplicative identity,

$$I_m A = A = A I_n \quad \text{for } A \in M_{m,n}(\mathbf{R}).$$

Proof. The right way to show these is intrinsic, by remembering that addition, scalar multiplication, and multiplication of matrices precisely mirror addition, scalar multiplication, and composition of mappings. For example, if A, B, C are the matrices of the linear mappings $S \in \mathcal{L}(\mathbf{R}^n, \mathbf{R}^m)$, $T \in \mathcal{L}(\mathbf{R}^p, \mathbf{R}^n)$, and $U \in \mathcal{L}(\mathbf{R}^q, \mathbf{R}^p)$, then $(AB)C$ and $A(BC)$ are the matrices of $(S \circ T) \circ U$ and $S \circ (T \circ U)$. But these two mappings are the same since in general the composition of mappings is associative. That is, for any $x \in \mathbf{R}^q$,

$$\begin{aligned} ((S \circ T) \circ U)(x) &= (S \circ T)(U(x)) = S(T(U(x))) = S((T \circ U)(x)) \\ &= (S \circ (T \circ U))(x). \end{aligned}$$

Consequently $(AB)C = A(BC)$.

Alternatively, one can verify the equalities elementwise by manipulating sums. Adopting the notation M_{ij} for the (i, j) th entry of a matrix M ,

$$\begin{aligned} (A(BC))_{ij} &= \sum_{k=1}^n A_{ik}(BC)_{kj} = \sum_{k=1}^n A_{ik} \sum_{l=1}^p B_{kl} C_{lj} = \sum_{k=1}^n \sum_{l=1}^p A_{ik} B_{kl} C_{lj} \\ &= \sum_{l=1}^p \sum_{k=1}^n A_{ik} B_{kl} C_{lj} = \sum_{l=1}^p (AB)_{il} C_{lj} = ((AB)C)_{ij}. \end{aligned}$$

The steps here are not explained in detail because the author finds this method as unenlightening as it is unnecessary: the coordinates work because they must, but their presence only clutters the argument. The other equalities are similar. \square

Composing mappings is most interesting when all the mappings in question take a set S back to the same set S , for the set of such mappings is closed under composition. In particular, $\mathcal{L}(\mathbf{R}^n, \mathbf{R}^n)$ is closed under composition. The corresponding statement about matrices is that $M_n(\mathbf{R})$ is closed under multiplication.

Exercises

3.2.1. Justify Definition 3.2.2 of scalar multiplication of matrices.

3.2.2. Carry out the matrix multiplications

$$\begin{bmatrix} a & b \\ c & d \end{bmatrix} \begin{bmatrix} d & -b \\ -c & a \end{bmatrix}, \quad [x_1 \ x_2 \ x_3] \begin{bmatrix} a_1 & b_1 \\ a_2 & b_2 \\ a_3 & b_3 \end{bmatrix}, \quad \begin{bmatrix} 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 \end{bmatrix}^2,$$

$$\begin{bmatrix} 1 & 1 & 1 \\ 0 & 1 & 1 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} 1 & 0 & 0 \\ 1 & 1 & 0 \\ 1 & 1 & 1 \end{bmatrix}, \quad \begin{bmatrix} 1 & 0 & 0 \\ 1 & 1 & 0 \\ 1 & 1 & 1 \end{bmatrix} \begin{bmatrix} 1 & 1 & 1 \\ 0 & 1 & 1 \\ 0 & 0 & 1 \end{bmatrix}.$$

3.2.3. Prove more of Proposition 3.2.5, that $A(B+C) = AB+AC$, $(\alpha A)B = A(\alpha B)$ and $I_m A = A$ for suitable matrices A, B, C and any scalar α .

3.2.4. Let $A = [a_{ij}]$ be a matrix in $M_{m,n}(\mathbf{R})$. Its **transpose** $A^t \in M_{n,m}(\mathbf{R})$ is the matrix obtained by flipping A about its Northwest–Southeast diagonal. Thus the rows of A^t are the columns of A , the columns of A^t are the rows of A , and the (i, j) th entry of A^t is a_{ji} . Show that

$$(AB)^t = B^t A^t \quad \text{for all } A \in M_{m,n}(\mathbf{R}) \text{ and } B \in M_{n,p}(\mathbf{R}).$$

3.2.5. For any matrix $A \in M_{m,n}(\mathbf{R})$ and column vector $a \in \mathbf{R}^m$ define the affine mapping (cf. exercise 3.1.14)

$$\text{Aff}_{A,a} : \mathbf{R}^n \longrightarrow \mathbf{R}^m$$

by the rule $\text{Aff}_{A,a}(x) = Ax + a$ for all $x \in \mathbf{R}^n$, viewing x as a column vector.

- (a) Explain why every affine mapping from \mathbf{R}^n to \mathbf{R}^m takes this form.
 (b) Given such A and a , define the matrix $A' \in M_{m+1,n+1}(\mathbf{R})$ to be

$$A' = \begin{bmatrix} A & a \\ \mathbf{0}_n & 1 \end{bmatrix}.$$

Show that for all $x \in \mathbf{R}^n$,

$$A' \begin{bmatrix} x \\ 1 \end{bmatrix} = \begin{bmatrix} \text{Aff}_{A,a}(x) \\ 1 \end{bmatrix}.$$

Thus, affine mappings, like linear mappings, behave as matrix-by-vector multiplications but where the vectors are the usual input and output vectors augmented with an extra “1” at the bottom.

(c) If the affine mapping $\text{Aff}_{B,b} : \mathbf{R}^p \longrightarrow \mathbf{R}^n$ determined by $B \in M_{n,p}(\mathbf{R})$ and $b \in \mathbf{R}^n$ has matrix

$$B' = \begin{bmatrix} B & b \\ \mathbf{0}_p & 1 \end{bmatrix}$$

show that $\text{Aff}_{A,a} \circ \text{Aff}_{B,b} : \mathbf{R}^p \longrightarrow \mathbf{R}^m$ has matrix $A'B'$. Thus, matrix multiplication is compatible with composition of affine mappings.

3.3 The Inverse of a Linear Mapping

Given a linear mapping $S : \mathbf{R}^n \rightarrow \mathbf{R}^m$, does it have an inverse? That is, is there a mapping $T : \mathbf{R}^m \rightarrow \mathbf{R}^n$ such that

$$S \circ T = id_m \quad \text{and} \quad T \circ S = id_n?$$

If so, what is T ?

The symmetry of the previous display shows that if T is an inverse of S then S is an inverse of T in turn. Also, the inverse T , if it exists, must be unique, for if $T' : \mathbf{R}^m \rightarrow \mathbf{R}^n$ also inverts S then

$$T' = T' \circ id_m = T' \circ (S \circ T) = (T' \circ S) \circ T = id_n \circ T = T.$$

Thus T can unambiguously be denoted S^{-1} . In fact, this argument has shown a little bit more than claimed: If T' inverts S from the left and T inverts S from the right then $T' = T$. On the other hand, the argument does *not* show that if T inverts S from the left then T also inverts S from the right—this is not true.

If the inverse T exists then it too is linear. To see this, note that the elementwise description of S and T being inverses of one another is that every $y \in \mathbf{R}^m$ takes the form $y = S(x)$ for some $x \in \mathbf{R}^n$, every $x \in \mathbf{R}^n$ takes the form $x = T(y)$ for some $y \in \mathbf{R}^m$, and

$$\text{for all } x \in \mathbf{R}^n \text{ and } y \in \mathbf{R}^m, \quad y = S(x) \iff x = T(y).$$

Now compute that for any $y_1, y_2 \in \mathbf{R}^m$,

$$\begin{aligned} T(y_1 + y_2) &= T(S(x_1) + S(x_2)) \quad \text{for some } x_1, x_2 \in \mathbf{R}^n \\ &= T(S(x_1 + x_2)) \quad \text{since } S \text{ is linear} \\ &= x_1 + x_2 \quad \text{since } T \text{ inverts } S \\ &= T(y_1) + T(y_2) \quad \text{since } y_1 = S(x_1) \text{ and } y_2 = S(x_2). \end{aligned}$$

Thus T satisfies (3.1). The argument that T satisfies (3.2) is similar.

Since matrices are more explicit than linear mappings, we replace the question at the beginning of this section with its matrix counterpart: Given a matrix $A \in M_{m,n}(\mathbf{R})$, does it have an inverse matrix, a matrix $B \in M_{n,m}(\mathbf{R})$ such that

$$AB = I_m \quad \text{and} \quad BA = I_n?$$

As above, if the inverse exists then it is unique, and so it can be denoted A^{-1} .

The first observation to make is that if the equation $Ax = \mathbf{0}_m$ has a nonzero solution $x \in \mathbf{R}^n$ then A has no inverse. Indeed, an inverse A^{-1} would give

$$x = I_n x = (A^{-1}A)x = A^{-1}(Ax) = A^{-1}\mathbf{0}_m = \mathbf{0}_n,$$

contradicting the fact that x is nonzero. This raises a subordinate question: When does the matrix equation

$$Ax = \mathbf{0}_m$$

have nonzero solutions $x \in \mathbf{R}^n$?

For example, let A be the 5-by-6 matrix

$$A = \begin{bmatrix} 5 & 1 & 17 & 26 & 1 & 55 \\ -3 & -1 & -13 & -20 & 0 & -28 \\ -2 & 1 & 3 & 5 & 0 & 3 \\ -2 & 0 & -4 & -6 & 0 & -10 \\ 5 & 0 & 10 & 15 & 1 & 42 \end{bmatrix}.$$

If there is a nonzero $x \in \mathbf{R}^6$ such that $Ax = \mathbf{0}_5$ then A is not invertible.

Left multiplication by certain special matrices will simplify the matrix A .

Definition 3.3.1 (Elementary Matrices). *There are three kinds of elementary matrices. For any $i, j \in \{1, \dots, m\}$ ($i \neq j$) and any $a \in \mathbf{R}$, the m -by- m $(i; j, a)$ recombine matrix is*

$$R_{i;j,a} = \begin{bmatrix} 1 & & & & & \\ & \ddots & & & & \\ & & 1 & & a & \\ & & & \ddots & & \\ & & & & 1 & \\ & & & & & \ddots \\ & & & & & & 1 \end{bmatrix}.$$

(Here the a sits in the (i, j) th position, the diagonal entries are 1 and all other entries are 0. The a is above the diagonal as shown only when $i < j$, otherwise it is below.)

For any $i \in \{1, \dots, m\}$ and any nonzero $a \in \mathbf{R}$, the m -by- m (i, a) scale matrix is

$$S_{i,a} = \begin{bmatrix} 1 & & & & & \\ & \ddots & & & & \\ & & 1 & & & \\ & & & a & & \\ & & & & 1 & \\ & & & & & \ddots \\ & & & & & & 1 \end{bmatrix}.$$

(Here the a sits in the i th diagonal position, all other diagonal entries are 1 and all other entries are 0.)

For any $i, j \in \{1, \dots, m\}$ ($i \neq j$), the m -by- m $(i; j)$ transposition matrix is

$$R_{1;2,3} \cdot \begin{bmatrix} 1 & 2 & 3 \\ 4 & 5 & 6 \end{bmatrix} = \begin{bmatrix} 13 & 17 & 21 \\ 4 & 5 & 6 \end{bmatrix},$$

because $R_{1;2,3}$ adds 3 times the second row to the first.

The next result is that performing row operations on A doesn't change the set of solutions x to the equation $Ax = \mathbf{0}_m$.

Lemma 3.3.3 (Invertibility of Products of the Elementary Matrices). *Products of elementary matrices are invertible. More specifically:*

(1) *The elementary matrices are invertible by other elementary matrices. Specifically,*

$$(R_{i;j,a})^{-1} = R_{i;j,-a}, \quad (S_{i,a})^{-1} = S_{i,a^{-1}}, \quad (T_{i;j})^{-1} = T_{i;j}.$$

(2) *If the m -by- m matrices M and N are invertible by M^{-1} and N^{-1} , then the product matrix MN is invertible by $N^{-1}M^{-1}$. (Note the order reversal.)*

(3) *Any product of elementary matrices is invertible by another such product, specifically the product of the inverses of the original matrices, but taken in reverse order.*

Proof. (1) To prove that $R_{i;j,-a}R_{i;j,a} = I_m$, note that $R_{i;j,a}$ is the identity matrix I_m with a times its j th row added to its i th row, and multiplying this from the left by $R_{i;j,-a}$ subtracts back off a times the j th row, restoring I_m . The proof that $R_{i;j,a}R_{i;j,-a} = I_m$ is either done similarly or by citing the proof just given with a replaced by $-a$. The rest of (1) is similar.

(2) Compute:

$$(MN)(N^{-1}M^{-1}) = M(NN^{-1})M^{-1} = MI_mM^{-1} = MM^{-1} = I_m.$$

Similarly for $(N^{-1}M^{-1})(MN) = I_m$.

(3) This is immediate from (1) and (2). □

Proposition 3.3.4 (Persistence of Solution). *Let A be an m -by- n matrix and let P be a product of m -by- m elementary matrices. Then the equations*

$$Ax = \mathbf{0}_m \quad \text{and} \quad (PA)x = \mathbf{0}_m$$

are satisfied by the same vectors x in \mathbf{R}^n .

Proof. Suppose that the vector $x \in \mathbf{R}^n$ satisfies the left equation, $Ax = \mathbf{0}_m$. Then

$$(PA)x = P(Ax) = P\mathbf{0}_m = \mathbf{0}_m.$$

Conversely, suppose that x satisfies $(PA)x = \mathbf{0}_m$. Lemma 3.3.3 says that P has an inverse P^{-1} , so

$$Ax = I_m Ax = (P^{-1}P)Ax = P^{-1}(PA)x = P^{-1}\mathbf{0}_m = \mathbf{0}_m.$$

□

The machinery is in place to solve the equation $Ax = \mathbf{0}_5$ where as before,

$$A = \begin{bmatrix} 5 & 1 & 17 & 26 & 1 & 55 \\ -3 & -1 & -13 & -20 & 0 & -28 \\ -2 & 1 & 3 & 5 & 0 & 3 \\ -2 & 0 & -4 & -6 & 0 & -10 \\ 5 & 0 & 10 & 15 & 1 & 42 \end{bmatrix}.$$

Scale A 's fourth row by $-1/2$; transpose A 's first and fourth rows:

$$T_{1;4}S_{4,-1/2}A = \begin{bmatrix} 1 & 0 & 2 & 3 & 0 & 5 \\ -3 & -1 & -13 & -20 & 0 & -28 \\ -2 & 1 & 3 & 5 & 0 & 3 \\ 5 & 1 & 17 & 26 & 1 & 55 \\ 5 & 0 & 10 & 15 & 1 & 42 \end{bmatrix} \stackrel{\text{call}}{=} B.$$

Note that B has a 1 as the leftmost entry of its first row. Recombine various multiples of the first row with the other rows to put 0's beneath the leading 1 of the first row:

$$R_{5;1,-5}R_{4;1,-5}R_{3;1,2}R_{2;1,3}B = \begin{bmatrix} 1 & 0 & 2 & 3 & 0 & 5 \\ 0 & -1 & -7 & -11 & 0 & -13 \\ 0 & 1 & 7 & 11 & 0 & 13 \\ 0 & 1 & 7 & 11 & 1 & 30 \\ 0 & 0 & 0 & 0 & 1 & 17 \end{bmatrix} \stackrel{\text{call}}{=} C.$$

Recombine various multiples of the second row with the others to put 0's above and below its leftmost nonzero entry; scale the second row to make its leading nonzero entry a 1:

$$S_{2,-1}R_{4;2,1}R_{3;2,1}C = \begin{bmatrix} 1 & 0 & 2 & 3 & 0 & 5 \\ 0 & 1 & 7 & 11 & 0 & 13 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 17 \\ 0 & 0 & 0 & 0 & 1 & 17 \end{bmatrix} \stackrel{\text{call}}{=} D.$$

Transpose the third and fifth rows; put 0's above and below the leading 1 in the third row:

$$R_{4;3,-1}T_{3;5}D = \begin{bmatrix} 1 & 0 & 2 & 3 & 0 & 5 \\ 0 & 1 & 7 & 11 & 0 & 13 \\ 0 & 0 & 0 & 0 & 1 & 17 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix} \stackrel{\text{call}}{=} E.$$

Matrix E is a prime example of a so-called *echelon matrix*. (The term will be defined precisely in a moment.) Its virtue is that the equation $Ex = \mathbf{0}_5$ is now easy to solve. This equation expands out to

$$Ex = \begin{bmatrix} 1 & 0 & 2 & 3 & 0 & 5 \\ 0 & 1 & 7 & 11 & 0 & 13 \\ 0 & 0 & 0 & 0 & 1 & 17 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \\ x_3 \\ x_4 \\ x_5 \\ x_6 \end{bmatrix} = \begin{bmatrix} x_1 + 2x_3 + 3x_4 + 5x_6 \\ x_2 + 7x_3 + 11x_4 + 13x_6 \\ x_5 + 17x_6 \\ 0 \\ 0 \\ 0 \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \end{bmatrix}.$$

Matching the components in the last equality gives

$$\begin{aligned} x_1 &= -2x_3 - 3x_4 - 5x_6 \\ x_2 &= -7x_3 - 11x_4 - 13x_6 \\ x_5 &= -17x_6. \end{aligned}$$

Thus, x_3, x_4 and x_6 are free variables that may take any values we wish, but then x_1, x_2 and x_5 are determined from these equations. For example, setting $x_3 = -5, x_4 = 3, x_6 = 2$ gives the solution $x = (-9, -24, -5, 3, -34, 2)$.

Definition 3.3.5 (Echelon Matrix). A matrix E is called **echelon** if it has the form

$$E = \begin{bmatrix} 0 \dots 0 & \boxed{1} * \dots * 0 * \dots * 0 * \dots \\ & & \boxed{1} * \dots * 0 * \dots \\ & & & \boxed{1} * \dots \end{bmatrix}.$$

Here the $*$'s are arbitrary entries and all entries below the stairway are 0. Thus each row's first nonzero entry is a 1, each row's leading 1 is farther right than that of the row above it, each leading 1 has a column of 0's above it, and any rows of 0's are at the bottom.

Note that the identity matrix I is a special case of an echelon matrix.

The algorithm for reducing any matrix A to echelon form by row operations should be fairly clear from the previous example. The interested reader may want to codify it more formally, perhaps in the form of a computer program. Although different sequences of row operations may reduce A to echelon form, the resulting echelon matrix E will always be the same. This result can be proved by induction on the number of columns of A , and its proof is in many linear algebra books.

Theorem 3.3.6 (Matrices Reduce to Echelon Form). Every matrix A row reduces to a unique echelon matrix E .

In an echelon matrix E , the columns with leading 1's are called **new columns**, and all others are **old columns**. The recipe for solving the equation $Ex = \mathbf{0}_m$ is then

1. Freely choose the entries in x that correspond to the old columns of E .

- Then each nonzero row of E will determine the entry of x corresponding to its leading 1 (which sits in a new column). This entry will be a linear combination of the free entries to its right.

Let's return to the problem of determining whether $A \in M_{m,n}(\mathbf{R})$ is invertible. The idea was to see if the equation $Ax = \mathbf{0}_m$ has any nonzero solutions x , in which case A is not invertible. Equivalently, we may check whether $Ex = \mathbf{0}_m$ has nonzero solutions, where E is the echelon matrix to which A row reduces. The recipe for solving $Ex = \mathbf{0}_m$ shows that there are nonzero solutions unless all of the columns are new.

If $A \in M_{m,n}(\mathbf{R})$ has more columns than rows then its echelon matrix E must have old columns. This is because each new column comes from the leading 1 in a distinct row, so

$$\text{new columns of } E \leq \text{rows of } E < \text{columns of } E,$$

showing that not all the columns are new. Thus A is not invertible when $m < n$. On the other hand, if $A \in M_{m,n}(\mathbf{R})$ has more rows than columns and it has an inverse matrix $A^{-1} \in M_{n,m}(\mathbf{R})$, then A^{-1} in turn has inverse A , but this is impossible since A^{-1} has more columns than rows. Thus A is also not invertible when $m > n$.

The remaining case is that A is square. The only square echelon matrix with all new columns is I , the identity matrix (exercise 3.3.10). Thus, unless A 's echelon matrix is I , A is not invertible. On the other hand, if A 's echelon matrix is I , then $PA = I$ for some product P of elementary matrices. Multiply from the left by P^{-1} to get $A = P^{-1}$; this is invertible by P , giving $A^{-1} = P$. Summarizing,

Theorem 3.3.7 (Invertibility and Echelon Form for Matrices). *A non-square matrix A is never invertible. A square matrix A is invertible if and only if its echelon form is the identity matrix.*

When A is square, the discussion above gives an algorithm that simultaneously checks whether it is invertible and finds its inverse when it is.

Proposition 3.3.8 (Matrix Inversion Algorithm). *Given $A \in M_n(\mathbf{R})$, set up the matrix*

$$B = [A \mid I_n]$$

in $M_{n,2n}(\mathbf{R})$. Carry out row operations on this matrix to reduce the left side to echelon form. If the left side reduces to I_n then A is invertible and the right side is A^{-1} . If the left side doesn't reduce to I_n then A is not invertible.

This works because if B is left multiplied by a product P of elementary matrices, the result is

$$PB = [PA \mid P].$$

As discussed, $PA = I_n$ exactly when $P = A^{-1}$.

For example, the calculation

$$R_{1;2,1}R_{2;3,1} \left[\begin{array}{ccc|ccc} 1 & -1 & 0 & 1 & 0 & 0 \\ 0 & 1 & -1 & 0 & 1 & 0 \\ 0 & 0 & 1 & 0 & 0 & 1 \end{array} \right] = \left[\begin{array}{ccc|ccc} 1 & 0 & 0 & 1 & 1 & 1 \\ 0 & 1 & 0 & 0 & 1 & 1 \\ 0 & 0 & 1 & 0 & 0 & 1 \end{array} \right]$$

shows that

$$\left[\begin{array}{ccc|ccc} 1 & -1 & 0 & 1 & 0 & 0 \\ 0 & 1 & -1 & 0 & 1 & 0 \\ 0 & 0 & 1 & 0 & 0 & 1 \end{array} \right]^{-1} = \left[\begin{array}{ccc|ccc} 1 & 1 & 1 & 1 & 1 & 1 \\ 0 & 1 & 1 & 0 & 1 & 1 \\ 0 & 0 & 1 & 0 & 0 & 1 \end{array} \right].$$

We now have an algorithmic answer to the question at the beginning of the section.

Theorem 3.3.9 (Invertibility and Echelon Form for Mappings). *The linear mapping $T : \mathbf{R}^n \rightarrow \mathbf{R}^m$ is invertible only when $m = n$ and its matrix A has echelon matrix I_n , in which case its inverse is the linear mapping with matrix A^{-1} .*

Exercises

3.3.1. Write down the following 3-by-3 elementary matrices and their inverses: $R_{3;2,\pi}$, $S_{3,3}$, $T_{3;2}$, $T_{2;3}$.

3.3.2. Finish the proof of Proposition 3.3.2.

3.3.3. Let $A = \begin{bmatrix} 1 & 2 \\ 3 & 4 \\ 5 & 6 \end{bmatrix}$. Evaluate the following products without actually multiplying matrices: $R_{3;2,\pi}A$, $S_{3,3}A$, $T_{3;2}A$, $T_{2;3}A$.

3.3.4. Finish the proof of Lemma 2.3.3, part (1).

3.3.5. What is the effect of right multiplying the m -by- n matrix M by an n -by- n matrix $R_{i;j,a}$? By $S_{i,a}$? By $T_{i;j}$?


3.3.6. Recall the transpose of a matrix M (cf. exercise LM:OM:ex4), denoted M^t . Prove: $R_{i;j,a}^t = R_{j;i,a}$; $S_{i,a}^t = S_{i,a}$; $T_{i;j}^t = T_{i;j}$. Use these results and the formula $(AB)^t = B^tA^t$ to redo the previous problem.

3.3.7. Are the following matrices echelon? For each matrix M , solve the equation $Mx = \mathbf{0}$.

$$\begin{bmatrix} 1 & 0 & 3 \\ 0 & 1 & 1 \\ 0 & 0 & 1 \end{bmatrix}, \quad \begin{bmatrix} 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 \end{bmatrix}, \quad \begin{bmatrix} 1 & 1 & 0 & 0 \\ 0 & 0 & 1 & 1 \end{bmatrix}, \quad \begin{bmatrix} 0 & 0 \\ 1 & 0 \\ 0 & 1 \\ 0 & 0 \end{bmatrix}, \quad \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 1 & 0 \\ 0 & 0 & 1 & 0 \end{bmatrix}, \quad \begin{bmatrix} 0 & 1 & 1 \\ 1 & 0 & 3 \\ 0 & 0 & 0 \end{bmatrix}.$$

3.3.8. For each matrix A solve the equation $Ax = \mathbf{0}$.

$$\begin{bmatrix} -1 & 1 & 4 \\ 1 & 3 & 8 \\ 1 & 2 & 5 \end{bmatrix}, \quad \begin{bmatrix} 2 & -1 & 3 & 2 \\ 1 & 4 & 0 & 1 \\ 2 & 6 & -1 & 5 \end{bmatrix}, \quad \begin{bmatrix} 3 & -1 & 2 \\ 2 & 1 & 1 \\ 1 & -3 & 0 \end{bmatrix}.$$

3.3.9. Balance the chemical equation

3.3.10. Prove by induction that the only square echelon matrix with all new columns is the identity matrix.

3.3.11. Are the following matrices invertible? Find the inverse when possible.

$$\begin{bmatrix} 1 & -1 & 1 \\ 2 & 0 & 1 \\ 3 & 0 & 1 \end{bmatrix}, \quad \begin{bmatrix} 2 & 5 & -1 \\ 4 & -1 & 2 \\ 6 & 4 & 1 \end{bmatrix}, \quad \begin{bmatrix} 1 & \frac{1}{2} & \frac{1}{3} \\ \frac{1}{2} & \frac{1}{3} & \frac{1}{4} \\ \frac{1}{3} & \frac{1}{4} & \frac{1}{5} \end{bmatrix}.$$

3.3.12. The matrix A is called **lower triangular** if $a_{ij} = 0$ whenever $i < j$. If A is a lower triangular square matrix with all diagonal entries equal to 1, show that A is invertible and A^{-1} takes the same form.

3.3.13. This exercise refers back to the Gram–Schmidt exercise in chapter 2. That exercise expresses the relation between the vectors $\{x'_j\}$ and the vectors $\{x_j\}$ formally as $x' = Ax$ where x' is a column vector whose entries are the vectors x'_1, \dots, x'_n , x is the corresponding column vector of x_j 's, and A is an n -by- n lower triangular matrix.

Show that each x_j has the form

$$x_j = a'_{j1}x'_1 + a'_{j2}x'_2 + \cdots + a'_{j,j-1}x'_{j-1} + x'_j,$$

and thus any linear combination of the original $\{x_j\}$ is also a linear combination of the new $\{x'_j\}$.

3.4 Inhomogeneous Linear Equations

The question of whether a linear mapping T is invertible led to solving the linear equation $Ax = \mathbf{0}$, where A was the matrix of T . Such a linear equation, with right side $\mathbf{0}$, is called **homogeneous**. An **inhomogeneous** linear equation has nonzero right side,

$$Ax = b, \quad A \in M_{m,n}(\mathbf{R}), \quad x \in \mathbf{R}^n, \quad b \in \mathbf{R}^m, \quad b \neq \mathbf{0}.$$

The methods of the homogeneous case apply here too. If P is a product of m -by- m elementary matrices such that PA is echelon (call it E), then multiplying the inhomogeneous equation from the left by P gives

$$Ex = Pb,$$

and since Pb is just a vector, the solutions to this can be read off as in the homogeneous case. There may not always be solutions, however.

Exercises

3.4.1. Solve the inhomogeneous equations

$$\begin{bmatrix} 1 & -1 & 2 \\ 2 & 0 & 2 \\ 1 & -3 & 4 \end{bmatrix} x = \begin{bmatrix} 1 \\ 1 \\ 2 \end{bmatrix}, \quad \begin{bmatrix} 1 & -2 & 1 & 2 \\ 1 & 1 & -1 & 1 \\ 1 & 7 & -5 & -1 \end{bmatrix} x = \begin{bmatrix} 1 \\ 2 \\ 3 \end{bmatrix}.$$

3.4.2. For what values b_1, b_2, b_3 does the equation

$$\begin{bmatrix} 3 & -1 & 2 \\ 2 & 1 & 1 \\ 1 & -3 & 0 \end{bmatrix} x = \begin{bmatrix} b_1 \\ b_2 \\ b_3 \end{bmatrix}$$

have a solution?

3.4.3. A parent has a son and a daughter. The parent is four times as old as the daughter, the daughter is four years older than the son. In three years the parent will be five times as old as the son. How old are the parent, daughter and son?

3.4.4. Show that to solve an inhomogeneous linear equation, one may solve a homogeneous system in one more variable and then restrict to solutions where the last variable is equal to -1 .

3.5 The Determinant: Characterizing Properties and Their Consequences

In this section all matrices are square, n -by- n . The goal is to define a function that takes such a matrix, with its n^2 entries, and returns a single number. This putative function is called the **determinant**,

$$\det : M_n(\mathbf{R}) \longrightarrow \mathbf{R}.$$

For any square matrix $A \in M_n(\mathbf{R})$, the scalar $\det(A)$ should contain as much algebraic and geometric information about the matrix as possible. Not surprisingly, so informative a function is complicated to encode.

This context nicely demonstrates a pedagogical principle already mentioned in section 3.1: characterizing a mathematical object illuminates its construction and its use. Rather than beginning with a definition of the determinant, we will stipulate a few natural behaviors for it, and then we will eventually see that

- there is a function with these behaviors (*existence*),
- there is only one such function (*uniqueness*), and, most importantly,
- these behaviors, rather than the definition, further show how the function works (*consequences*).

We could start at the first bullet and proceed from the construction of the determinant to its properties, but when a construction is complicated (as the determinant's construction is) it fails to communicate intent, and pulling it out of thin air as the starting point of a long discussion is an obstacle to understanding. A few naturally gifted readers will see what the unexplained idea really is, enabling them to skim the ensuing technicalities and go on to start using the determinant effectively; some other tough-minded readers can work through the machinery and then see its operational consequences; but it is all too easy for the rest of us to be defeated by disoriented detail-fatigue before the presentation gets to the consequential points and provides any energizing clarity.

Another option would be to start at the second bullet, letting the desired properties of the determinant guide our construction of it. This process wouldn't be as alienating as starting at the first bullet, but deriving the determinant's necessary construction has only short-term benefit since we intend to use the construction as little as possible. Working through the derivation would still squander our limited energy on the internal mechanisms of the determinant before getting to its behavior, when its behavior is what gives us perspective in understanding it. We first want to learn to use the determinant easily and artfully. This will make its internals feel of secondary importance, as they should.

The upshot is that in this section we will pursue the third bullet (consequences), and then the next section will proceed to the second bullet (uniqueness) and finally the first one (existence).

Instead of viewing the determinant only as a function of a matrix $A \in M_n(\mathbf{R})$ with n^2 scalar entries, view it also as a function of A 's n rows, each of which is an n -vector. If A has rows r_1, \dots, r_n , write $\det(r_1, \dots, r_n)$ for $\det(A)$. Thus, \det is now being interpreted as a function of n vectors, i.e., the domain of \det is n copies of \mathbf{R}^n ,

$$\det : \mathbf{R}^n \times \cdots \times \mathbf{R}^n \longrightarrow \mathbf{R}.$$

The advantage of this view is that now we can impose conditions on the determinant, using language already at our disposal in a natural way. Specifically, we make three requirements:

- (1) The determinant is **multilinear**, meaning that it is linear as a function of each of its vector variables. That is, for any vectors $r_1, \dots, r_k, r'_k, \dots, r_n$ and any scalars α, α' ,

$$\begin{aligned} \det(r_1, \dots, \alpha r_k + \alpha' r'_k, \dots, r_n) &= \alpha \det(r_1, \dots, r_k, \dots, r_n) \\ &\quad + \alpha' \det(r_1, \dots, r'_k, \dots, r_n). \end{aligned}$$

- (2) The determinant is **skew-symmetric** as a function of its vector variables, meaning that exchanging any two inputs negates the determinant,

$$\det(r_1, \dots, r_j, \dots, r_i, \dots, r_n) = -\det(r_1, \dots, r_i, \dots, r_j, \dots, r_n).$$

(Here $i \neq j$.) Consequently, the determinant is also **alternating**, meaning that if two inputs r_i and r_j are equal then $\det(r_1, \dots, r_n) = 0$.

- (3) The determinant is **normalized**, meaning that the standard basis has determinant 1,

$$\det(e_1, \dots, e_n) = 1.$$

Condition (1) does not say that $\det(\alpha A + \alpha' A') = \alpha \det(A) + \alpha' \det(A')$ for scalars α , α' and square matrices A , A' . Especially, the determinant is not additive,

$$\det(A + B) \text{ is in general not } \det(A) + \det(B). \quad (3.5)$$

What the condition does say is that if all rows but one of a square matrix are held fixed, then the determinant of the matrix varies linearly as a function of the one row.

We will prove the following theorem in the next section.

Theorem 3.5.1 (Existence and Uniqueness of the Determinant). *One, and only one, multilinear skew-symmetric normalized function from the n -fold product of \mathbf{R}^n to \mathbf{R} exists. This function is the determinant,*

$$\det : \mathbf{R}^n \times \dots \times \mathbf{R}^n \longrightarrow \mathbf{R}.$$

Furthermore, all multilinear skew-symmetric functions from the n -fold product of \mathbf{R}^n to \mathbf{R} are scalar multiples of the determinant. That is, any multilinear skew-symmetric function $\delta : \mathbf{R}^n \times \dots \times \mathbf{R}^n \longrightarrow \mathbf{R}$ is

$$\delta = c \cdot \det \quad \text{where} \quad c = \delta(e_1, \dots, e_n).$$

In more structural language, Theorem 3.5.1 says that the multilinear skew-symmetric functions from the n -fold product of \mathbf{R}^n to \mathbf{R} form a 1-dimensional vector space over \mathbf{R} , and $\{\det\}$ is a basis.

The reader may object that even if the conditions of multilinearity, skew-symmetry, and normalization are grammatically natural, they are conceptually opaque. Indeed they reflect considerable hindsight, since the idea of a determinant originally emerged from explicit calculations. But again, the payoff is that characterizing the determinant rather than constructing it illuminates its many useful properties. The rest of the section can be viewed as an amplification of this idea.

For one quick application of the existence of the determinant, consider the standard basis of \mathbf{R}^n taken in order,

$$(e_1, \dots, e_n).$$

Suppose that some succession of m pair-exchanges of the vectors in this ordered n -tuple has no net effect, i.e., after the m pair-exchanges, the vectors are

back in their original order. By skew-symmetry each pair-exchange negates the determinant, and so after all m pair-exchanges the net result is

$$(-1)^m \det(e_1, \dots, e_n) = \det(e_1, \dots, e_n).$$

Since \det is normalized, it follows that $(-1)^m = 1$, i.e., m is even. That is, no odd number of pair-exchanges can leave an ordered n -tuple in its initial order. Consequently, if two different sequences of pair-exchanges have the same net effect then their lengths are both odd or both even—this is because running one sequence forwards and then the other back has no net effect and hence comes to an even number of moves. In other words, although a net rearrangement of an n -tuple does not determine a unique succession of pair-exchanges to bring it about, or even a unique number of such exchanges, it does determine the parity of any such number: the net rearrangement requires an odd number of exchanges, or it requires an even number. (For reasons related to this, an old puzzle involving fifteen squares that slide in a 4-by-4 grid can be made unsolvable by popping two pieces out and exchanging them.) The fact that the parity of a rearrangement is well defined may be easy to believe, perhaps so easy that the need for a proof is hard to see. But in fact any proof requires some auxiliary measuring tool. The determinant measures much more than necessary, giving the proof as a byproduct. The determinant's uniqueness is irrelevant here. On the other hand, the uniqueness will play a significant role in chapter 8.

A weightier example follows.

Theorem 3.5.2 (The Determinant is Multiplicative). *For all matrices $A, B \in M_n(\mathbf{R})$,*

$$\det(AB) = \det(A) \det(B).$$

In particular, if A is invertible then the determinant of the matrix inverse is the scalar inverse of the determinant,

$$\det(A^{-1}) = (\det(A))^{-1}.$$

Multilinearity says that the determinant behaves well additively and scalar-multiplicatively as a function of each of n vectors, while (3.5) says that the determinant does not behave well additively as a function of one matrix. Theorem 3.5.2 says that the determinant behaves perfectly well multiplicatively as a function of one matrix. Also, the theorem tacitly says that if A is invertible then $\det(A)$ is nonzero. Soon we will establish the converse as well.

Proof. Let $B \in M_n(\mathbf{R})$ be fixed. Consider the function

$$\delta : M_n(\mathbf{R}) \longrightarrow \mathbf{R}, \quad \delta(A) = \det(AB).$$

As a function of the rows of A , δ is the determinant of the rows of AB ,

$$\delta : \mathbf{R}^n \times \cdots \times \mathbf{R}^n \longrightarrow \mathbf{R}, \quad \delta(r_1, \dots, r_n) = \det(r_1 B, \dots, r_n B).$$

The function δ is multilinear and skew-symmetric. To show multilinearity, compute (using the definition of δ , properties of vector–matrix algebra, the multilinearity of \det , and the definition of δ again),

$$\begin{aligned}\delta(r_1, \dots, \alpha r_k + \alpha' r'_k, \dots, r_n) &= \det(r_1 B, \dots, (\alpha r_k + \alpha' r'_k) B, \dots, r_n B) \\ &= \det(r_1 B, \dots, \alpha r_k B + \alpha' r'_k B, \dots, r_n B) \\ &= \alpha \det(r_1 B, \dots, r_k B, \dots, r_n B) \\ &\quad + \alpha' \det(r_1 B, \dots, r'_k B, \dots, r_n B) \\ &= \alpha \delta(r_1, \dots, r_k, \dots, r_n) \\ &\quad + \alpha' \delta(r_1, \dots, r'_k, \dots, r_n).\end{aligned}$$

To show skew-symmetry, take two distinct indices $i, j \in \{1, \dots, n\}$ and compute similarly,

$$\begin{aligned}\delta(r_1, \dots, r_j, \dots, r_i, \dots, r_n) &= \det(r_1 B, \dots, r_j B, \dots, r_i B, \dots, r_n B) \\ &= -\det(r_1 B, \dots, r_i B, \dots, r_j B, \dots, r_n B) \\ &= -\delta(r_1, \dots, r_i, \dots, r_j, \dots, r_n).\end{aligned}$$

Also compute that

$$\delta(e_1, \dots, e_n) = \det(e_1 B, \dots, e_n B) = \det(B).$$

It follows from Theorem 3.5.1 that $\delta(A) = \det(B) \det(A)$, and this is the desired main result $\det(AB) = \det(A) \det(B)$ of the theorem. Finally, if A is invertible then

$$\det(A) \det(A^{-1}) = \det(AA^{-1}) = \det(I) = 1.$$

That is, $\det(A^{-1}) = (\det(A))^{-1}$. This completes the proof. \square

More facts about the determinant are immediate consequences of its characterizing properties.

Proposition 3.5.3 (Determinants of Elementary and Echelon Matrices).

- (1) $\det(R_{i;j,a}) = 1$ for all $i, j \in \{1, \dots, n\}$ ($i \neq j$) and $a \in \mathbf{R}$.
- (2) $\det(S_{i,a}) = a$ for all $i \in \{1, \dots, n\}$ and nonzero $a \in \mathbf{R}$.
- (3) $\det(T_{i;j}) = -1$ for all $i, j \in \{1, \dots, n\}$ ($i \neq j$).
- (4) If E is n -by- n echelon then

$$\det(E) = \begin{cases} 1 & \text{if } E = I, \\ 0 & \text{otherwise.} \end{cases}$$

Proof. (1) Compute,

$$\begin{aligned}\det(R_{i,j,a}) &= \det(e_1, \dots, e_i + ae_j, \dots, e_j, \dots, e_n) \\ &= \det(e_1, \dots, e_i, \dots, e_j, \dots, e_n) + a \det(e_1, \dots, e_j, \dots, e_j, \dots, e_n) \\ &= 1 + 0 = 1.\end{aligned}$$

(2) and (3) are similar. For (4), if $E = I$ then $\det(E) = 1$ since the determinant is normalized. Otherwise the bottom row of E is $\mathbf{0}$, and since a linear function takes $\mathbf{0}$ to 0, it follows that $\det(E) = 0$. \square

For one consequence of Theorem 3.5.2 and Proposition 3.5.3, recall that every matrix $A \in M_n(\mathbf{R})$ has a transpose matrix A^t , obtained by flipping A about its Northwest–Southeast diagonal. The next theorem (whose proof is exercise 3.5.3) says that all statements about the determinant as a function of the rows of A also apply to the columns. This fact will be used without comment from now on. In particular, $\det(A)$ is the unique multilinear skew-symmetric normalized function of the columns of A .

Theorem 3.5.4 (Determinant and Transpose). *For all $A \in M_n(\mathbf{R})$, $\det(A^t) = \det(A)$.*

A far more important consequence of Theorem 3.5.2 and Proposition 3.5.3 is one of the main results of this chapter. Recall that any matrix A row-reduces as

$$R_1 \cdots R_N A = E$$

where the R_k are elementary, E is echelon, and A is invertible if and only if $E = I$. Since the determinant is multiplicative,

$$\det(R_1) \cdots \det(R_N) \det(A) = \det(E).$$

But each $\det(R_k)$ is nonzero, and $\det(E)$ is 1 if $E = I$ and 0 otherwise, so this gives the algebraic significance of the determinant:

Theorem 3.5.5 (Linear Invertibility Theorem). *The matrix $A \in M_n(\mathbf{R})$ is invertible if and only if $\det(A) \neq 0$.*

That is, the zeroness or nonzeroness of the determinant says whether the matrix is invertible. Once the existence and uniqueness of the determinant are established in the next section, we will continue to use the determinant properties to interpret the magnitude and the sign of the determinant as well.

Exercises

3.5.1. Consider a scalar-valued function of pairs of vectors,

$$\text{ip} : \mathbf{R}^n \times \mathbf{R}^n \longrightarrow \mathbf{R},$$

satisfying the following three properties.

(1) The function is bilinear,

$$\begin{aligned}\text{ip}(\alpha x + \alpha' x', y) &= \alpha \text{ip}(x, y) + \alpha' \text{ip}(x', y), \\ \text{ip}(x, \beta y + \beta' y') &= \beta \text{ip}(x, y) + \beta' \text{ip}(x, y')\end{aligned}$$

for all $\alpha, \alpha', \beta, \beta' \in \mathbf{R}$ and $x, x', y, y' \in \mathbf{R}^n$.

(2) The function is symmetric,

$$\text{ip}(x, y) = \text{ip}(y, x) \quad \text{for all } x, y \in \mathbf{R}^n.$$

(3) The function is normalized,

$$\text{ip}(e_i, e_j) = \delta_{ij} \quad \text{for all } i, j \in \{1, \dots, n\}.$$

(The Kronecker delta δ_{ij} was defined in section 2.2.)

Compute that this function, if it exists at all, must be the inner product. On the other hand, we already know that the inner product has these three properties, so this exercise has shown that it is characterized by them.

3.5.2. Let $f : \mathbf{R}^n \times \dots \times \mathbf{R}^n \rightarrow \mathbf{R}$ be a multilinear skew-symmetric function, and let c be any real number. Show that the function cf is again multilinear and skew-symmetric.

3.5.3. This exercise shows that for any matrix A , $\det(A^t) = \det(A)$.

(a) Show that $\det(R^t) = \det(R)$ for any elementary matrix R . (That is, R can be a recombine matrix, a scale matrix, or a transposition matrix.)

(b) If E is an echelon matrix then either $E = I$ or the bottom row of E is $\mathbf{0}$. In either case, show that $\det(E^t) = \det(E)$. (For the case $E \neq I$, we know that E is not invertible. What is $E^t e_n$, and what does this say about the invertibility of E^t ?)

(c) Use the formula $(MN)^t = N^t M^t$, Theorem 3.5.2, and Proposition 3.5.3 to show that $\det(A^t) = \det(A)$ for all $A \in M_n(\mathbf{R})$.

3.5.4. The square matrix A is **orthogonal** if $A^t A = I$. Show that if A is orthogonal then $\det(A) = \pm 1$. Give an example with determinant -1 .

3.5.5. The matrix A is **skew-symmetric** if $A^t = -A$. Show that if A is n -by- n skew-symmetric with n odd then $\det(A) = 0$.

3.6 The Determinant: Uniqueness and Existence

Recall that Theorem 3.5.1 asserts that exactly one multilinear skew-symmetric normalized function from the n -fold product of \mathbf{R}^n to \mathbf{R} exists. That is, a unique determinant exists.

We warm up for the proof of the theorem by using the three defining conditions of the determinant to show that only one formula is possible for the determinant of a general 2-by-2 matrix,

$$A = \begin{bmatrix} a & b \\ c & d \end{bmatrix}.$$

The first row of this matrix is

$$r_1 = (a, b) = a(1, 0) + b(0, 1) = ae_1 + be_2,$$

and similarly its second row is $r_2 = ce_1 + de_2$. Thus, since we view the determinant as a function of rows, its determinant must be

$$\det(A) = \det(r_1, r_2) = \det(ae_1 + be_2, ce_1 + de_2).$$

Since \det is linear in its first vector variable, this expands to

$$\det(ae_1 + be_2, ce_1 + de_2) = a \det(e_1, ce_1 + de_2) + b \det(e_2, ce_1 + de_2),$$

and since \det is also linear in its second vector variable, this expands further,

$$\begin{aligned} a \det(e_1, ce_1 + de_2) + b \det(e_2, ce_1 + de_2) \\ = ac \det(e_1, e_1) + ad \det(e_1, e_2) \\ + bc \det(e_2, e_1) + bd \det(e_2, e_2). \end{aligned}$$

But since \det is skew-symmetric and alternating, the expanded expression simplifies considerably,

$$\begin{aligned} ac \det(e_1, e_1) + ad \det(e_1, e_2) + bc \det(e_2, e_1) + bd \det(e_2, e_2) \\ = (ad - bc) \det(e_1, e_2). \end{aligned}$$

And finally, since \det is normalized, we have found the only possible formula for the determinant of a 2-by-2 matrix,

$$\det(A) = ad - bc.$$

All three characterizing properties of the determinant were required to derive this formula.

(As a brief digression, the reader can use the matrix inversion algorithm from section 3.3 to verify that the 2-by-2 matrix A is invertible if and only if $ad - bc$ is nonzero, showing that the formula for the 2-by-2 determinant arises from considerations of invertibility as well as from our three conditions. However, the argument requires cases, e.g., $a \neq 0$ or $a = 0$, making this approach uninviting for larger matrices.)

Returning to the main line of exposition, nothing here has yet shown that a determinant function exists at all for 2-by-2 matrices. What it has shown is that there is only one possibility,

$$\det((a, b), (c, d)) = ad - bc.$$

But now that we have the only possible formula, checking that indeed it satisfies the desired properties is purely mechanical. For example, to verify linearity in the first vector variable, compute

$$\begin{aligned} \det((a, b) + (a', b'), (c, d)) &= \det((a + a', b + b'), (c, d)) \\ &= (a + a')d - (b + b')c \\ &= (ad - bc) + (a'd - b'c) \\ &= \det((a, b), (c, d)) + \det((a', b'), (c, d)). \end{aligned}$$

For skew-symmetry,

$$\det((c, d), (a, b)) = cb - da = -(ad - bc) = -\det((a, b), (c, d)).$$

And for normalization,

$$\det(1, 0), (0, 1) = 1 \cdot 1 - 0 \cdot 0 = 1.$$

We should also verify linearity in the second vector variable, but this no longer requires the defining formula. Instead, since the formula is skew-symmetric and is linear in the first variable,

$$\begin{aligned} \det(r_1, r_2 + r'_2) &= -\det(r_2 + r'_2, r_1) \\ &= -(\det(r_2, r_1) + \det(r'_2, r_1)) \\ &= -(-\det(r_1, r_2)) - \det(r_1, r'_2) \\ &= \det(r_1, r_2) + \det(r_1, r'_2). \end{aligned}$$

This little trick illustrates the value of thinking in general terms: a slight modification, inserting a few occurrences of “...” and replacing the subscripts 1 and 2 by i and j , shows that for any n , the three required conditions for the determinant are redundant—linearity in one vector variable combines with skew-symmetry to ensure linearity in all the vector variables.

One can similarly show that for a 1-by-1 matrix,

$$A = [a],$$

the only possible formula for its determinant is

$$\det(A) = a,$$

and that indeed this works. The result is perhaps silly, but the exercise of working through a piece of language and logic in the simplest instance can help one to understand its more elaborate cases. As another exercise, the same techniques show that the only possible formula for a 3-by-3 determinant is

$$\det \begin{bmatrix} a & b & c \\ d & e & f \\ g & h & k \end{bmatrix} = aek + bfg + cdh - afh - bdk - ceg.$$

This formula is complicated enough that we should rethink it in a more systematic way before verifying that it has the desired properties. And we may as well generalize it to arbitrary n in the process. Here are some observations about the 3-by-3 formula:

- It is a sum of 3-fold products of matrix entries.
- Every 3-fold product contains one element from each row of the matrix.
- Every 3-fold product also contains one element from each column of the matrix. So every 3-fold product arises from the positions of three rooks that don't threaten each other on a 3-by-3 chessboard.
- Every 3-fold product comes weighted by a "+" or a "-".

Similar observations apply to the 1-by-1 and 2-by-2 formulas. Our general formula should encode them. Making it do so is partly a matter of notation, but also an idea is needed to describe the appropriate distribution of plus signs and minus signs among the terms. The following language provides all of this.

Definition 3.6.1 (Permutation). A permutation of $\{1, 2, \dots, n\}$ is a vector

$$\pi = (\pi(1), \pi(2), \dots, \pi(n))$$

whose entries are $\{1, 2, \dots, n\}$, each appearing once, in any order. An **inversion** in the permutation π is a pair of entries with the larger one to the left. The **sign** of the permutation π , written $(-1)^\pi$, is -1 raised to the number of inversions in π . The set of permutations of $\{1, \dots, n\}$ is denoted S_n .

Examples are the permutations $\pi = (1, 2, 3, \dots, n)$, $\sigma = (2, 1, 3, \dots, n)$, and $\tau = (5, 4, 3, 2, 1)$ (here $n = 5$). In these examples π has no inversions, σ has one, and τ has ten. Thus $(-1)^\pi = 1$, $(-1)^\sigma = -1$, and $(-1)^\tau = 1$. In general, the sign of a permutation with an even number of inversions is 1 and the sign of a permutation with an odd number of inversions is -1 . There are $n!$ permutations of $\{1, 2, \dots, n\}$; that is, the set S_n contains $n!$ elements.

As advertised, permutations and their signs provide the notation for the only possible n -by- n determinant formula. Consider any n vectors

$$r_1 = \sum_{i=1}^n a_{1i}e_i, \quad r_2 = \sum_{j=1}^n a_{2j}e_j, \quad \dots, \quad r_n = \sum_{p=1}^n a_{np}e_p.$$

By multilinearity, any multilinear function δ (if it exists at all) must satisfy

$$\begin{aligned}\delta(r_1, r_2, \dots, r_n) &= \delta \left(\sum_{i=1}^n a_{1i} e_i, \sum_{j=1}^n a_{2j} e_j, \dots, \sum_{p=1}^n a_{np} e_p \right) \\ &= \sum_{i=1}^n \sum_{j=1}^n \cdots \sum_{p=1}^n a_{1i} a_{2j} \cdots a_{np} \delta(e_i, e_j, \dots, e_p).\end{aligned}$$

If δ is also alternating then for any $i, j, \dots, p \in \{1, \dots, n\}$,

$$\delta(e_i, e_j, \dots, e_p) = 0 \quad \text{if any two subscripts agree.}$$

Thus we may sum only over permutations,

$$\delta(r_1, r_2, \dots, r_n) = \sum_{(i,j,\dots,p) \in S_n} a_{1i} a_{2j} \cdots a_{np} \det(e_i, e_j, \dots, e_p).$$

Consider any permutation $\pi = (i, j, \dots, p)$. Suppose that π contains an inversion, i.e., two elements are out of order. Then necessarily two elements in adjacent slots are out of order. (For example, if $i > p$ then either $i > j$, giving adjacent elements out of order as desired; or $j > i > p$, so that j and p are an out of order pair in closer slots than i and p , and so on.) If a permutation contains any inversions, then exchanging a suitable adjacent pair decreases the number of inversions by one, negating the sign of the permutation, while exchanging the corresponding two input vectors negates the determinant. Repeating this process until the permutation has no remaining inversions shows that

$$\delta(e_i, e_j, \dots, e_p) = (-1)^\pi \delta(e_1, e_2, \dots, e_n).$$

That is, the only possible formula for a multilinear skew-symmetric function δ is

$$\delta(r_1, r_2, \dots, r_n) = \sum_{\pi=(i,j,\dots,p)} (-1)^\pi a_{1i} a_{2j} \cdots a_{np} \cdot c$$

where

$$c = \delta(e_1, \dots, e_n).$$

Especially, the only possible formula for a multilinear skew-symmetric normalized function is

$$\det(r_1, r_2, \dots, r_n) = \sum_{\pi=(i,j,\dots,p)} (-1)^\pi a_{1i} a_{2j} \cdots a_{np}.$$

Definition 3.6.2 (Determinant). *The determinant function,*

$$\det : M_n(\mathbf{R}) \longrightarrow \mathbf{R},$$

is defined as follows. For any $A \in M_n(\mathbf{R})$ with entries (a_{ij}) ,

$$\det(A) = \sum_{\pi \in S_n} (-1)^\pi a_{1\pi(1)} a_{2\pi(2)} \cdots a_{n\pi(n)}.$$

The formula in the definition is indeed the formula computed a moment ago, since for any permutation $\pi = (i, j, \dots, p) \in S_n$ we have $\pi(1) = i$, $\pi(2) = j$, \dots , $\pi(n) = p$.

As an exercise to clarify the formula, we use it to reproduce the 3-by-3 determinant. Each permutation in S_3 determines a rook-placement, and the sign of the permutation is the parity of the number of Northeast–Southwest segments joining any two of its rooks. For example, the permutation $(2, 3, 1)$ specifies that the rooks in the top, middle, and bottom rows are respectively in columns 2, 3, and 1, and the sign is positive since there are two Northeast–Southwest segments. (See figure 3.9.) The following table lists each permutation in S_3 followed by the corresponding term in the determinant formula. For each permutation, the term is its sign times the product of the three matrix entries where its rooks are placed.

π	$(-1)^\pi a_{1\pi(1)} a_{2\pi(2)} a_{3\pi(3)}$
$(1, 2, 3)$	$ae k$
$(1, 3, 2)$	$-af h$
$(2, 1, 3)$	$-bd k$
$(2, 3, 1)$	$bf g$
$(3, 1, 2)$	$cd h$
$(3, 2, 1)$	$-ce g$

The sum of the right column entries is the anticipated formula from before,

$$\det \begin{bmatrix} a & b & c \\ d & e & f \\ g & h & k \end{bmatrix} = ae k + bf g + cd h - af h - bd k - ce g.$$

The same procedure also reproduces the 2-by-2 determinant as well,

$$\det \begin{bmatrix} a & b \\ c & d \end{bmatrix} = ad - bc,$$

and even the silly 1-by-1 formula $\det[a] = a$. The 2-by-2 and 3-by-3 cases are worth memorizing. They can be visualized as adding the products along Northwest–Southeast diagonals of the matrix and then subtracting the products along Southwest–Northeast diagonals, where the word “diagonal” connotes wraparound in the 3-by-3 case. (See figure 3.10.) But be aware that this pattern of the determinant as the Northwest–Southeast diagonals minus the Southwest–Northeast diagonals is valid only for $n = 2$ and $n = 3$.

This completes the program of the second bullet at the beginning of the previous section, finding the only possible formula (the one in Definition 3.6.2) that could satisfy the three desired determinant properties. We don’t yet know that it does, just that it is the only formula that could. That is, we have now proved the uniqueness but not yet the existence of the determinant in Theorem 3.5.1.




Figure 3.9. The rook placement for (2, 3, 1), showing the two inversions




Figure 3.10. The 3-by-3 determinant

The first bullet tells us to prove the existence by verifying that the computed determinant formula indeed does satisfy the three stipulated determinant properties. Similarly to the 2-by-2 case, this is a mechanical exercise. The impediments are purely notational, but the notation is admittedly cumbersome, and so the reader is encouraged to skim the next proof.

Proposition 3.6.3 (Properties of the Determinant).

- (1) *The determinant is linear as a function of each row of A .*
- (2) *The determinant is skew-symmetric as a function of the rows of A .*
- (3) *The determinant is normalized.*

Proof. (1) If A has rows $r_i = (a_{i1}, \dots, a_{in})$ except that its k th row is the linear combination $\alpha r_k + \alpha' r'_k$ where $r_k = (a_{k1}, \dots, a_{kn})$ and $r'_k = (a'_{k1}, \dots, a'_{kn})$, then its (i, j) th entry is

$$\begin{cases} a_{ij} & \text{if } i \neq k, \\ \alpha a_{kj} + \alpha' a'_{kj} & \text{if } i = k. \end{cases}$$

Thus

$$\begin{aligned}
 & \det(r_1, \dots, \alpha r_k + \alpha' r'_k, \dots, r_n) \\
 &= \sum_{\pi \in S_n} (-1)^\pi a_{1\pi(1)} \cdots (\alpha a_{k\pi(j)} + \alpha' a'_{k\pi(j)}) \cdots a_{n\pi(n)} \\
 &= \alpha \sum_{\pi \in S_n} (-1)^\pi a_{1\pi(1)} \cdots a_{k\pi(j)} \cdots a_{n\pi(n)} \\
 &\quad + \alpha' \sum_{\pi \in S_n} (-1)^\pi a_{1\pi(1)} \cdots a'_{k\pi(j)} \cdots a_{n\pi(n)} \\
 &= \alpha \det(r_1, \dots, r_k, \dots, r_n) + \alpha' \det(r_1, \dots, r'_k, \dots, r_n),
 \end{aligned}$$

as desired.

(2) Let A have rows r_1, \dots, r_n where $r_i = (a_{i1}, \dots, a_{in})$. Suppose that rows k and $k + 1$ are exchanged. The resulting matrix has (i, j) th entry

$$\begin{cases} a_{ij} & \text{if } i \notin \{k, k + 1\}, \\ a_{k+1,j} & \text{if } i = k, \\ a_{kj} & \text{if } i = k + 1. \end{cases}$$

For each permutation $\pi \in S_n$, define a companion permutation π' by exchanging the k th and $(k + 1)$ st entries,

$$\pi' = (\pi(1), \dots, \pi(k + 1), \pi(k), \dots, \pi(n)).$$

Thus $\pi'(k) = \pi(k + 1)$, $\pi'(k + 1) = \pi(k)$, and $\pi'(i) = \pi(i)$ for all other i . As π varies through S_n , so does π' , and for each π we have the relation $(-1)^{\pi'} = -(-1)^\pi$ (exercise 3.6.5). The defining formula of the determinant gives

$$\begin{aligned}
 & \det(r_1, \dots, r_{k+1}, r_k, \dots, r_n) \\
 &= \sum_{\pi} (-1)^\pi a_{1\pi(1)} \cdots a_{k+1,\pi(k)} a_{k\pi(k+1)} \cdots a_{n\pi(n)} \\
 &= - \sum_{\pi'} (-1)^{\pi'} a_{1\pi'(1)} \cdots a_{k+1,\pi'(k+1)} a_{k\pi'(k)} \cdots a_{n\pi'(n)} \\
 &= - \det(r_1, \dots, r_k, r_{k+1}, \dots, r_n).
 \end{aligned}$$

This establishes the result when adjacent rows of A are exchanged. To exchange rows k and ℓ in A where $\ell > k$, carry out the following adjacent row exchanges:

rows k and $k + 1$,	k and $k + 1$.
$k + 1$ and $k + 2$,	$k + 1$ and $k + 2$,
\dots ,	\dots ,
$\ell - 2$ and $\ell - 1$,	$\ell - 2$ and $\ell - 1$,
$\ell - 1$ and ℓ ,	

This process trickles the k th row down to the ℓ th and then bubbles the ℓ th row back up to the k th, bobbing each row in between them up one position and then back down. And the display shows that the process carries out an odd number of exchanges (all but the bottom one come in pairs), each of which negates the determinant.

(3) This is left to the reader. □

So a unique determinant function with the stipulated behavior exists. And we have seen that any multilinear skew-symmetric function must be a scalar multiple of the determinant. The last comment necessary to complete the proof of Theorem 3.5.1 is that since the determinant is multilinear and skew-symmetric, so are its scalar multiples. This was shown in exercise 3.5.2.

The reader is encouraged to contemplate how tortuous it would have been to prove the various facts about the determinant in the previous section by using the unwieldy determinant formula, with its $n!$ terms.

A type of matrix that has an easily calculable determinant is a **triangular** matrix, meaning a matrix all of whose subdiagonal entries are 0 or all of whose superdiagonal entries are 0. (Lower triangular matrices have already been introduced in exercise 3.3.12.) For example, the matrices

$$\begin{bmatrix} a_{11} & a_{12} & a_{13} \\ 0 & a_{22} & a_{23} \\ 0 & 0 & a_{33} \end{bmatrix} \quad \text{and} \quad \begin{bmatrix} a_{11} & 0 & 0 \\ a_{21} & a_{22} & 0 \\ a_{31} & a_{32} & a_{33} \end{bmatrix}$$

are triangular. The determinant formula immediately shows that the determinant of a triangular matrix is the product of the diagonal entries (exercise 3.6.7). This fact should be cited freely to save time.

An algorithm for computing $\det(A)$ for any $A \in M_n(\mathbf{R})$ is now at hand. Algebraically, the idea is that if

$$P_1 A P_2 = T$$

where P_1 and P_2 are products of elementary matrices and T is a triangular matrix (here T does not denote a transposition matrix), then since the determinant is multiplicative,

$$\det(A) = \det(P_1)^{-1} \det(T) \det(P_2)^{-1}.$$

Multiplying A by P_2 on the right carries out a sequence of column operations on A , just as multiplying A by P_1 on the left carries out row operations. Recall that the determinants of the elementary matrices are

$$\begin{aligned} \det(R_{i;j,a}) &= 1, \\ \det(S_{i,a}) &= a, \\ \det(T_{i;j}) &= -1. \end{aligned}$$

Procedurally, this all plays out as follows.

Proposition 3.6.4 (Determinant Algorithm). *Given $A \in M_n(\mathbf{R})$, use row and column operations—recombines, scales, transpositions—to reduce A to a triangular matrix T . Then $\det(A)$ is $\det(T)$ times the reciprocal of each scale factor and times -1 for each transposition.*

The only role that the determinant formula (as compared to the determinant properties) played in obtaining this algorithm is that it gave the determinant of a triangular matrix easily.

For example, the matrix

$$A = \begin{bmatrix} 1/0! & 1/1! & 1/2! & 1/3! \\ 1/1! & 1/2! & 1/3! & 1/4! \\ 1/2! & 1/3! & 1/4! & 1/5! \\ 1/3! & 1/4! & 1/5! & 1/6! \end{bmatrix}$$

becomes, after scaling the first row by $3!$, the second row by $4!$, the third row by $5!$, and the fourth row by $6!$,

$$B = \begin{bmatrix} 6 & 6 & 3 & 1 \\ 24 & 12 & 4 & 1 \\ 60 & 20 & 5 & 1 \\ 120 & 30 & 6 & 1 \end{bmatrix}.$$

Subtract the first row from each of the others to get

$$C = \begin{bmatrix} 6 & 6 & 3 & 1 \\ 18 & 6 & 1 & 0 \\ 54 & 14 & 2 & 0 \\ 114 & 24 & 3 & 0 \end{bmatrix},$$

and then scale the third row by $1/2$ and the fourth row by $1/3$, yielding

$$D = \begin{bmatrix} 6 & 6 & 3 & 1 \\ 18 & 6 & 1 & 0 \\ 27 & 7 & 1 & 0 \\ 38 & 8 & 1 & 0 \end{bmatrix}.$$

Next subtract the second row from the third row and the fourth rows, and scale the fourth row by $1/2$ to get

$$E = \begin{bmatrix} 6 & 6 & 3 & 1 \\ 18 & 6 & 1 & 0 \\ 9 & 1 & 0 & 0 \\ 10 & 1 & 0 & 0 \end{bmatrix}.$$

Subtract the third row from the fourth, transpose the first and fourth columns, and transpose the second and third columns, leading to

$$T = \begin{bmatrix} 1 & 3 & 6 & 6 \\ 0 & 1 & 6 & 18 \\ 0 & 0 & 1 & 9 \\ 0 & 0 & 0 & 1 \end{bmatrix}.$$

This triangular matrix has determinant 1, and so according to the algorithm,

$$\det(A) = \frac{2 \cdot 3 \cdot 2}{6! 5! 4! 3!} = \frac{1}{1036800}.$$

In the following exercises, feel free to use the determinant properties and the determinant formula in whatever combined way gives you the least work.

Exercises

3.6.1. Use the three desired determinant properties to derive the formulas in the section for 1-by-1 and 3-by-3 determinant. Verify that the 1-by-1 formula satisfies the properties.

3.6.2. For each permutation, count the inversions and compute the sign: $(2, 3, 4, 1)$, $(3, 4, 1, 2)$, $(5, 1, 4, 2, 3)$.

3.6.3. Explain why there are $n!$ permutations of $\{1, \dots, n\}$.

3.6.4. Define the permutation $\mu = (n, n-1, n-2, \dots, 1) \in S_n$. Show that μ has $(n-1)n/2$ inversions and that

$$(-1)^\mu = \begin{cases} 1 & \text{if } n \text{ has the form } 4k \text{ or } 4k+1 \text{ } (k \in \mathbf{Z}), \\ -1 & \text{otherwise.} \end{cases}$$

3.6.5. Explain why $(-1)^{\pi'} = -(-1)^\pi$ in the proof of part (2) of Proposition 3.6.3.

3.6.6. Use the defining formula of the determinant to reproduce the result that $\det(I_n) = 1$.

3.6.7. Prove that the determinant of a triangular matrix is the product of its diagonal entries.

3.6.8. Calculate the determinants of the following matrices:

$$\begin{bmatrix} 4 & 3 & -1 & 2 \\ 0 & 1 & 2 & 3 \\ 1 & 0 & 4 & 1 \\ 2 & 0 & 3 & 0 \end{bmatrix}, \quad \begin{bmatrix} 1 & -1 & 2 & 3 \\ 2 & 2 & 0 & 2 \\ 4 & 1 & -1 & -1 \\ 1 & 2 & 3 & 0 \end{bmatrix}.$$

3.6.9. Show that the **Vandermonde matrix**,

$$\begin{bmatrix} 1 & a & a^2 \\ 1 & b & b^2 \\ 1 & c & c^2 \end{bmatrix},$$

has determinant $(b-a)(c-a)(c-b)$. For what values of a, b, c is the Vandermonde matrix invertible? (The idea is to do the problem conceptually rather than writing out the determinant and then factoring it, so that the same ideas would work for larger matrices. The determinant formula shows that the determinant in the problem is a polynomial in a, b , and c . What is its degree in each variable? Why must it vanish if any two variables are equal? Once you have argued that the determinant is as claimed, don't forget to finish the problem.)

3.6.10. Consider the following n -by- n matrix based on Pascal's triangle:

$$A = \begin{bmatrix} 1 & 1 & 1 & 1 & \cdots & 1 \\ 1 & 2 & 3 & 4 & \cdots & n \\ 1 & 3 & 6 & 10 & \cdots & \frac{n(n+1)}{2} \\ 1 & 4 & 10 & 20 & \cdots & \cdot \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots \\ 1 & n & \frac{n(n+1)}{2} & \cdot & \cdots & \cdot \end{bmatrix}.$$

Find $\det(A)$. (Hint: Row and column reduce.)

3.7 An Explicit Formula for the Inverse

Consider an invertible linear mapping

$$T : \mathbf{R}^n \longrightarrow \mathbf{R}^n$$

having matrix

$$A \in M_n(\mathbf{R}).$$

In section 3.3 we discussed a process to invert A and thereby invert T . Now, with the determinant in hand, we can also write the inverse of A explicitly in closed form. Because the pending formula giving the inverse involves many determinants, it is hopelessly inefficient for computation. Nonetheless, it is of interest to us for a theoretical reason (the pending Corollary 3.7.3) that we will need in chapter 5.

Definition 3.7.1 (Classical Adjoint). Let $n \geq 2$ be an integer, and let $A \in M_n(\mathbf{R})$ be an n -by- n matrix. For any $i, j \in \{1, \dots, n\}$, let

$$A^{i,j} \in M_{n-1}(\mathbf{R})$$

be the $(n - 1)$ -by- $(n - 1)$ matrix obtained by deleting the i th row and the j th column of A . The **classical adjoint** of A is the n -by- n matrix whose (i, j) th entry is $(-1)^{i+j}$ times the determinant of $A^{j,i}$,

$$A^{\text{adj}} = [(-1)^{i+j} \det(A^{j,i})] \in M_n(\mathbf{R}).$$

The factor $(-1)^{i+j}$ in the formula produces an alternating checkerboard pattern of plus and minus signs, starting with a plus sign in the upper left corner of A^{adj} . Note that the (i, j) th entry of A^{adj} involves $A^{j,i}$ rather than $A^{i,j}$. For instance, in the 2-by-2 case,

$$\begin{bmatrix} a & b \\ c & d \end{bmatrix}^{\text{adj}} = \begin{bmatrix} d & -b \\ -c & a \end{bmatrix}.$$

Already for a 3-by-3 matrix the formula for the classical adjoint is daunting,

$$\begin{aligned} \begin{bmatrix} a & b & c \\ d & e & f \\ g & h & k \end{bmatrix}^{\text{adj}} &= \begin{bmatrix} \det \begin{bmatrix} e & f \\ h & k \end{bmatrix} & -\det \begin{bmatrix} b & c \\ h & k \end{bmatrix} & \det \begin{bmatrix} b & c \\ e & f \end{bmatrix} \\ -\det \begin{bmatrix} d & f \\ g & k \end{bmatrix} & \det \begin{bmatrix} a & c \\ g & k \end{bmatrix} & -\det \begin{bmatrix} a & c \\ d & f \end{bmatrix} \\ \det \begin{bmatrix} d & e \\ g & h \end{bmatrix} & -\det \begin{bmatrix} a & b \\ g & h \end{bmatrix} & \det \begin{bmatrix} a & b \\ d & e \end{bmatrix} \end{bmatrix} \\ &= \begin{bmatrix} ek - fh & ch - bk & bf - ce \\ fg - dk & ak - cg & cd - af \\ dh - eg & bg - ah & ae - bd \end{bmatrix}. \end{aligned}$$

Returning to the 2-by-2 case, where

$$A = \begin{bmatrix} a & b \\ c & d \end{bmatrix} \quad \text{and} \quad A^{\text{adj}} = \begin{bmatrix} d & -b \\ -c & a \end{bmatrix},$$

compute that

$$A A^{\text{adj}} = \begin{bmatrix} ad - bc & 0 \\ 0 & ad - bc \end{bmatrix} = (ad - bc) \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} = \det(A)I_2.$$

The same result holds in general:

Proposition 3.7.2 (Classical Adjoint Identity). *Let $n \geq 2$ be an integer, let $A \in M_n(\mathbf{R})$ be an n -by- n matrix, and let A^{adj} be its classical adjoint. Then*

$$A A^{\text{adj}} = \det(A)I_n.$$

Especially, if A is invertible then

$$A^{-1} = \frac{1}{\det(A)} A^{\text{adj}}.$$

The proof is purely formal but notationally tedious, and so we omit it. In the 2-by-2 case the proposition gives us a slogan:

To invert a 2-by-2 matrix, exchange the diagonal elements, negate the off-diagonal elements, and divide by the determinant.

Again, for $n > 2$ the explicit formula for the inverse is rarely of calculational use. We care about it for the following reason.

Corollary 3.7.3. *Let $A \in M_n(\mathbf{R})$ be an invertible n -by- n matrix. Then each entry of the inverse matrix A^{-1} is a continuous function of the entries of A .*

Proof. Specifically, the (i, j) th entry of A^{-1} is

$$(A^{-1})_{i,j} = (-1)^{i+j} \det(A^{j,i}) / \det(A).$$

This is a rational function (ratio of polynomials) of the entries of A , and as such it varies continuously in the entries of A so long as A remains invertible. \square

Exercise

3.7.1. Verify at least one diagonal entry and at least one off-diagonal entry in the formula $AA^{\text{adj}} = \det(A)I_n$ for $n = 3$.

3.8 Geometry of the Determinant: Volume

Consider a linear mapping from n -space back to n -space,

$$T : \mathbf{R}^n \longrightarrow \mathbf{R}^n.$$

This section discusses two ideas:

- The mapping T magnifies volume by a constant factor. (Here *volume* is a pan-dimensional term that in particular means *length* when $n = 1$, *area* when $n = 2$, and the usual notion of volume when $n = 3$.) That is, there is some number $t \geq 0$ such that if one takes a set,

$$\mathcal{E} \subset \mathbf{R}^n,$$

and passes it through the mapping to get another set,

$$T\mathcal{E} \subset \mathbf{R}^n,$$

then the set's volume is multiplied by t ,

$$\text{vol } T\mathcal{E} = t \cdot \text{vol } \mathcal{E}.$$

The magnification factor t depends on T but is independent of the set \mathcal{E} .

- Furthermore, if the matrix of T is A then the magnification factor associated to T is

$$t = |\det A|.$$

That is, the absolute value of $\det A$ has a geometric interpretation as the factor by which T magnifies volume.

(The geometric interpretation of the sign of $\det A$ will be discussed in the next section.)

An obstacle to pursuing these ideas is that we don't have a theory of volume in \mathbf{R}^n readily at hand. In fact, volume presents real difficulties. For instance, any notion of volume that has sensible properties can not apply to all sets; so either volume behaves unreasonably or some sets don't have well defined volumes at all. Here we have been tacitly assuming that volume does behave well and that the sets \mathcal{E} under consideration do have volumes. This section will investigate volume informally by considering how it ought to behave, stating assumptions as they arise and arriving only at a partial description. The resulting arguments will be heuristic, and the skeptical reader will see gaps in the reasoning. Volume will be discussed further in chapter 6, but a full treatment of the subject (properly called *measure*) is beyond the range of this text.

The standard basis vectors e_1, \dots, e_n in \mathbf{R}^n span the **unit box**,

$$\mathcal{B} = \{\alpha_1 e_1 + \dots + \alpha_n e_n : 0 \leq \alpha_1 \leq 1, \dots, 0 \leq \alpha_n \leq 1\}.$$

Thus *box* means *interval* when $n = 1$, *rectangle* when $n = 2$, and the usual notion of box when $n = 3$. Let p be a point in \mathbf{R}^n , let a_1, \dots, a_n be positive real numbers, and let \mathcal{B}' denote the box spanned by the vectors $a_1 e_1, \dots, a_n e_n$ and translated by p ,

$$\mathcal{B}' = \{\alpha_1 a_1 e_1 + \dots + \alpha_n a_n e_n + p : 0 \leq \alpha_1 \leq 1, \dots, 0 \leq \alpha_n \leq 1\}.$$

(See figure 3.11. The figures of this section are set in two dimensions, but the ideas are general and hence so are the figure captions.) A **face** of a box is the set of its points such that some particular α_i is held fixed at 0 or at 1 while the others vary. A box in \mathbf{R}^n has $2n$ faces.

A natural definition is that the unit box has unit volume,

$$\text{vol } \mathcal{B} = 1.$$

We assume that volume is unchanged by translation. Also, we assume that box volume is *finitely additive*, meaning that given finitely many boxes $\mathcal{B}_1, \dots, \mathcal{B}_M$ that are disjoint except possibly for shared faces or shared subsets of faces, the volume of their union is the sum of their volumes,

$$\text{vol } \bigcup_{i=1}^M \mathcal{B}_i = \sum_{i=1}^M \text{vol } \mathcal{B}_i. \quad (3.6)$$

And we assume that scaling any spanning vector of a box affects the box's volume continuously in the scaling factor. It follows that scaling any spanning vector of a box by a real number a magnifies the volume by $|a|$. To see this, first note that scaling a spanning vector by an integer ℓ creates $|\ell|$ abutting translated copies of the original box, and so the desired result follows in this case from finite additivity. A similar argument applies to scaling a spanning vector by a reciprocal integer $1/m$ ($m \neq 0$), since the original box is now $|m|$ copies of the scaled one. These two special cases show that the result holds for scaling a spanning vector by any rational number $r = \ell/m$. Finally, the continuity assumption extends the result from the rational numbers to the real numbers, since every real number is approached by a sequence of rational numbers. Since the volume of the unit box is normalized to 1, since volume is unchanged by translation, and since scaling any spanning vector of a box by a magnifies its volume by $|a|$, the volume of the general box is (recalling that a_1, \dots, a_n are assumed to be positive)

$$\text{vol } \mathcal{B}' = a_1 \cdots a_n.$$




Figure 3.11. Scaling and translating the unit box

A subset of \mathbf{R}^n that is well approximated by boxes plausibly has a volume. To be more specific, a subset \mathcal{E} of \mathbf{R}^n is well approximated by boxes if for any $\varepsilon > 0$ there exist boxes $\mathcal{B}_1, \dots, \mathcal{B}_N, \mathcal{B}_{N+1}, \dots, \mathcal{B}_M$ that are disjoint except possibly for shared faces, such that \mathcal{E} is contained between a partial union of the boxes and the full union,

$$\bigcup_{i=1}^N \mathcal{B}_i \subset \mathcal{E} \subset \bigcup_{i=1}^M \mathcal{B}_i, \quad (3.7)$$

and such that the boxes that complete the partial union to the full union have a small sum of volumes,

$$\sum_{i=N+1}^M \text{vol } \mathcal{B}_i < \varepsilon. \quad (3.8)$$

(See figure 3.12, where \mathcal{E} is an elliptical region, the boxes \mathcal{B}_1 through \mathcal{B}_N that it contains are dark, and the remaining boxes \mathcal{B}_{N+1} through \mathcal{B}_M are light.)

To see that \mathcal{E} should have a volume, note that the first containment of (3.7) says that a number at most big enough to serve as $\text{vol } \mathcal{E}$ (a lower bound) is $L = \text{vol } \bigcup_{i=1}^N \mathcal{B}_i$, and the second containment says that a number at least big enough (an upper bound) is $U = \text{vol } \bigcup_{i=1}^M \mathcal{B}_i$. By the finite additivity condition (3.6), the lower and upper bounds are $L = \sum_{i=1}^N \text{vol } \mathcal{B}_i$ and $U = \sum_{i=1}^M \text{vol } \mathcal{B}_i$. Thus they are close to each other by (3.8),

$$U - L = \sum_{i=N+1}^M \text{vol } \mathcal{B}_i < \varepsilon.$$

Since ε is arbitrarily small, the bounds should be squeezing down on a unique value that is the actual volume of \mathcal{E} , and so indeed \mathcal{E} should have a volume. For now this is only a plausibility argument, but it is essentially the idea of integration and it will be quantified in chapter 6.




Figure 3.12. Inner and outer approximation of \mathcal{E} by boxes

Any n vectors v_1, \dots, v_n in \mathbf{R}^n span an n -dimensional **parallelepiped**

$$\mathcal{P}(v_1, \dots, v_n) = \{\alpha_1 v_1 + \dots + \alpha_n v_n : 0 \leq \alpha_1 \leq 1, \dots, 0 \leq \alpha_n \leq 1\},$$

abbreviated to \mathcal{P} when the vectors are firmly fixed. Again the terminology is pan-dimensional, meaning in particular *interval*, *parallelogram*, and *parallelepiped* in the usual sense for $n = 1, 2, 3$. We will also consider translations of parallelepipeds away from the origin by offset vectors p ,

$$\mathcal{P}' = \mathcal{P} + p = \{v + p : v \in \mathcal{P}\}.$$

(See figure 3.13.) A face of a parallelepiped is the set of its points such that some particular α_i is held fixed at 0 or at 1 while the others vary. A parallelepiped in \mathbf{R}^n has $2n$ faces. Boxes are special cases of parallelepipeds. The methods of chapter 6 will show that parallelepipeds are well approximated by boxes, and so they have well defined volumes. We assume that parallelepiped volume is finitely additive, and we assume that any finite union of parallelepipeds each having volume zero again has volume zero.




Figure 3.13. Parallelepipeds

To begin the argument that the linear mapping $T : \mathbf{R}^n \rightarrow \mathbf{R}^n$ magnifies volume by a constant factor, we pass the unit box \mathcal{B} and the scaled translated box \mathcal{B}' from earlier in the section through T . The image of \mathcal{B} under T is a parallelepiped $T\mathcal{B}$ spanned by $T(e_1), \dots, T(e_n)$, and the image of \mathcal{B}' is a parallelepiped $T\mathcal{B}'$ spanned by $T(a_1e_1), \dots, T(a_n e_n)$ and translated by $T(p)$. (See figure 3.14.) Since $T(a_1e_1) = a_1T(e_1), \dots, T(a_n e_n) = a_nT(e_n)$, it follows that scaling the sides of $T\mathcal{B}$ by a_1, \dots, a_n and then translating the scaled parallelepiped by $T(p)$ gives $T\mathcal{B}'$. As for boxes, scaling any spanning vector of a parallelepiped by a real number a magnifies the volume by $|a|$, and so we have

$$\text{vol } T\mathcal{B}' = a_1 \cdots a_n \cdot \text{vol } T\mathcal{B}.$$

But recall that also

$$\text{vol } \mathcal{B}' = a_1 \cdots a_n.$$

The two displays combine to give

$$\frac{\text{vol } T\mathcal{B}'}{\text{vol } \mathcal{B}'} = \text{vol } T\mathcal{B}.$$

That is, the volume of the T -image of a box divided by the volume of the box is constant, regardless of the box's location or side lengths, the constant being the volume of $T\mathcal{B}$, the T -image of the unit box \mathcal{B} . Call this constant magnification factor t . That is,

$$\text{vol } T\mathcal{B}' = t \cdot \text{vol } \mathcal{B}' \quad \text{for all boxes } \mathcal{B}'. \quad (3.9)$$

We need one last preliminary result about volume. Again let \mathcal{E} be a subset of \mathbf{R}^n that is well approximated by boxes. Fix a linear mapping $T : \mathbf{R}^n \rightarrow \mathbf{R}^n$. Very similarly to the argument for \mathcal{E} , the set $T\mathcal{E}$ also should have a volume, because it is well approximated by parallelepipeds. Indeed, the set containments (3.7) are preserved under the linear mapping T ,

$$T \bigcup_{i=1}^N \mathcal{B}_i \subset T\mathcal{E} \subset T \bigcup_{i=1}^M \mathcal{B}_i.$$




Figure 3.14. Linear image of the unit box and of a scaled translated box

In general, the image of a union is the union of the images, so this rewrites as

$$\bigcup_{i=1}^N T\mathcal{B}_i \subset T\mathcal{E} \subset \bigcup_{i=1}^M T\mathcal{B}_i.$$

(See figure 3.15.) As before, numbers at most big enough and at least big enough for the volume of $T\mathcal{E}$ are

$$L = \text{vol} \bigcup_{i=1}^N T\mathcal{B}_i = \sum_{i=1}^N \text{vol} T\mathcal{B}_i, \quad U = \text{vol} \bigcup_{i=1}^M T\mathcal{B}_i = \sum_{i=1}^M \text{vol} T\mathcal{B}_i.$$

The only new wrinkle is that citing the finite additivity of parallelepiped volume here assumes that the parallelepipeds $T\mathcal{B}_i$ either inherit from the original boxes \mathcal{B}_i the property of being disjoint except possibly for shared faces, or they all have volume zero. The assumption is valid because if T is invertible then the inheritance holds, while if T is not invertible then we will see later in this section that the $T\mathcal{B}_i$ have volume zero as desired. With this point established, let t be the factor by which T magnifies box-volume. The previous display and (3.9) combine to show that the difference of the bounds is

$$U - L = \sum_{i=N+1}^M \text{vol} T\mathcal{B}_i = \sum_{i=N+1}^M t \cdot \text{vol} \mathcal{B}_i = t \cdot \sum_{i=N+1}^M \text{vol} \mathcal{B}_i \leq t\varepsilon.$$

The inequality is strict if $t > 0$, and it collapses to $U - L = 0$ if $t = 0$. In either case, since ε is arbitrarily small, the argument that $T\mathcal{E}$ should have a volume is the same as for \mathcal{E} .

To complete the argument that the linear mapping $T : \mathbf{R}^n \rightarrow \mathbf{R}^n$ magnifies volume by a constant factor, we argue that for any subset \mathcal{E} of \mathbf{R}^n that is well approximated by boxes, $\text{vol} T\mathcal{E}$ is t times the volume of \mathcal{E} . Let $V = \text{vol} \bigcup_{i=1}^N \mathcal{B}_i$. Then \mathcal{E} is contained between a set of volume V and a set of volume less than $V + \varepsilon$ (again see figure 3.12, where V is the shaded area and $V + \varepsilon$ is the total area), and $T\mathcal{E}$ is contained between a set of volume tV and a set of volume at most $t(V + \varepsilon)$ (again see figure 3.15, where tV is the shaded




Figure 3.15. Inner and outer approximation of $T\mathcal{E}$ by parallelepipeds

area and $t(V + \varepsilon)$ is the total area). Thus the volumes $\text{vol } \mathcal{E}$ and $\text{vol } T\mathcal{E}$ satisfy the condition

$$\frac{tV}{V + \varepsilon} \leq \frac{\text{vol } T\mathcal{E}}{\text{vol } \mathcal{E}} \leq \frac{t(V + \varepsilon)}{V}.$$

Since ε can be arbitrarily small, the left and right quantities in the display can be arbitrarily close to t , and so the only possible value for the quantity in the middle (which is independent of ε) is t . This gives the desired equality announced at the beginning of the section,

$$\text{vol } T\mathcal{E} = t \cdot \text{vol } \mathcal{E}.$$

In sum, subject to various assumptions about volume, T magnifies the volumes of all boxes and of all figures that are well approximated by boxes by the same factor, which we have denoted t .

Now we investigate the magnification factor t associated to the linear mapping T , with the goal of showing that it is $|\det A|$ where A is the matrix of T . As a first observation, if the linear mappings $S, T : \mathbf{R}^n \rightarrow \mathbf{R}^n$ magnify volume by s and t respectively, then ST magnifies volume by st . In other words, the magnification of linear mappings is multiplicative. Also, recall that the mapping T is simply multiplication by the matrix A . Since any matrix is a product of elementary matrices times an echelon matrix, we only need to study the magnification of multiplying by such matrices. Temporarily let $n = 2$.

The 2-by-2 recombine matrices take the form $R = \begin{bmatrix} 1 & a \\ 0 & 1 \end{bmatrix}$ and $R' = \begin{bmatrix} 1 & 0 \\ a & 1 \end{bmatrix}$ with $a \in \mathbf{R}$. The standard basis vectors e_1 and e_2 are taken by R to its columns, e_1 and $ae_1 + e_2$. Thus R acts geometrically as a shear by a in the e_1 -direction, magnifying volume by 1. (See figure 3.16.) Note that $1 = |\det R|$ as desired. The geometry of R' is left as an exercise.

The scale matrices are $S = \begin{bmatrix} a & 0 \\ 0 & 1 \end{bmatrix}$ and $S' = \begin{bmatrix} 1 & 0 \\ 0 & a \end{bmatrix}$. The standard basis gets taken by S to ae_1 and e_2 , so S acts geometrically as a scale in the e_1 -direction, magnifying volume by $|a|$; this is $|\det S|$, again as desired. (See figure 3.17.) The situation for S' is similar.




Figure 3.16. Shear




Figure 3.17. Scale

The transposition matrix is $T = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}$. This exchanges e_1 and e_2 , acting as a reflection through the diagonal, magnifying volume by 1. (See figure 3.18.) Since $\det T = -1$, the magnification factor is the absolute value of the determinant.




Figure 3.18. Reflection

Finally, the identity matrix $E = I$ has no effect, magnifying volume by 1, and any other echelon matrix E has bottom row $(0, 0)$ and hence squashes e_1 and e_2 to vectors whose last component is 0, magnifying volume by 0. (See figure 3.19.) The magnification factor is $|\det E|$ in both cases.

The discussion for scale matrices, transposition matrices, and echelon matrices generalizes effortlessly from 2 to n dimensions, but generalizing the discussion for recombine matrices $R_{i,j,a}$ takes a small argument. Since trans-




Figure 3.19. Squash

position matrices have no effect on volume, we may multiply $R_{i;j,a}$ from the left and from the right by various transposition matrices to obtain $R_{1;2,a}$ and study it instead. Multiplication by $R_{1;2,a}$ preserves all of the standard basis vectors except e_2 , which is taken to $ae_1 + e_2$ as before. The resulting parallelepiped $\mathcal{P}(e_1, ae_1 + e_2, e_3, \dots, e_n)$ consists of the parallelogram shown in the right side of figure 3.16, extended one unit in each of the remaining orthogonal $n - 2$ directions of \mathbf{R}^n . The n -dimensional volume of the parallelepiped is its base (the area of the parallelogram, 1) times its height (the $(n - 2)$ -dimensional volume of the unit box over each point of the parallelogram, again 1). That is, the n -by- n recombine matrix still magnifies volume by 1, the absolute value of its determinant, as desired. The base times height property of volume is yet another invocation here, but it is a consequence of a theorem to be proved in chapter 6, Fubini's Theorem. Summarizing,

Theorem 3.8.1 (Geometry of Linear Mappings). *Any linear mapping $T : \mathbf{R}^n \rightarrow \mathbf{R}^n$ is the composition of a possible squash followed by shears, scales and reflections. If the matrix of T is A then T magnifies volume by $|\det A|$.*

Proof. The matrix A of T is a product of elementary matrices and an echelon matrix. The elementary matrices act as shears, scales and reflections, and if the echelon matrix is not the identity it acts as a squash. This proves the first statement. Each elementary or echelon matrix magnifies volume by the absolute value of its determinant. The second statement follows since magnification and $|\det|$ are both multiplicative. \square

The work of this section has given a geometric interpretation of the magnitude of $\det A$: it is the magnification factor of multiplication by A . If the columns of A are denoted c_1, \dots, c_n then $Ae_j = c_j$ for $j = 1, \dots, n$, so that even more explicitly $|\det A|$ is the volume of the parallelepiped spanned by the columns of A .

Exercises

3.8.1. (a) The section states that the image of a union is the union of the images. More specifically, let A and B be any sets, let $f : A \rightarrow B$ be any

mapping, and let A_1, \dots, A_N be any subsets of A . Show that

$$f\left(\bigcup_{i=1}^N A_i\right) = \bigcup_{i=1}^N f(A_i).$$

(This exercise is purely set-theoretic, making no reference to our working environment of \mathbf{R}^n .)

(b) Consider a two-point set $A = \{a_1, a_2\}$ where $a_1 \neq a_2$, a one-point set $B = \{b\}$, and the only possible mapping $f : A \rightarrow B$, given by $f(a_1) = f(a_2) = b$. Let $A_1 = \{a_1\}$ and $A_2 = \{a_2\}$, subsets of A . What is the intersection $A_1 \cap A_2$? What is the image of the intersection, $f(A_1 \cap A_2)$? What are the images $f(A_1)$ and $f(A_2)$? What is the intersection of the images, $f(A_1) \cap f(A_2)$? Is the image of an intersection in general the intersection of the images?

3.8.2. Describe the geometric effect of multiplying by the matrices R' and S' in the section. Describe the effect of multiplying by R and S if $a < 0$.

3.8.3. Describe the geometric effect of multiplying by the 3-by-3 elementary matrices $R_{2;3,1}$, $R_{3;1,2}$, and $S_{2,-3}$.

3.8.4. Describe counterclockwise rotation of the plane by angle $\pi/2$ as a composition of shears and scales.

3.8.5. Describe counterclockwise rotation of the plane by angle θ (where $\cos \theta \neq 0$ and $\sin \theta \neq 0$) as a composition of shears and scales.

3.8.6. In \mathbf{R}^3 , describe the linear mapping that takes e_1 to e_2 , e_2 to e_3 , and e_3 to e_1 as a composition of shears, scales, and transpositions.

3.8.7. Let \mathcal{P} be the parallelogram in \mathbf{R}^2 spanned by (a, c) and (b, d) . Calculate directly that $|\det \begin{bmatrix} a & b \\ c & d \end{bmatrix}| = \text{area } \mathcal{P}$. (Hint: area = base \times height = $|(a, c)| |(b, d)| |\sin \theta_{(a,c), (b,d)}|$. It may be cleaner to find the square of the area.)

3.8.8. This exercise shows directly that $|\det| = \text{volume in } \mathbf{R}^3$. Let \mathcal{P} be the parallelepiped in \mathbf{R}^3 spanned by v_1, v_2, v_3 , let \mathcal{P}' be spanned by the vectors v'_1, v'_2, v'_3 obtained from performing the Gram-Schmidt process on the v_j 's, let $A \in M_3(\mathbf{R})$ have rows v_1, v_2, v_3 and let $A' \in M_3(\mathbf{R})$ have rows v'_1, v'_2, v'_3 .

(a) Explain why $\det A' = \det A$.

(b) Give a plausible geometric argument that $\text{vol } \mathcal{P}' = \text{vol } \mathcal{P}$.

(c) Show that

$$A'A'^t = \begin{bmatrix} |v'_1|^2 & 0 & 0 \\ 0 & |v'_2|^2 & 0 \\ 0 & 0 & |v'_3|^2 \end{bmatrix}.$$

Explain why therefore $|\det A'| = \text{vol } \mathcal{P}'$. It follows from parts (a) and (b) that that $|\det A| = \text{vol } \mathcal{P}$.

3.9 Geometry of the Determinant: Orientation

Recall from section 2.1 that a basis of \mathbf{R}^n is a set of vectors $\{f_1, \dots, f_p\}$ such that any vector in \mathbf{R}^n is a unique linear combination of the $\{f_j\}$. Though strictly speaking a basis is only a set, we adopt here the convention that the basis vectors are given in the specified order indicated. Given such a basis, view the vectors as columns and let F denote the matrix in $M_{n,p}(\mathbf{R})$ with columns f_1, \dots, f_p . Thus the order of the basis vectors is now relevant. For a standard basis vector e_j of \mathbf{R}^p , the matrix-by-vector product Fe_j gives the j th column f_j of F . Therefore, for any vector $x = (x_1, \dots, x_p) \in \mathbf{R}^p$ (viewed as a column),

$$Fx = F \cdot \left(\sum_{j=1}^p x_j e_j \right) = \sum_{j=1}^p x_j Fe_j = \sum_{j=1}^p x_j f_j.$$

This shows that multiplying all column vectors $x \in \mathbf{R}^p$ by the matrix F gives precisely the linear combinations of f_1, \dots, f_p . Hence we have the equivalences

$$\begin{aligned} & \{f_1, \dots, f_p\} \text{ is a basis for } \mathbf{R}^n \\ \iff & \left(\begin{array}{l} \text{each } y \in \mathbf{R}^n \text{ is uniquely expressible} \\ \text{as a linear combination of the } \{f_j\} \end{array} \right) \\ \iff & \left(\begin{array}{l} \text{each } y \in \mathbf{R}^n \text{ takes the form} \\ y = Fx \text{ for a unique } x \in \mathbf{R}^p \end{array} \right) \\ \iff & F \text{ is invertible} \\ \iff & F \text{ is square (i.e., } p = n) \text{ and } \det F \neq 0. \end{aligned}$$

These considerations have proved

Theorem 3.9.1 (Characterization of Bases). *Any basis of \mathbf{R}^n has n elements. The vectors $\{f_1, \dots, f_n\}$ form a basis exactly when the matrix F formed by their columns has nonzero determinant.*

Let $\{f_1, \dots, f_n\}$ be a basis of \mathbf{R}^n , and let F be the matrix formed by their columns. Abuse terminology and call $\det F$ the **determinant of the basis**, written $\det\{f_1, \dots, f_n\}$. Again, this depends on the order of the $\{f_j\}$. There are then two kinds of basis of \mathbf{R}^n , **positive** and **negative bases**, according to the sign of their determinants. The standard basis $\{e_1, \dots, e_n\}$ forms the columns of the identity matrix I and is therefore positive.

Since every linear mapping is continuous, the multilinear function $\det F$ is continuous in the n^2 entries of f_1, \dots, f_n . If a basis $\{f_1, \dots, f_n\}$ can be smoothly deformed via other bases to the standard basis then the corresponding determinants must change continuously to 1 without passing through 0. Such a basis must therefore be positive. Similarly, a negative basis can not be

smoothly deformed via other bases to the standard basis. It is also true but less clear (and not proved here) that every positive basis deforms smoothly to the standard basis.

The plane \mathbf{R}^2 is usually drawn with $\{e_1, e_2\}$ forming a counterclockwise angle of $\pi/2$. Two vectors $\{f_1, f_2\}$ form a basis if they are not collinear. Therefore the basis $\{f_1, f_2\}$ can be deformed via bases to $\{e_1, e_2\}$ exactly when the angle θ_{f_1, f_2} goes counterclockwise from f_1 to f_2 . (Recall from equation (2.2) that the angle between two nonzero vectors is between 0 and π .) This shows that in \mathbf{R}^2 , the basis $\{f_1, f_2\}$ is positive exactly when the angle from f_1 to f_2 is counterclockwise. (See figure 3.20.)




Figure 3.20. Positive and negative bases of \mathbf{R}^2

Three-space \mathbf{R}^3 is usually drawn with $\{e_1, e_2, e_3\}$ forming a right-handed triple, meaning that when the fingers of your right hand curl from e_1 to e_2 your thumb forms an acute angle with e_3 . Three vectors $\{f_1, f_2, f_3\}$ form a basis if they are not coplanar. In other words they must form a right- or left-handed triple. Only right-handed triples deform via other nonplanar triples to $\{e_1, e_2, e_3\}$. Therefore in \mathbf{R}^3 , the basis $\{f_1, f_2, f_3\}$ is positive exactly when it forms a right-handed triple. (See figure 3.21.)




Figure 3.21. Positive and negative bases of \mathbf{R}^3

The geometric generalization to \mathbf{R}^n of a counterclockwise angle in the plane and a right handed triple in space is not so clear, but the algebraic notion of positive basis is the same for all n .

Consider any invertible mapping $T : \mathbf{R}^n \rightarrow \mathbf{R}^n$ with matrix $A \in M_n(\mathbf{R})$, and any basis $\{f_1, \dots, f_n\}$ of \mathbf{R}^n . If F again denotes the matrix with columns f_1, \dots, f_n , then AF has columns $\{Af_1, \dots, Af_n\} = \{T(f_1), \dots, T(f_n)\}$. These form a new basis for \mathbf{R}^n with determinant

$$\det\{T(f_1), \dots, T(f_n)\} = \det AF = \det A \det F = \det A \det\{f_1, \dots, f_n\}.$$

This gives a geometric interpretation of the sign of $\det A$: If $\det A > 0$ then T preserves the orientation of bases, and if $\det A < 0$ then T reverses orientation. For example, the mapping with matrix

$$\begin{bmatrix} 0 & 0 & 0 & 1 \\ 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \end{bmatrix}$$

reverses orientation in \mathbf{R}^4 .

To summarize: Let A be an n -by- n matrix. Whether $\det A$ is nonzero says whether A is invertible, the magnitude of $\det A$ is the factor by which A magnifies volume, and (assuming that $\det A \neq 0$) the sign of $\det A$ determines how A affects orientation.

Exercises

3.9.1. Any invertible mapping $T : \mathbf{R}^n \rightarrow \mathbf{R}^n$ is a composition of scales, shears and transpositions. Give conditions on such a composition to make the mapping orientation-preserving, orientation-reversing.

3.9.2. Does the linear mapping $T : \mathbf{R}^n \rightarrow \mathbf{R}^n$ that takes e_1 to e_2 , e_2 to e_3 , \dots , e_n to e_1 preserve or reverse orientation? (The answer depends on n .) More generally, if π is a permutation in S_n , does the linear mapping taking e_1 to $e_{\pi(1)}$, \dots , e_n to $e_{\pi(n)}$ preserve or reverse orientation? (This depends on π .)

3.9.3. Argue geometrically in \mathbf{R}^2 that any basis can be smoothly deformed via other bases to the standard basis or to $\{e_1, -e_2\}$. Do the same for \mathbf{R}^3 and $\{e_1, e_2, -e_3\}$.

3.10 The Cross Product, Lines, and Planes in \mathbf{R}^3

Generally in \mathbf{R}^n there is no natural way to associate to a pair of vectors u and v a third vector. In \mathbf{R}^3 , however, the plane specified by u and v has only one orthogonal direction, i.e., dimension 3 is special because $3 - 2 = 1$. In \mathbf{R}^3

a normal vector to u and v can be specified by making suitable conventions on its orientation viz a viz the other two vectors, and on its length. This will give a vector-valued product of two vectors that is special to three-dimensional space, called the *cross product*. The first part of this section develops these ideas.

Given any two vectors $u, v \in \mathbf{R}^3$, we want their cross product $u \times v \in \mathbf{R}^3$ to be orthogonal to u and v ,

$$u \times v \perp u \quad \text{and} \quad u \times v \perp v. \quad (3.10)$$

There is the question of which way $u \times v$ should point along the line orthogonal to the plane spanned by u and v . The natural answer is that the direction should be chosen to make the ordered triple of vectors $\{u, v, u \times v\}$ positive unless it is degenerate,

$$\det(u, v, u \times v) \geq 0. \quad (3.11)$$

Also there is the question of how long $u \times v$ should be. With hindsight, we assert that specifying the length to be the area of the parallelogram spanned by u and v will work well. That is,

$$|u \times v| = \text{area } \mathcal{P}(u, v). \quad (3.12)$$

The three desired geometric properties (3.10) through (3.12) seem to describe the cross product completely. (See figure 3.22.)




Figure 3.22. The cross product of u and v

The three geometric properties also seem disparate. However, they combine into a uniform algebraic property, as follows. Since the determinant in (3.11) is nonnegative, it is the volume of the parallelepiped spanned by u , v , and $u \times v$. The volume is the base times the height, and since $u \times v$ is normal to u and v the base is the area of $\mathcal{P}(u, v)$ and the height is $|u \times v|$. Thus

$$\det(u, v, u \times v) = \text{area } \mathcal{P}(u, v) |u \times v|.$$

It follows from the previous display and (3.12) that

$$|u \times v|^2 = \det(u, v, u \times v).$$

Since orthogonal vectors have inner product 0, since the determinant is 0 when two rows agree, and since the square of the absolute value is the vector's inner product with itself, we can rewrite (3.10) and this last display (obtained from (3.11) and (3.12)) uniformly as equalities of the form $\langle u \times v, w \rangle = \det(u, v, w)$ for various w ,

$$\begin{aligned} \langle u \times v, u \rangle &= \det(u, v, u), \\ \langle u \times v, v \rangle &= \det(u, v, v), \\ \langle u \times v, u \times v \rangle &= \det(u, v, u \times v). \end{aligned} \tag{3.13}$$

Instead of saying what the cross product *is*, as an equality of the form $u \times v = f(u, v)$ would, the three equalities of (3.13) say how the cross product *interacts* with certain vectors—including itself—via the inner product. Again, the idea is to characterize rather than construct.

(The reader may object to the argument just given that $\det(u, v, u \times v) = \text{area } \mathcal{P}(u, v) |u \times v|$, on the grounds that we don't really understand the area of a 2-dimensional parallelogram in 3-dimensional space to start with, that in \mathbf{R}^3 we measure volume rather than area, and the parallelogram surely has volume zero. In fact, the argument can be viewed as motivating the formula as the *definition* of the area. This idea will be discussed more generally in section 8.1.)

We continue the discussion of (3.13). In general, a vector's inner products with other vectors completely describe the vector itself. The observation to make is that for any vector $x \in \mathbf{R}^n$ (n need not be 3 in this paragraph),

$$\text{if } \langle x, w \rangle = 0 \text{ for all } w \in \mathbf{R}^n \text{ then } x = \mathbf{0}_n.$$

This is because specializing w to x shows that $\langle x, x \rangle = 0$. It follows from the observation that for any two vectors $x, x' \in \mathbf{R}^n$,

$$\text{if } \langle x, w \rangle = \langle x', w \rangle \text{ for all } w \in \mathbf{R}^n \text{ then } x = x'.$$

That is, the inner product values $\langle x, w \rangle$ for all $w \in \mathbf{R}^n$ specify x , as anticipated. The following definition relies on this fact.

Definition 3.10.1 (Cross Product). *Let u and v be any two vectors in \mathbf{R}^3 . Their cross product is defined by the property*

$$\langle u \times v, w \rangle = \det(u, v, w) \quad \text{for all } w \in \mathbf{R}^3.$$

That is, $u \times v$ is the unique vector $x \in \mathbf{R}^3$ such that $\langle x, w \rangle = \det(u, v, w)$ for all $w \in \mathbf{R}^3$.

The problem with this definition is that we don't yet know that the cross product exists at all, only that if it exists then its intrinsic algebraic defining property characterizes it uniquely. Soon we will prove that the cross product

exists by writing its formula in coordinates and verifying that the formula satisfies the property, but for now we defer this in order to continue working intrinsically. Granting that the cross product does exist, graceful arguments with its defining condition show that it has all the properties that we want it to have.

Proposition 3.10.2 (Properties of the Cross Product).

(CP1) *The cross product is skew-symmetric: $v \times u = -u \times v$ for all $u, v \in \mathbf{R}^3$.*

(CP2) *The cross product is bilinear: For all scalars $a, a', b, b' \in \mathbf{R}$ and all vectors $u, u', v, v' \in \mathbf{R}^3$,*

$$\begin{aligned}(au + a'u') \times v &= a(u \times v) + a'(u' \times v), \\ u \times (bv + b'v') &= b(u \times v) + b'(u \times v').\end{aligned}$$

(CP3) *The cross product $u \times v$ is orthogonal to u and v .*

(CP4) *$u \times v = \mathbf{0}$ if and only if u and v are collinear (meaning that $u = av$ or $v = au$ for some $a \in \mathbf{R}$).*

(CP5) *If u and v are not collinear then the triple $\{u, v, u \times v\}$ is right handed.*

(CP6) *The magnitude $|u \times v|$ is the area of the parallelogram spanned by u and v .*

Proof. (1) This follows from the skew-symmetry of the determinant. For any $w \in \mathbf{R}^3$,

$$\langle v \times u, w \rangle = \det(v, u, w) = -\det(u, v, w) = -\langle u \times v, w \rangle = \langle -u \times v, w \rangle.$$

Since w is arbitrary, $v \times u = -u \times v$.

(2) For the first variable, this follows from the linearity of the determinant in its first row-vector variable and the linearity of the inner product in its first vector variable. Fix $a, a' \in \mathbf{R}$, $u, u', v \in \mathbf{R}^3$. For any $w \in \mathbf{R}^3$,

$$\begin{aligned}\langle (au + a'u') \times v, w \rangle &= \det(au + a'u', v, w) \\ &= a \det(u, v, w) + a' \det(u', v, w) \\ &= a \langle u \times v, w \rangle + a' \langle u' \times v, w \rangle \\ &= \langle a(u \times v) + a'(u' \times v), w \rangle.\end{aligned}$$

Since w is arbitrary, $(au + a'u') \times v = a(u \times v) + a'(u' \times v)$. The proof for the second variable follows from the result for the first variable and from (1).

(3) $\langle u \times v, u \rangle = \det(u, v, u) = 0$ since the determinant of a matrix with two equal rows vanishes. Similarly, $\langle u \times v, v \rangle = 0$.

(4) If $u = av$ then for any $w \in \mathbf{R}^3$,

$$\langle u \times v, w \rangle = \langle av \times v, w \rangle = \langle a(v \times v), w \rangle = a \langle v \times v, w \rangle = a \det(v, v, w) = 0.$$

Since w is arbitrary, $u \times v = \mathbf{0}$. Similarly if $v = au$.

Conversely, suppose that u and v are not collinear. Then they are linearly independent, and so no element of \mathbf{R}^3 can be written as a linear combination of u and v in more than one way. The set $\{u, v\}$ is not a basis of \mathbf{R}^3 , because every basis consists of three elements. Since no elements of \mathbf{R}^3 can be written as a linear combination of u and v in more than one way, and since $\{u, v\}$ is not a basis, the only possibility is that some $w \in \mathbf{R}^3$ can not be written as a linear combination of u and v at all. Thus the set $\{u, v, w\}$ is a linearly independent set of three elements, making it a basis of \mathbf{R}^3 . Compute that since $\{u, v, w\}$ is a basis,

$$\langle u \times v, w \rangle = \det(u, v, w) \neq 0.$$

Therefore $u \times v \neq \mathbf{0}$.

(5) By (4), $u \times v \neq \mathbf{0}$, so $0 < \langle u \times v, u \times v \rangle = \det(u, v, u \times v)$. By the results on determinants and orientation, $\{u, v, u \times v\}$ is right-handed.

(6) By definition, $|u \times v|^2 = \langle u \times v, u \times v \rangle = \det(u, v, u \times v)$. As discussed earlier in the section, $\det(u, v, u \times v) = \text{area}\mathcal{P}(u, v) |u \times v|$. The result follows from dividing by $|u \times v|$ if it is nonzero, and from (4) otherwise. \square

To prove that the cross product exists, it suffices to write a formula for it that satisfies Definition 3.10.1. The formula is

$$u \times v = (\det(u, v, e_1), \det(u, v, e_2), \det(u, v, e_3)).$$

This formula indeed satisfies the definition because by the linearity of the determinant in its third argument we have for any $w = (x_w, y_w, z_w) \in \mathbf{R}^3$,

$$\begin{aligned} \langle u \times v, w \rangle &= \det(u, v, e_1) \cdot x_w + \det(u, v, e_2) \cdot y_w + \det(u, v, e_3) \cdot z_w \\ &= \det(u, v, x_w e_1) + \det(u, v, y_w e_2) + \det(u, v, z_w e_3) \\ &= \det(u, v, x_w e_1 + y_w e_2 + z_w e_3) \\ &= \det(u, v, w). \end{aligned}$$

In coordinates, the formula for the cross product is

$$\begin{aligned} u \times v &= \left(\det \begin{bmatrix} x_u & y_u & z_u \\ x_v & y_v & z_v \\ 1 & 0 & 0 \end{bmatrix}, \det \begin{bmatrix} x_u & y_u & z_u \\ x_v & y_v & z_v \\ 0 & 1 & 0 \end{bmatrix}, \det \begin{bmatrix} x_u & y_u & z_u \\ x_v & y_v & z_v \\ 0 & 0 & 1 \end{bmatrix} \right) \\ &= (y_u z_v - z_u y_v, z_u x_v - x_u z_v, x_u y_v - y_u x_v). \end{aligned}$$

A bit more conceptually, the cross product formula in coordinates is

$$u \times v = \det \begin{bmatrix} x_u & y_u & z_u \\ x_v & y_v & z_v \\ e_1 & e_2 & e_3 \end{bmatrix}.$$


This is only a mnemonic device—strictly speaking, it doesn't lie within our grammar because the entries of the bottom row are vectors rather than scalars.

But even so, its two terms $x_u y_v e_3 - y_u x_v e_3$ do give the third entry of the cross product, and similarly for the others. In chapter 8, where we will have to compromise our philosophy of working intrinsically rather than in coordinates, this formula will be cited and generalized. In the meantime its details are not important except for mechanical calculations, and we want to use it as little as possible, as with the determinant earlier. Indeed, the display shows that the cross product is essentially a special case of the determinant.

It is worth knowing the cross products of the standard basis pairs,

$$\begin{aligned} e_1 \times e_1 &= \mathbf{0}_3, & e_1 \times e_2 &= e_3, & e_1 \times e_3 &= -e_2, \\ e_2 \times e_1 &= -e_3, & e_2 \times e_2 &= \mathbf{0}_3, & e_2 \times e_3 &= e_1, \\ e_3 \times e_1 &= e_2, & e_3 \times e_2 &= -e_1, & e_3 \times e_3 &= \mathbf{0}_3. \end{aligned}$$

Here $e_i \times e_j$ is $\mathbf{0}_3$ if $i = j$, and $e_i \times e_j$ is the third standard basis vector if $i \neq j$ and i and j are in order in the diagram



and $e_i \times e_j$ is minus the third standard basis vector if $i \neq j$ and i and j are out of order in the diagram.

The remainder of this section describes lines and planes in \mathbf{R}^3 .

A line in \mathbf{R}^3 is determined by a point p and a direction vector d . (See figure 3.23.) A point q lies on the line exactly when it is a translation from p by some multiple of d . Therefore,

$$\ell(p, d) = \{p + td : t \in \mathbf{R}\}.$$

In coordinates, a point (x, y, z) lies on $\ell((x_p, y_p, z_p), (x_d, y_d, z_d))$ exactly when

$$x = x_p + tx_d, \quad y = y_p + ty_d, \quad z = z_p + tz_d \quad \text{for some } t \in \mathbf{R}.$$

If the components of d are all nonzero then the relation between the coordinates can be expressed without the parameter t ,

$$\frac{x - x_p}{x_d} = \frac{y - y_p}{y_d} = \frac{z - z_p}{z_d}.$$

For example, the line through $(1, 1, 1)$ in the direction $(1, 2, 3)$ consists of all points (x, y, z) satisfying $x = 1 + t$, $y = 1 + 2t$, $z = 1 + 3t$ for $t \in \mathbf{R}$, or, equivalently, satisfying $x - 1 = (y - 1)/2 = (z - 1)/3$.

A plane in \mathbf{R}^3 is determined by a point p and a normal (orthogonal) vector n . (See figure 3.24.) A point x lies on the plane exactly when the vector from p to x is orthogonal to n . Therefore,




Figure 3.23. Line in \mathbf{R}^3

$$P(p, n) = \{x \in \mathbf{R}^3 : \langle x - p, n \rangle = 0\}.$$

In coordinates, a point (x, y, z) lies on $P((x_p, y_p, z_p), (x_n, y_n, z_n))$ exactly when

$$(x - x_p)x_n + (y - y_p)y_n + (z - z_p)z_n = 0.$$




Figure 3.24. Plane in \mathbf{R}^3

Exercises

3.10.1. Evaluate $(2, 0, -1) \times (1, -3, 2)$.

3.10.2. Suppose that $v \times e_1 = v \times e_2 = \mathbf{0}$. Describe v .

3.10.3. True or false: For all u, v, w in \mathbf{R}^3 , $(u \times v) \times w = u \times (v \times w)$.

3.10.4. Express $(u + v) \times (u - v)$ as a scalar multiple of $u \times v$.

3.10.5. For fixed u, v in \mathbf{R}^3 with $u \neq \mathbf{0}$, describe the vectors w satisfying the condition $u \times v = u \times w$.

3.10.6. What is the line specified by two points p and p' ?

3.10.7. Give conditions on the points p , p' and the directions d , d' so that $\ell(p, d) = \ell(p', d')$.

3.10.8. Express the relation between the coordinates of a point on $\ell(p, d)$ if the x -component of d is 0.

3.10.9. What can you conclude about the lines

$$\frac{x - x_p}{x_d} = \frac{y - y_p}{y_d} = \frac{z - z_p}{z_d} \quad \text{and} \quad \frac{x - x_p}{x_D} = \frac{y - y_p}{y_D} = \frac{z - z_p}{z_D}$$

given that $x_d x_D + y_d y_D + z_d z_D = 0$? What can you conclude if $x_d/x_D = y_d/y_D = z_d/z_D$?

3.10.10. Show that $\ell(p, d)$ and $\ell(p', d')$ intersect if and only if the linear equation $Dt = \Delta p$ is solvable, where $D \in M_{3,2}(\mathbf{R})$ has columns d and d' , t is the column vector $\begin{bmatrix} t_1 \\ t_2 \end{bmatrix}$, and $\Delta p = p' - p$. For what points p and p' do $\ell(p, (1, 2, 2))$ and $\ell(p', (2, -1, 4))$ intersect?

3.10.11. Use vector geometry to show that the distance from the point q to the line $\ell(p, d)$ is

$$\frac{|(q - p) \times d|}{|d|}.$$

(Hint: what is the area of the parallelogram spanned by $q - p$ and d ?) Find the distance from the point $(3, 4, 5)$ to the line $\ell((1, 1, 1), (1, 2, 3))$.

3.10.12. Show that the time of nearest approach of two particles whose positions are $s(t) = p + tv$, $\bar{s}(t) = \bar{p} + t\bar{v}$ is $t = -\langle \Delta p, \Delta v \rangle / |\Delta v|^2$. (You may assume that the particles are at their nearest approach when the difference of their velocities is orthogonal to the difference of their positions.)

3.10.13. Write the equation of the plane through $(1, 2, 3)$ with normal direction $(1, 1, 1)$.

3.10.14. Where does the plane $x/a + y/b + z/c = 1$ intersect each axis?

3.10.15. Specify the plane containing the point p and spanned by directions d and d' . Specify the plane containing the three points p , q , and r .

3.10.16. Use vector geometry to show that the distance from the point q to the plane $P(p, n)$ is

$$\frac{|\langle q - p, n \rangle|}{|n|}.$$

(Hint: Resolve $q - p$ into components parallel and normal to n .) Find the distance from the point $(3, 4, 5)$ to the plane $P((1, 1, 1), (1, 2, 3))$.

3.11 Summary

Linear algebra is easy in the abstract since the vector space operations pass through linear mappings, and it is easy in the concrete since mechanical matrix manipulations are straightforward. While the matrix methods from this chapter are handy computational tools, it is also crucial to understand the intrinsic notion of a linear mapping: this is the idea that we will use to define multivariable differentiation in the next chapter.

The Derivative

In one-variable calculus the derivative is a limit of difference quotients, but this idea does not generalize to many variables. The multivariable definition of the derivative to be given in this chapter has three noteworthy features:

- The derivative is defined as a linear mapping.
- The derivative is characterized intrinsically rather than constructed in coordinates.
- The derivative is characterized by the property of closely approximating the original function near the point of approximation.

Section 4.1 defines the multivariable derivative in a way that captures these ideas. Section 4.2 obtains some basic results about the derivative intrinsically, notably the Chain Rule. Section 4.3 computes with coordinates to calculate the derivative by considering one variable at a time and using the techniques of one-variable calculus. This section also obtains a coordinate-based version of the Chain Rule. Section 4.4 studies the multivariable counterparts of higher order derivatives from one-variable calculus. Section 4.5 discusses optimization of functions of many variables. Finally, section 4.6 discusses the rate of change of a function of many variables as its input moves in any fixed direction, not necessarily parallel to a coordinate axis.

4.1 The Derivative Redefined

In one variable calculus, the derivative of a function $f : \mathbf{R} \rightarrow \mathbf{R}$ at a point $a \in \mathbf{R}$ is defined as a limit:

$$f'(a) = \lim_{h \rightarrow 0} \frac{f(a+h) - f(a)}{h}. \quad (4.1)$$

But for any integer $n > 1$, the corresponding expression makes no sense for a mapping $f : \mathbf{R}^n \rightarrow \mathbf{R}^m$ and for a point a of \mathbf{R}^n . Indeed, the expression is

$$\text{“ } \lim_{h \rightarrow \mathbf{0}_n} \frac{f(a+h) - f(a)}{h}, \text{ ”}$$

but this is not even grammatically admissible—there is no notion of division by the vector h . That is, the definition of derivative does not generalize transparently to more than one input variable. We need to think about the derivative in different terms.

In the one variable case $n = m = 1$, the scalar $f'(a)$ is interpreted geometrically as the slope of the tangent line to the graph of f at the point $(a, f(a))$. The graph of the tangent line consists of the points

$$(a+h, f(a) + f'(a)h) \quad \text{for all } h \in \mathbf{R},$$

whereas the graph of f itself consists of the points

$$(a+h, f(a+h)) \quad \text{for all } h \in \mathbf{R}.$$

The tangent line is distinguished by the property that the second coordinate of its points approximates the second coordinate of the corresponding points of the original curve extremely well when h is small. To explicate “extremely well,” rewrite (4.1) as

$$\lim_{h \rightarrow 0} \frac{f(a+h) - f(a) - f'(a)h}{h} = 0. \quad (4.2)$$

The numerator here is the difference between the y -coordinates in the previous two displays. (See figure 4.1.) The rewritten equation (4.2) shows that when h is small, not only is the vertical distance $f(a+h) - f(a) - f'(a)h$ from the tangent line to the curve small as well, but it is small even relative to the horizontal distance h .

This is the right idea. Instead of viewing the one-variable derivative as the scalar $f'(a)$, think of it as the corresponding linear mapping $T_a : \mathbf{R} \rightarrow \mathbf{R}$, multiplication by $f'(a)$. That is, think of it as the mapping

$$T_a(h) = f'(a)h \quad \text{for all } h \in \mathbf{R},$$

with its characterizing property

$$\lim_{h \rightarrow 0} \frac{|f(a+h) - f(a) - T_a(h)|}{|h|} = 0.$$

Since we are working with one variable, the absolute value signs have no effect on the condition (i.e., it is the same condition as (4.2)), but they will become essential when we generalize to many variables. See figure 4.2 for an illustration of the condition—the figure is similar to figure 4.1, but it shows that the condition is best understood in the coordinate system centered at the point of tangency. The shaded axis-portions in the figure are the numerator and denominator in the condition, illustrating that the numerator is much




Figure 4.1. Vertical distance from tangent line to curve




Figure 4.2. Vertical distance in local coordinates

smaller. The condition is being called a “characterizing property” because it describes how T_a behaves rather than simply saying what T_a is. The issue here is similar to the previous discussions of the determinant and the cross product.

The characterizing property is **local**, meaning that it depends only on the behavior of $f(x)$ for x -values $x = a + h$ near a and for $x = a$ itself. On the other hand, the characterizing property depends on the behavior of $f(x)$ for *all* x -values near a along with $x = a$. Thus the following definition is handy for our purposes.

Definition 4.1.1 (Interior Point). Let A be a subset of \mathbf{R}^n , and let a be a point of A . Then a is an **interior point** of A if some ε -ball about a is a subset of A . That is, a is an interior point of A if $B(a, \varepsilon) \subset A$ for some $\varepsilon > 0$.

Now we can define the derivative in a way that encompasses many variables and is suitably local.

Definition 4.1.2 (Derivative). Let A be a subset of \mathbf{R}^n , let $f : A \rightarrow \mathbf{R}^m$ be a mapping, and let a be an interior point of A . Then f is **differentiable at a** if there exists a linear mapping $T_a : \mathbf{R}^n \rightarrow \mathbf{R}^m$ satisfying the condition

$$\lim_{h \rightarrow \mathbf{0}_n} \frac{|f(a+h) - f(a) - T_a(h)|}{|h|} = 0. \quad (4.3)$$

This T_a is called **the derivative of f at a** , written Df_a or $(Df)_a$. When f is differentiable at a , the matrix of the linear mapping Df_a is written $f'(a)$ and is called **the Jacobian matrix of f at a** .

Here are some points to note about Definition 4.1.2:

- Any assertion that a mapping is differentiable at a point has the connotation that the point is an interior point of the mapping's domain. That is, if f is differentiable at a then $B(a, \varepsilon) \subset A$ for some $\varepsilon > 0$. In the special case $n = 1$ we are not allowing the derivative at an endpoint of the domain.
- The limit in (4.3) is a function limit as in Definition 2.5.1. Specifically, the function is

$$g : B(\mathbf{0}_n, \varepsilon) - \{\mathbf{0}_n\} \rightarrow \mathbf{R}, \quad g(h) = \frac{|f(a+h) - f(a) - T_a(h)|}{|h|},$$

for the same ε as in the first bullet. Thus indeed $\mathbf{0}_n$ is a limit point of the domain of g , as required for the limit to exist. In (4.3) the numerator is absolute value on \mathbf{R}^m , the denominator is absolute value on \mathbf{R}^n , and the quotient is real.

- The domain of the linear mapping T_a is unrestricted even if f itself is defined only locally about a . Indeed, the definition of linearity requires that the linear mapping have all of \mathbf{R}^n as its domain. Any linear mapping is so uniform that in any case its behavior on all of \mathbf{R}^n is determined by its behavior on any ε -ball about $\mathbf{0}_n$ (exercise 4.1.1). In geometric terms, the graph of T , the tangent object approximating the graph of f at $(a, f(a))$, extends without bound, even if the graph of f itself is restricted to points near $(a, f(a))$. But the approximation of the graph by the tangent object needs to be close only near the point of tangency.

Returning to the idea of the derivative as a linear mapping, when $n = 2$ and $m = 1$ a function $f : A \rightarrow \mathbf{R}$ is differentiable at an interior point (a, b) of A if for small scalar values h and k , $f(a+h, b+k) - f(a, b)$ is well approximated by a linear function

$$T(h, k) = \alpha h + \beta k$$

where α and β are scalars. Since the equation $z = f(a, b) + \alpha h + \beta k$ describes a plane in (x, y, z) -space (where $h = x - a$ and $k = y - b$), f is differentiable at (a, b) if its graph has a well-fitting tangent plane through $(a, b, f(a, b))$. (See figure 4.3.) Here the derivative of f at (a, b) is the linear mapping taking (h, k) to $\alpha h + \beta k$ and the Jacobian matrix of f at a is therefore $[\alpha, \beta]$. The tangent plane in the figure is not the graph of the derivative $Df_{(a,b)}$, but rather a translation of the graph. Another way to say this is that the $(h, k, Df_{(a,b)}(h, k))$ -coordinate system has its origin at the point $(a, b, f(a, b))$ in the figure.




Figure 4.3. Graph and tangent plane

When $n = 1$ and $m = 3$, a mapping $f : A \rightarrow \mathbf{R}^3$ is differentiable at an interior point a of A if $f(a+h) - f(a)$ is closely approximated for small real h by a linear mapping

$$T(h) = \begin{bmatrix} \alpha \\ \beta \\ \gamma \end{bmatrix} h$$

for some scalars α , β , and γ . As h varies through \mathbf{R} , $f(a) + T(h)$ traverses the line $\ell = \ell(f(a), (\alpha, \beta, \gamma))$ in \mathbf{R}^3 that is tangent at $f(a)$ to the output curve of f . (See figure 4.4.) Here $Df_a(h) = \begin{bmatrix} \alpha \\ \beta \\ \gamma \end{bmatrix} h$ and the corresponding Jacobian matrix is $\begin{bmatrix} \alpha \\ \beta \\ \gamma \end{bmatrix}$. Note that the figure does not show the domain of f , so it may help to think of f as a time-dependent traversal of the curve rather than as the curve itself. The figure does not have room for the $(h, Df_a(h))$ -coordinate

system (which is 4-dimensional), but the $Df_a(h)$ -coordinate system has its origin at the point $f(a)$.




Figure 4.4. Tangent to a parameterized curve

For an example, let $A = B((0, 0), 1)$ be the unit disk in \mathbf{R}^2 , and consider the function

$$f : A \longrightarrow \mathbf{R}, \quad f(x, y) = x^2 - y^2.$$

We show that for any point $(a, b) \in A$, f is differentiable at (a, b) and its derivative is the linear mapping

$$T_{(a,b)} : \mathbf{R}^2 \longrightarrow \mathbf{R}^2, \quad T_{(a,b)}(h, k) = 2ah - 2bk.$$

To verify this, we need to check Definition 4.1.2. The point that is written in the definition intrinsically as a (where a is a vector) is written here in coordinates as (a, b) (where a and b are scalars), and similarly the vector h in the definition is written (h, k) here, because the definition is conceptual whereas here we are going to compute. To check the definition, first note that every point (a, b) of A is an interior point. This doesn't deserve a detailed proof right now, only a quick comment. Second, confirm (4.3) by calculating that

$$\begin{aligned}
0 &\leq \frac{|f(a+h, b+k) - f(a, b) - T_{(a,b)}(h, k)|}{|(h, k)|} \\
&= \frac{|(a+h)^2 - (b+k)^2 - a^2 + b^2 - 2ah + 2bk|}{|(h, k)|} \\
&= \frac{|h^2 - k^2|}{|(h, k)|} \\
&\leq \frac{|h|^2 + |k|^2}{|(h, k)|} \\
&\leq \frac{|(h, k)|^2 + |(h, k)|^2}{|(h, k)|} \\
&= 2|(h, k)|.
\end{aligned}$$

This shows that by the squeeze theorem,

$$\lim_{(h,k) \rightarrow (0,0)} \frac{|f(a+h, b+k) - f(a, b) - T_{(a,b)}(h, k)|}{|(h, k)|} = 0.$$

Also, it tacitly shows how the derivative was found for us to verify: the difference $f(a+h, b+k) - f(a, b)$ is $2ah - 2bk + h^2 - k^2$, which as a function of h and k has a linear part $2ah - 2bk$ and a quadratic part $h^2 - k^2$ that is much smaller when h and k are small. The linear approximation of the difference is the derivative.

Before continuing, we need to settle a grammatical issue. Definition 4.1.2 refers to *any* linear mapping that satisfies condition (4.3) as *the* derivative of f at a . Fortunately, the derivative, if it exists, is unique, justifying the definite article. The uniqueness is geometrically plausible: if two “straight” objects (e.g., lines or planes) approximate the graph of f well near $(a, f(a))$, then they should also approximate each other well enough that straightness forces them to coincide. Here is the quantitative argument.

Proposition 4.1.3 (Uniqueness of the Derivative). *Let $f : A \rightarrow \mathbf{R}^m$ (where $A \subset \mathbf{R}^n$) be differentiable at a . Then there is only one linear mapping satisfying the definition of Df_a .*

Proof. Suppose that the linear mappings $T_a, \tilde{T}_a : \mathbf{R}^n \rightarrow \mathbf{R}^m$ are both derivatives of f at a . To show that $T_a = \tilde{T}_a$, it suffices to show that

$$\frac{|T_a(x) - \tilde{T}_a(x)|}{|x|} = 0 \quad \text{for all nonzero } x \in \mathbf{R}^n.$$

Fix a nonzero vector x , and let t be a nonzero variable scalar. Note that

$$T_a(x) - \tilde{T}_a(x) = \frac{T_a(tx) - \tilde{T}_a(tx)}{t},$$

so that taking absolute values and dividing by $|x|$ gives

$$\begin{aligned} \frac{|T_a(x) - \tilde{T}_a(x)|}{|x|} &= \frac{|T_a(tx) - \tilde{T}_a(tx)|}{|tx|} \\ &= \frac{|T_a(h) - \tilde{T}_a(h)|}{|h|} \quad \text{where } h = tx. \end{aligned}$$

(Note the interplay among the three absolute values on \mathbf{R}^n , on \mathbf{R}^m , and on \mathbf{R} .) That is,

$$\frac{|T_a(x) - \tilde{T}_a(x)|}{|x|} = \frac{|T_a(h) + f(a) - f(a+h) + f(a+h) - f(a) - \tilde{T}_a(h)|}{|h|}.$$

By the Triangle Inequality,

$$\begin{aligned} \frac{|T_a(x) - \tilde{T}_a(x)|}{|x|} &\leq \frac{|T_a(h) + f(a) - f(a+h)|}{|h|} + \frac{|f(a+h) - f(a) - \tilde{T}_a(h)|}{|h|} \\ &= \frac{|f(a+h) - f(a) - T_a(h)|}{|h|} + \frac{|f(a+h) - f(a) - \tilde{T}_a(h)|}{|h|}. \end{aligned}$$

Now let the scalar t go to 0 in \mathbf{R} . This takes the vector $h = tx$ to $\mathbf{0}_n$ in \mathbf{R}^n and thus takes the right side to 0 since T_a and \tilde{T}_a both satisfy condition (4.3). Consequently the left side $|T_a(x) - \tilde{T}_a(x)|/|x|$, which is independent of t , must be 0, and the proof is complete. \square

We will study the derivative via two routes. On the one hand, the linear mapping $Df_a : \mathbf{R}^n \rightarrow \mathbf{R}^m$ is specified by mn scalar entries of its matrix $f'(a)$, and so calculating the derivative is tantamount to determining these scalars by using coordinates. On the other hand, developing conceptual theorems without getting lost in coefficients and indices requires the intrinsic idea of the derivative as a well-approximating linear mapping.

Exercises

4.1.1. Let $T : \mathbf{R}^n \rightarrow \mathbf{R}^m$ be a linear mapping. Show that for any $\varepsilon > 0$, the behavior of T on $B(\mathbf{0}_n, \varepsilon)$ determines the behavior of T everywhere.

4.1.2. Give a geometric interpretation of the derivative when $n = m = 2$. Give a geometric interpretation of the derivative when $n = 1$ and $m = 2$.

4.1.3. Prove: if $f : A \rightarrow \mathbf{R}^m$ (where $A \subset \mathbf{R}^n$) is differentiable at a then f is continuous at a .

4.1.4. Prove the **componentwise nature of differentiability**: Let $f : A \rightarrow \mathbf{R}^m$ (where $A \subset \mathbf{R}^n$) have component functions f_1, \dots, f_m , and let a be a point of A . Show that f is differentiable at a if and only if each component f_i is, in which case Df_a has component functions $(Df_1)_a, \dots, (Df_m)_a$.

4.1.5. Let $f(x, y) = (x^2 - y^2, 2xy)$. Show that $Df_{(a,b)}(h, k) = (2ah - 2bk, 2bh + 2ak)$ for all $(a, b) \in \mathbf{R}^2$. (By the previous problem, you may work componentwise.)

4.1.6. Let $g(x, y) = xe^y$. Show that $Dg_{(a,b)}(h, k) = he^b + kae^b$ for all $(a, b) \in \mathbf{R}^2$. (You may need to use the Mean Value Theorem. Note that $1 = e^0$.)

4.1.7. Show that if $f : \mathbf{R}^n \rightarrow \mathbf{R}^m$ satisfies $|f(x)| \leq |x|^2$ for all $x \in \mathbf{R}^n$ then f is differentiable at $\mathbf{0}_n$.

4.1.8. Show that the function $f(x, y) = \sqrt{|xy|}$ for all $(x, y) \in \mathbf{R}^2$ is not differentiable at $(0, 0)$. (First see what $Df_{(0,0)}(h, 0)$ and $Df_{(0,0)}(0, k)$ need to be.)

4.2 Basic Results

Before digging into the derivative coordinatewise via the Jacobian matrix, we derive some results intrinsically from the characterizing definition. We begin by computing two explicit derivatives.

Proposition 4.2.1 (Derivatives of Constant and Linear Mappings).

- (1) Let $C : A \rightarrow \mathbf{R}^m$ (where $A \subset \mathbf{R}^n$) be the constant mapping $C(x) = c$ for all $x \in A$, where c is some fixed value in \mathbf{R}^m . Then the derivative of C at any interior point a of A is the zero mapping.
- (2) The derivative of a linear mapping $T : \mathbf{R}^n \rightarrow \mathbf{R}^m$ at any point $a \in \mathbf{R}^n$ is again T .

Proof. Both of these results hold essentially by grammar. In general the derivative of a mapping f at a is the linear mapping that well approximates $f(a+h) - f(a)$ for h near $\mathbf{0}_n$. But $C(a+h) - C(a)$ is the zero mapping for all $h \in A$, so it is well approximated near $\mathbf{0}_n$ by the zero mapping on \mathbf{R}^n . Similarly, $T(a+h) - T(a)$ is $T(h)$ for all $h \in \mathbf{R}^n$, and this linear mapping is well approximated by itself near $\mathbf{0}_n$.

To prove (1) more symbolically, let $Z : \mathbf{R}^n \rightarrow \mathbf{R}^m$ denote the zero mapping, $Z(h) = \mathbf{0}_m$ for all $h \in \mathbf{R}^n$. Then

$$\lim_{h \rightarrow \mathbf{0}_n} \frac{|C(a+h) - C(a) - Z(h)|}{|h|} = \lim_{h \rightarrow \mathbf{0}_n} \frac{|c - c - \mathbf{0}_m|}{|h|} = \lim_{h \rightarrow \mathbf{0}_n} 0 = 0.$$

Thus Z meets the condition to be DC_a . And (2) is similar, left as exercise 4.2.1. \square

A natural question is what effect addition and scalar multiplication of mappings have on their derivatives. The answer is unsurprising. Differentiation passes through the operations: the derivative of a sum is the sum of the derivatives and the derivative of a scalar multiple is the scalar multiple of the derivative.

Proposition 4.2.2 (Linearity of the Derivative). *Let $f : A \rightarrow \mathbf{R}^m$ (where $A \subset \mathbf{R}^n$) and $g : B \rightarrow \mathbf{R}^m$ (where $B \subset \mathbf{R}^n$) be mappings, and let a be a point of $A \cap B$. Suppose that f and g are differentiable at a with derivatives Df_a and Dg_a . Then*

- (1) *The sum $f + g : A \cap B \rightarrow \mathbf{R}^m$ is differentiable at a with derivative $D(f + g)_a = Df_a + Dg_a$.*
- (2) *For any $\alpha \in \mathbf{R}$, the scalar multiple $\alpha f : A \rightarrow \mathbf{R}^m$ is differentiable at a with derivative $D(\alpha f)_a = \alpha Df_a$.*

Proof. Since f and g are differentiable at a , some ball about a lies in A and some ball about a lies in B . The smaller of these two balls lies in $A \cap B$. That is, a is an interior point of the domain of $f + g$. With this topological issue settled, proving the proposition reduces to direct calculation. For (1),

$$\begin{aligned} & \lim_{h \rightarrow \mathbf{0}_n} \frac{|(f + g)(a + h) - (f + g)(a) - (Df_a + Dg_a)(h)|}{|h|} \\ &= \lim_{h \rightarrow \mathbf{0}_n} \frac{|f(a + h) + g(a + h) - f(a) - g(a) - Df_a(h) - Dg_a(h)|}{|h|} \\ &\leq \lim_{h \rightarrow \mathbf{0}_n} \frac{|f(a + h) - f(a) - Df_a(h)|}{|h|} + \lim_{h \rightarrow \mathbf{0}_n} \frac{|g(a + h) - g(a) - Dg_a(h)|}{|h|} \\ &= 0. \end{aligned}$$

And (2) is similar (exercise 4.2.2). □

You may want to contrast how nicely our topological setup worked at the beginning of this proof to the irritating example that we encountered in connection with the Sum Rule for mappings back in section 2.5.

Elaborate mappings are often built by composing simpler ones. The next theorem yields the important result that the derivative of a composition is the composition of the derivatives. Like all of our results thus far, this seems plausible—the best linear approximation of a composition could well be the composition of the best linear approximations—but it is deeper than the previous facts about derivatives, with a correspondingly denser proof.

Theorem 4.2.3 (Chain Rule). *Let $f : A \rightarrow \mathbf{R}^m$ be a mapping, and let $B \subset \mathbf{R}^m$ be a set containing $f(A)$, and let $g : B \rightarrow \mathbf{R}^l$ be a mapping. Thus the composition $g \circ f : A \rightarrow \mathbf{R}^l$ is defined. If f is differentiable at the point $a \in A$, and g is differentiable at the point $f(a) \in B$, then the composition $g \circ f$ is differentiable at the point a , and its derivative there is*

$$D(g \circ f)_a = Dg_{f(a)} \circ Df_a.$$

In terms of Jacobian matrices, since the matrix of a composition is the product of the matrices, the Chain Rule is

$$(g \circ f)'(a) = g'(f(a))f'(a).$$

The proof uses two lemmas, the first a restatement of exercise 3.1.15(a) (and now exercise 4.2.3).

Lemma 4.2.4 (Linear Magnification Boundedness). *Consider a linear mapping $T : \mathbf{R}^n \rightarrow \mathbf{R}^m$. There exists a nonnegative constant K such that $|T(h)| \leq K|h|$ for all $h \in \mathbf{R}^n$.*

The second lemma is a quick consequence of the first and of the definition of derivative.

Lemma 4.2.5 (Absolute Difference Quotient Boundedness). *Let the mapping $f : A \rightarrow \mathbf{R}^m$ (where $A \subset \mathbf{R}^n$) be differentiable at the point $a \in A$. Then the quantity*

$$\frac{|f(a+h) - f(a)|}{|h|}$$

is bounded as $h \rightarrow \mathbf{0}_n$.

Proof. Let $S = Df_a$. By the triangle inequality,

$$|f(a+h) - f(a)| \leq |f(a+h) - f(a) - S(h)| + |S(h)|,$$

and so by the first lemma there is a nonnegative constant K such that

$$\frac{|f(a+h) - f(a)|}{|h|} \leq \frac{|f(a+h) - f(a) - S(h)|}{|h|} + K.$$

Since f is differentiable at a , the right side goes to K as $h \rightarrow \mathbf{0}_n$, and therefore the left side is bounded. \square

In particular, this lemma shows that if f is differentiable at a then $\lim_{h \rightarrow \mathbf{0}_n} f(a+h) = f(a)$. That is, if f is differentiable at a then f is continuous at a .

Now we can prove the Chain Rule.

Proof. The Chain Rule follows from the definition of derivative and the two lemmas, via straightforward but delicate symbol juggling and an auxiliary variable. Let

$$S = Df_a : \mathbf{R}^n \rightarrow \mathbf{R}^m \quad \text{and} \quad T = Dg_{f(a)} : \mathbf{R}^m \rightarrow \mathbf{R}^\ell.$$

Then we must show that

$$T \circ S : \mathbf{R}^n \rightarrow \mathbf{R}^\ell$$

satisfies the defining property of $D(g \circ f)_a$.

Some ε -ball about a lies in A since Df_a exists. Let $b = f(a)$, and for all $h \in \mathbf{R}^n$ small enough that $a+h$ lies in the ball, let $k = f(a+h) - f(a)$, so that $b+k = f(a+h)$. Then by substitution,

$$\begin{aligned}(g \circ f)(a + h) - (g \circ f)(a) &= g(f(a + h)) - g(f(a)) \\ &= g(b + k) - g(b).\end{aligned}$$

That is, letting $r(k) = g(b + k) - g(b) - T(k)$,

$$(g \circ f)(a + h) - (g \circ f)(a) = T(k) + r(k).$$

Similarly letting $q(h) = f(a + h) - f(a) - S(h) = k - S(h)$ gives

$$(g \circ f)(a + h) - (g \circ f)(a) = T(S(h) + q(h)) + r(k).$$

And since T is linear it follows that

$$(g \circ f)(a + h) - (g \circ f)(a) - (T \circ S)(h) = T(q(h)) + r(k).$$

So we need to show that

$$\lim_{h \rightarrow \mathbf{0}_n} \frac{|T(q(h)) + r(k)|}{|h|} = 0,$$

where

$$\begin{aligned}q(h) &= f(a + h) - f(a) - S(h), \\ r(k) &= g(b + k) - g(b) - T(k), \\ k &= f(a + h) - f(a).\end{aligned}$$

By the Triangle Inequality, it suffices to show instead that

$$\lim_{h \rightarrow \mathbf{0}_n} \frac{|T(q(h))|}{|h|} = 0 \quad \text{and} \quad \lim_{h \rightarrow \mathbf{0}_n} \frac{|r(k)|}{|h|} = 0. \quad (4.4)$$

Note that the respective characterizing properties of S and T are

$$\lim_{h \rightarrow \mathbf{0}_n} \frac{|q(h)|}{|h|} = 0 \quad \text{and} \quad \lim_{k \rightarrow \mathbf{0}_m} \frac{|r(k)|}{|k|} = 0.$$

The Linear Magnification Boundedness Lemma, with $q(h)$ in place of h , says that $|T(q(h))| \leq K|q(h)|$ for a nonnegative constant K . By the characterizing property of S , the first of the two desired limits in (4.4) follows. Also, the Absolute Difference Quotient Boundedness Lemma says that for some positive constant C we have $|k| \leq C|h|$ for all small enough h . That is, if h is small then so is k . Thus, by the characterizing property of T , given any $\varepsilon > 0$ we have,

$$\text{for all small enough } h, \quad |r(k)| \leq \frac{\varepsilon}{C} \cdot |k| \leq \varepsilon|h|,$$

and the second of the two desired limits in (4.4) follows as well. This completes the proof of the Chain Rule. \square

Two quick applications of the Chain Rule arise naturally for scalar-valued functions. Given two such functions, not only is their sum defined, but since \mathbf{R} is a field (unlike \mathbf{R}^m for $m > 1$), so is their product and so is their quotient at points where g is nonzero. With some help from the Chain Rule, the derivative laws for product and quotient follow easily from elementary calculations.

Lemma 4.2.6 (Derivatives of the Product and Reciprocal Functions). *Define the product function,*

$$p : \mathbf{R}^2 \longrightarrow \mathbf{R}, \quad p(x, y) = xy,$$

and define the reciprocal function

$$r : \mathbf{R} - \{0\} \longrightarrow \mathbf{R}, \quad r(x) = 1/x.$$

Then

(1) *The derivative of p at the point $(a, b) \in \mathbf{R}^2$ is*

$$Dp_{(a,b)}(h, k) = ak + bh.$$

(2) *The derivative of r at the nonzero real number a is*

$$Dr_a(h) = -h/a^2.$$

Proof. (1) Compute,

$$\begin{aligned} \frac{|p(a+h, b+k) - p(a, b) - ak - bh|}{|(h, k)|} &= \frac{|(a+h)(b+k) - ab - ak - bh|}{|(h, k)|} \\ &= \frac{|hk|}{|(h, k)|}. \end{aligned}$$

For any value of (h, k) , $|h| \leq |(h, k)|$ and $|k| \leq |(h, k)|$ by the Size Bounds, so $|hk| = |h||k| \leq |(h, k)|^2$. Therefore,

$$\begin{aligned} \lim_{(h,k) \rightarrow (0,0)} \frac{|p(a+h, b+k) - p(a, b) - ak - bh|}{|(h, k)|} &\leq \lim_{(h,k) \rightarrow (0,0)} \frac{|(h, k)|^2}{|(h, k)|} \\ &= \lim_{(h,k) \rightarrow (0,0)} |(h, k)| = 0. \end{aligned}$$

(2) is left as exercise 4.2.4. □

Proposition 4.2.7 (Multivariable Product and Quotient Rules). *Let $f : A \longrightarrow \mathbf{R}$ (where $A \subset \mathbf{R}^n$) and $g : B \longrightarrow \mathbf{R}$ (where $B \subset \mathbf{R}^n$) be functions, and let f and g differentiable at a . Then*

(1) *fg is differentiable at a with derivative*

$$D(fg)_a = f(a)Dg_a + g(a)Df_a.$$

(2) If $g(a) \neq 0$ then f/g is differentiable at a with derivative

$$D\left(\frac{f}{g}\right)_a = \frac{g(a)Df_a - f(a)Dg_a}{g(a)^2}.$$

Proof. (1) As explained in the proof of Proposition 4.2.2, a is an interior point of the domain $A \cap B$ of fg , so we need only to compute. The product function fg is the composition $p \circ (f, g)$, where $(f, g) : A \cap B \rightarrow \mathbf{R}^2$ is the mapping with component functions f and g . For any $h \in \mathbf{R}^n$, the Chain Rule and the componentwise nature of differentiation (this was exercise 4.1.4) give

$$\begin{aligned} D(fg)_a(h) &= D(p \circ (f, g))_a(h) = (Dp_{(f,g)(a)} \circ D(f, g)_a)(h) \\ &= Dp_{(f(a), g(a))}(Df_a(h), Dg_a(h)), \end{aligned}$$

and by the lemma,

$$\begin{aligned} Dp_{(f(a), g(a))}(Df_a(h), Dg_a(h)) &= f(a)Dg_a(h) + g(a)Df_a(h) \\ &= (f(a)Dg_a + g(a)Df_a)(h). \end{aligned}$$

This proves (1) since h is arbitrary. (2) is similar (exercise 4.2.5) but with the wrinkle that one needs to show that since $g(a) \neq 0$ and since Dg_a exists, it follows that a is an interior point of the domain of f/g . Here it is relevant that (as noted after Lemma 4.2.5) g must be continuous at a , and so by the Persistence of Inequality principle (Proposition 2.3.9), g is nonzero on some ε -ball at a as desired. \square

With the results accumulated so far, we can compute the derivative of any mapping whose component functions are given by rational expressions in its component input scalars. By the componentwise nature of differentiability, it suffices to find the derivatives of the component functions. Since these are compositions of sums, products, and reciprocals of constants and linear functions, their derivatives are calculable with the existing machinery.

Suppose, for instance, that $f(x, y) = (x^2 - y)/(y + 1)$ for all $(x, y) \in \mathbf{R}^2$ such that $y \neq -1$. Note that every point of the domain of f is an interior point. Rewrite f as

$$f = \frac{X^2 - Y}{Y + 1}$$

where X is the linear function $X(x, y) = x$ on \mathbf{R}^2 and similarly $Y(x, y) = y$. Applications of the Chain Rule and virtually every other result on derivatives so far shows that at any point (a, b) in the domain of f , the derivative $Df_{(a,b)}$ is given by (justify the steps)

$$\begin{aligned}
Df_{(a,b)}(h, k) &= \frac{(Y+1)(a,b)D(X^2 - Y)_{(a,b)} - (X^2 - Y)(a,b)D(Y+1)_{(a,b)}}{((Y+1)(a,b))^2}(h, k) \\
&= \frac{(b+1)(D(X^2)_{(a,b)} - DY_{(a,b)}) - (a^2 - b)(DY_{(a,b)} + D1_{(a,b)})}{(b+1)^2}(h, k) \\
&= \frac{(b+1)(2X(a,b)DX_{(a,b)} - Y) - (a^2 - b)Y}{(b+1)^2}(h, k) \\
&= \frac{(b+1)(2aX - Y) - (a^2 - b)Y}{(b+1)^2}(h, k) \\
&= \frac{(b+1)(2ah - k) - (a^2 - b)k}{(b+1)^2} \\
&= \frac{2a}{b+1}h - \frac{a^2 + 1}{(b+1)^2}k.
\end{aligned}$$

In practice this method is too unwieldy for any functions beyond the simplest, and in any case it applies only to mappings with rational component functions. But on the other hand, there is no reason to expect much in the way of computational results from our methods so far, since we have been studying the derivative based on its intrinsic characterization. In the next section we will construct the derivative in coordinates, enabling us to compute easily by drawing on the results of one-variable calculus.

For another application of the Chain Rule, let A and B be subsets of \mathbf{R}^n , and suppose that $f : A \rightarrow B$ is invertible with inverse $g : B \rightarrow A$. Suppose further that f is differentiable at $a \in A$ and that g is differentiable at $f(a)$. The composition $g \circ f$ is the identity mapping $id_A : A \rightarrow A$ which, being the restriction of a linear mapping, has the linear mapping as its derivative at a . Therefore,

$$id = D(id_A)_a = D(g \circ f)_a = Dg_{f(a)} \circ Df_a.$$

This partly shows that for invertible f as described, the linear mapping Df_a is also invertible. (A symmetric argument completes the proof by showing that also $id = Df_a \circ Dg_{f(a)}$.) Since we have methods available to check the invertibility of a linear map, we can apply this criterion once we know how to compute derivatives.

Not too much should be made of this result, however; its hypotheses are too strong. Even in the one-variable case the function $f(x) = x^3$ from \mathbf{R} to \mathbf{R} is invertible and yet has the noninvertible derivative 0 at $x = 0$. (The inverse, $g(x) = \sqrt[3]{x}$ is not differentiable at 0, so the conditions above are not met.) Besides, we would prefer a converse statement, that if the derivative is invertible then so is the mapping. This is not true, but we will see in chapter 5 that it is *locally* true, i.e., it is true in the small.

Exercises

4.2.1. Prove part (2) of Proposition 4.2.1.

4.2.2. Prove part (2) of Proposition 4.2.2.

4.2.3. Prove Lemma 4.2.4. (See exercise 3.1.15 for an outline.)

4.2.4. Prove part (2) of Lemma 4.2.6.

4.2.5. Prove the Quotient Rule.

4.2.6. Let $f(x, y, z) = xyz$. Find $Df_{(a,b,c)}$ for arbitrary $(a, b, c) \in \mathbf{R}^3$. (Hint: f is the product XYZ where X is the linear function $X(x, y, z) = x$ similarly for Y and Z .)

4.2.7. Define $f(x, y) = xy^2/(y - 1)$ on $\{(x, y) \in \mathbf{R}^2 : y \neq 1\}$. Find $Df_{(a,b)}$ where (a, b) is a point in the domain of f .

4.2.8. (A generalization of the product rule.) Recall that a function

$$f : \mathbf{R}^n \times \mathbf{R}^n \longrightarrow \mathbf{R}$$

is called bilinear if for all $x, x', y, y' \in \mathbf{R}^n$ and all $\alpha \in \mathbf{R}$,

$$\begin{aligned} f(x + x', y) &= f(x, y) + f(x', y), \\ f(x, y + y') &= f(x, y) + f(x, y'), \\ f(\alpha x, y) &= \alpha f(x, y) = f(x, \alpha y). \end{aligned}$$

(a) Show that if f is bilinear then $\lim_{(h,k) \rightarrow (\mathbf{0}_n, \mathbf{0}_n)} \frac{|f(h, k)|}{|(h, k)|} = 0$.

(b) Show that if f is bilinear then f is differentiable with $Df_{(a,b)}(h, k) = f(a, k) + f(h, b)$.

(c) What does this exercise say about the inner product?

4.2.9. (A bigger generalization of the product rule.) A function

$$f : \mathbf{R}^n \times \dots \times \mathbf{R}^n \longrightarrow \mathbf{R}$$

(there are k copies of \mathbf{R}^n) is called **multilinear** if for each $j \in \{1, \dots, k\}$, for all $x_1, \dots, x_j, x'_j, \dots, x_k \in \mathbf{R}^n$ and all $\alpha \in \mathbf{R}$,

$$\begin{aligned} f(x_1, \dots, x_j + x'_j, \dots, x_k) &= f(x_1, \dots, x_j, \dots, x_k) + f(x_1, \dots, x'_j, \dots, x_k) \\ f(x_1, \dots, \alpha x_j, \dots, x_k) &= \alpha f(x_1, \dots, x_j, \dots, x_k). \end{aligned}$$

(a) Show that if f is multilinear and $a_1, \dots, a_k \in \mathbf{R}^n$, $h_1, \dots, h_k \in \mathbf{R}^n$ then for any distinct indices $i, j \in \{1, \dots, k\}$,

$$\lim_{(h_1, \dots, h_k) \rightarrow (\mathbf{0}_n, \dots, \mathbf{0}_n)} \frac{|f(a_1, \dots, h_i, \dots, h_j, \dots, a_k)|}{|(h_1, \dots, h_n)|} = 0.$$

(Use the previous problem.)

(b) Show that if f is multilinear then f is differentiable with

$$Df_{(a_1, \dots, a_k)}(h_1, \dots, h_k) = \sum_{j=1}^k f(a_1, \dots, a_{j-1}, h_j, a_{j+1}, \dots, a_k).$$

(c) When $k = n$, what does this exercise say about the determinant?

4.3 Calculating the Derivative

Working directly from Definition 4.1.2 of the multivariable derivative without using coordinates has yielded some easy results and one harder one—the Chain Rule—but no explicit description of the derivative except in the simplest cases. We don't even know that any multivariable derivatives exist except for mappings with rational coefficient functions.

Following the general principle that necessary conditions are more easily obtained than sufficient ones, we assume that the derivative exists and determine what it then must be. Geometry provides the insight. By the usual componentwise argument, there is no loss in studying a function f with scalar output, i.e., way may take $m = 1$. Setting $n = 2$ fits the graph of f in \mathbf{R}^3 where we can see it. Thus take $f : A \rightarrow \mathbf{R}$ where $A \subset \mathbf{R}^2$.

Suppose that f is differentiable at the point (a, b) . This means that the graph of f has a well-fitting tangent plane \mathcal{P} at the point $(a, b, f(a, b))$, as shown earlier in figure 4.3. To determine this plane, we need two of its lines through $(a, b, f(a, b))$. The natural lines to consider are those whose (x, y) -shadows run in the x and y directions. Call them ℓ_x and ℓ_y . (See figure 4.5.)

The line ℓ_x is tangent to a cross section of the graph of f . To see this cross section, freeze the variable y at the value b and look at the resulting function of one variable, $\varphi(x) = f(x, b)$. The slope of ℓ_x in the vertical (x, b, z) -plane is precisely $\varphi'(a)$. A slightly subtle issue here is that since (a, b) is an interior point of A , also a is an interior point of the domain of φ .

Similarly, ℓ_y has slope $\psi'(b)$ where $\psi(y) = f(a, y)$. The linear function approximating $f(a + h, b + k) - f(a, b)$ for small (h, k) is now specified as $T(h, k) = \varphi'(a)h + \psi'(b)k$. Thus $Df_{(a, b)}$ has matrix $[\varphi'(a), \psi'(b)]$. Since the entries of this matrix are simply one variable derivatives, this is something we can compute.

Definition 4.3.1 (Partial Derivative). *Let A be a subset of \mathbf{R}^n , let $f : A \rightarrow \mathbf{R}$ be a function, and let $a = (a_1, \dots, a_n)$ be an interior point of A . Fix $j \in \{1, \dots, n\}$. Define*

$$\varphi(t) = f(a_1, \dots, a_{j-1}, t, a_{j+1}, \dots, a_n) \quad \text{for } t \text{ near } a_j.$$




Figure 4.5. Cross-sectional lines

Then the j th partial derivative of f at a is defined as

$$D_j f(a) = \varphi'(a_j)$$

if $\varphi'(a_j)$ exists. Here the prime signifies ordinary one variable differentiation. Equivalently,

$$D_j f(a) = \lim_{t \rightarrow 0} \frac{f(a + te_j) - f(a)}{t}$$

if the limit exists and it is not being taken at an endpoint of the domain of the difference quotient.

Partial derivatives are easy to compute: fix all but one of the variables, and then take the one-variable derivative with respect to the variable that remains. For example if

$$f(x, y, z) = e^y \cos x + z$$

then

$$D_1 f(a, b, c) = \frac{d}{dx}(e^b \cos x + c)|_{x=a} = -e^b \sin a,$$

$$D_2 f(a, b, c) = e^b \cos a,$$

$$D_3 f(a, b, c) = 1.$$

Theorem 4.3.2 (The Derivative in Coordinates: Necessity). Let the mapping $f : A \rightarrow \mathbf{R}^m$ (where $A \subset \mathbf{R}^n$) be differentiable at the point $a \in A$. Then for each $i \in \{1, \dots, m\}$ and $j \in \{1, \dots, n\}$, the partial derivative $D_j f_i(a)$ exists. Furthermore, each $D_j f_i(a)$ is the (i, j) th entry of the Jacobian matrix of f at a . Thus the Jacobian matrix is

$$f'(a) = [D_j f_i(a)]_{i=1, \dots, m, j=1, \dots, n}.$$

Proof. The idea is to read off the (i, j) th entry of $f'(a)$ by studying the i th component function of f and letting $h \rightarrow \mathbf{0}_n$ along the j th coordinate direction in the defining property (4.3) of the derivative. The ensuing calculation will essentially reverse the steps that led us from the one-variable derivative to the definition of the multivariable derivative at the very beginning of this chapter.

The derivative of the component function f_i at a is described by the i th row of $f'(a)$. Call the row entries (d_{i1}, \dots, d_{in}) . Denoting the rows in this fashion has the operational consequence that

$$(Df_i)_a(te_j) = d_{ij}t \quad \text{for all } t \in \mathbf{R}.$$

Let $h = te_j$ with t a variable real number, so that $h \rightarrow \mathbf{0}_n$ as $t \rightarrow 0_{\mathbf{R}}$. Since $(Df_i)_a$ exists, we have as a particular instance of the defining characterization (4.3),

$$\begin{aligned} 0 &= \lim_{t \rightarrow 0} \frac{|f_i(a + te_j) - f_i(a) - (Df_i)_a(te_j)|}{|te_j|} \\ &= \lim_{t \rightarrow 0} \left| \frac{f_i(a + te_j) - f_i(a) - d_{ij}t}{t} \right| \\ &= \lim_{t \rightarrow 0} \left| \frac{f_i(a + te_j) - f_i(a)}{t} - d_{ij} \right|. \end{aligned}$$

That is,

$$\lim_{t \rightarrow 0} \frac{f_i(a + te_j) - f_i(a)}{t} = d_{ij}.$$

This says precisely that $D_j f_i(a)$ exists and equals d_{ij} . □

So the existence of the derivative Df_a makes *necessary* the existence of all partial derivatives of all component functions of f at a . The natural question is whether their existence is also *sufficient* for the existence of Df_a . Regrettably it is not. For example, the function

$$f : \mathbf{R}^2 \longrightarrow \mathbf{R}, \quad f(x, y) = \begin{cases} \frac{2xy}{x^2 + y^2} & \text{if } (x, y) \neq (0, 0), \\ 0 & \text{if } (x, y) = (0, 0) \end{cases}$$

has for its first partial derivative at the origin

$$D_1 f(0, 0) = \lim_{t \rightarrow 0} \frac{f(t, 0) - f(0, 0)}{t} = \lim_{t \rightarrow 0} \frac{0 - 0}{t} = 0,$$

and similarly $D_2 f(0, 0) = 0$; but as discussed in chapter 2, f is not continuous at the origin, much less differentiable there. However, this example is contrived, the sort of function that one sees only in a mathematics class, and in fact with slightly stronger hypotheses a statement in the spirit of the converse to Theorem 4.3.2 does hold.

Theorem 4.3.3 (The Derivative in Coordinates: Sufficiency). *Let $f : A \rightarrow \mathbf{R}^m$ (where $A \subset \mathbf{R}^n$) be a mapping, and let a be an interior point of A . Suppose that for each $i \in \{1, \dots, m\}$ and $j \in \{1, \dots, n\}$, the partial derivative $D_j f_i$ exists not only at a but at all points in some ε -ball about a , and the partial derivative $D_j f_i$ is continuous at a . Then f is differentiable at a .*

Note that if f meets the conditions of Theorem 4.3.3 then the theorem's conclusion shows that f also meets the condition of Theorem 4.3.2, so that we can use the latter theorem to calculate the derivative of f .

The difference between the necessary conditions in Theorem 4.3.2 and the sufficient conditions in Theorem 4.3.3 has a geometric interpretation when $n = 2$ and $m = 1$. The necessary conditions in Theorem 4.3.2 are that if a graph has a well-fitting plane at some point, then at that point we see well-fitting lines in the cross sections parallel to the coordinate axes. The sufficient conditions in Theorem 4.3.3 are that if we see well-fitting lines in the cross sections at and near the point, and if those lines don't change much as we move among cross sections at and near the point, then the graph has a well-fitting plane.

Proof. It suffices to show the differentiability of each component function f_i , so we may assume that $m = 1$, i.e., that f is scalar-valued. To thin out the notation, the proof will be done for $n = 2$ (so for example $a = (a_1, a_2)$), but its generality should be clear.

The Mean Value Theorem will play a crucial role. Recall its statement that if the continuous one-variable function $\varphi : [\alpha, \beta] \rightarrow \mathbf{R}$ is differentiable on (α, β) , then there exists some point $t \in (\alpha, \beta)$ such that $\varphi(\beta) - \varphi(\alpha) = \varphi'(t)(\beta - \alpha)$.

Theorem 4.3.2 says that if the derivative Df_a exists then it is defined by the matrix of partial derivatives $D_j f(a)$. The goal therefore is to show that the linear mapping

$$T_a(h_1, h_2) = D_1 f(a)h_1 + D_2 f(a)h_2$$

satisfies the defining property (4.3) of the derivative. As usual, we need to study the quantity

$$\frac{|f(a+h) - f(a) - T_a(h)|}{|h|}$$

for small h . In particular, h may be assumed small enough that the partial derivatives $D_j f$ are defined at all points within distance $|h|$ of a . This is where the hypothesis that the partial derivatives are defined everywhere near a is used. Now the idea is to move from a to $a+h$ in steps, changing one coordinate at a time, and to apply the Mean Value Theorem in each direction. Specifically,

$$\begin{aligned} |f(a+h) - f(a) - T_a(h)| &= |f(a_1 + h_1, a_2 + h_2) - f(a_1, a_2 + h_2) - D_1 f(a)h_1 \\ &\quad + f(a_1, a_2 + h_2) - f(a_1, a_2) - D_2 f(a)h_2|, \end{aligned}$$

and two applications of the Mean Value Theorem to the right side give

$$|f(a+h) - f(a) - T_a(h)| = |D_1f(a_1+t_1, a_2+h_2)h_1 - D_1f(a)h_1 \\ + D_2f(a_1, a_2+t_2)h_2 - D_2f(a)h_2|$$

where $|t_j| \leq |h_j| \leq |h|$ for $j = 1, 2$, so that now by the Triangle Inequality,

$$|f(a+h) - f(a) - T_a(h)| \leq |D_1f(a_1+t_1, a_2+h_2) - D_1f(a)| |h| \\ + |D_2f(a_1, a_2+t_2) - D_2f(a)| |h|.$$

Therefore for small h ,

$$\frac{|f(a+h) - f(a) - T_a(h)|}{|h|} \leq |D_1f(a_1+t_1, a_2+h_2) - D_1f(a)| \\ + |D_2f(a_1, a_2+t_2) - D_2f(a)|.$$

As $h \rightarrow \mathbf{0}_2$ the h_j and t_j are squeezed to 0 in \mathbf{R} , and the continuity of the partial derivatives D_jf at a forces the right side to 0. This proves the theorem. \square

Thus,

- the differentiability of f at a implies the existence of all partial derivatives at a ,
- while the existence of all partial derivatives at and about a , and their continuity at a , combine to imply the differentiability of f at a .

Note how this all compares to the discussion of the determinant in the previous chapter. There we wanted the determinant to satisfy characterizing properties, we found the only function that could possibly satisfy them, and then we verified that it did. Here we wanted the derivative to satisfy a characterizing property, and we found the only possibility for the derivative: the linear mapping whose matrix consists of the partial derivatives, which must exist if the derivative does. But analysis is more subtle than algebra: this linear mapping need not satisfy the characterizing property of the derivative unless we add further assumptions. Theorem 4.3.3 is the most substantial result so far in this chapter, since it guarantees that the derivative exists. Its proof was correspondingly the deepest, requiring the abstract existence statement supplied by the Mean Value Theorem. The converse to Theorem 4.3.3 is not true (exercise 4.3.3).

For an example of all this, consider the function

$$f(x, y) = \begin{cases} \frac{x^2y}{x^2+y^2} & \text{if } (x, y) \neq (0, 0), \\ 0 & \text{if } (x, y) = (0, 0). \end{cases}$$

The top formula in the definition describes a rational function of x and y on the punctured plane $\mathbf{R}^2 - \{(0, 0)\}$. Any rational function and all of its

partial derivatives are continuous on its domain (feel free to invoke this result), and furthermore every point (a, b) away from $(0, 0)$ lies in some ε -ball that is also away from $(0, 0)$. That is, for any point $(a, b) \neq (0, 0)$ the partial derivatives of f exist at and about (a, b) and they are continuous at (a, b) . Thus the conditions for Theorem 4.3.3 are met, and so its conclusion follows: f is differentiable at (a, b) . Now Theorem 4.3.2 says that the derivative matrix at (a, b) is the matrix of partial derivatives,

$$f'(a, b) = [D_1f(a, b) \ D_2f(a, b)] = \left[\frac{2ab^3}{(a^2+b^2)^2} \ \frac{a^2(a^2-b^2)}{(a^2+b^2)^2} \right].$$

Consequently the derivative of f at any nonzero (a, b) is the corresponding linear map

$$Df_{(a,b)}(h, k) = \frac{2ab^3}{(a^2+b^2)^2}h + \frac{a^2(a^2-b^2)}{(a^2+b^2)^2}k.$$

However, this analysis breaks down at the point $(a, b) = (0, 0)$. Here our only recourse is to figure out whether a candidate derivative exists and then test whether it works. The first partial derivative of f at $(0, 0)$ is

$$D_1f(0, 0) = \lim_{t \rightarrow 0} \frac{f(t, 0) - f(0, 0)}{t} = \lim_{t \rightarrow 0} \frac{0 - 0}{t} = 0,$$

and similarly $D_2f(0, 0) = 0$. So by Theorem 4.3.2, the only possibility for the derivative of f at $(0, 0)$ is the zero map. Now the question is whether the limit

$$\lim_{(h,k) \rightarrow (0,0)} \frac{|f(h, k) - f(0, 0) - 0|}{|(h, k)|}$$

exists and is 0. Compute that since the denominator $h^2 + k^2$ of f away from the origin is $|(h, k)|^2$,

$$\frac{|f(h, k) - f(0, 0) - 0|}{|(h, k)|} = \frac{|h|^2|k|}{|(h, k)|^3}.$$

As in section 2.3, the straight line test can only suggest that this quantity does go to 0, but it can prove that the quantity doesn't go to 0; and the Size Bounds can only suggest that this quantity doesn't go to 0, but they can prove that the quantity does go to 0. In this case we use the straight line test. Let $k = h$ and compute that

$$\frac{|h|^2|h|}{|(h, h)|^3} = \frac{|h|^3}{2^{3/2}|h|^3} = \frac{1}{2^{3/2}}.$$

This does not go to 0 as h goes to 0. Therefore the function f is not differentiable at $(0, 0)$. And indeed, the graph of f near $(0, 0)$ shows a shape that isn't well approximated by any plane through its center, no matter how closely we zoom in. (See figure 4.6.)




Figure 4.6. The manta ray is differentiable everywhere except at the origin

Returning to the discussion of invertibility of a mapping and invertibility of its derivative, let $f : \mathbf{R}^2 - \{(0, 0)\} \rightarrow \mathbf{R}^2 - \{(0, 0)\}$ be given by $f(x, y) = (x^2 - y^2, 2xy)$. At any (x, y) where f is defined, the partial derivatives are $D_1 f_1(x, y) = 2x$, $D_2 f_1(x, y) = -2y$, $D_1 f_2(x, y) = 2y$, and $D_2 f_2(x, y) = 2x$. These are continuous functions of (x, y) , so for any $(a, b) \neq (0, 0)$, $Df_{(a,b)}$ exists and its matrix is

$$f'(a, b) = \begin{bmatrix} D_1 f_1(a, b) & D_2 f_1(a, b) \\ D_1 f_2(a, b) & D_2 f_2(a, b) \end{bmatrix} = \begin{bmatrix} 2a & -2b \\ 2b & 2a \end{bmatrix}.$$

This has determinant $4(a^2 + b^2) > 0$, and hence it is always invertible. On the other hand, the mapping f takes the same value at points (x, y) and $-x, -y$, so it is definitely not invertible.

With the Jacobian matrix described explicitly, a more calculational version of the Chain Rule is available.

Theorem 4.3.4 (Chain Rule in Coordinates). *Let $f : A \rightarrow \mathbf{R}^m$ (where $A \subset \mathbf{R}^n$) be differentiable at the point a of A , and let $g : f(A) \rightarrow \mathbf{R}$ be differentiable at the point $b = f(a)$. Then the composition $g \circ f : A \rightarrow \mathbf{R}$ is differentiable at a , and its partial derivatives are*

$$D_j(g \circ f)(a) = \sum_{k=1}^m D_k g(b) D_j f_k(a) \quad \text{for } j = 1, \dots, n.$$

Proof. The composition is differentiable by the intrinsic Chain Rule. The Jacobian matrix of g at b is

$$g'(b) = [D_1 g(b) \cdots D_m g(b)],$$

and the Jacobian matrix of f at a is

$$f'(a) = \begin{bmatrix} D_1 f_1(a) & \cdots & D_n f_1(a) \\ \vdots & \ddots & \vdots \\ D_1 f_m(a) & \cdots & D_n f_m(a) \end{bmatrix},$$

and the Jacobian matrix of $g \circ f$ at a is

$$(g \circ f)'(a) = [D_1(g \circ f)(a) \cdots D_n(g \circ f)(a)].$$

By the intrinsic Chain Rule,

$$(g \circ f)'(a) = g'(b)f'(a).$$

Equate the j th entries to obtain the result. \square

With slightly stronger assumptions about the partial derivatives, the formula for the partial derivatives of a composition (which may exist without the composition being differentiable) can be proved without recourse to the intrinsic version of the Chain Rule. One studies the difference quotient

$$\frac{(g \circ f)(a + te_j) - (g \circ f)(a)}{t},$$

inserting intermediate terms and using the Mean Value Theorem as in the proof of Theorem 4.3.3. This is exercise 4.3.5.


Notations for the partial derivative vary. A function is often described by a formula such as $w = f(x, y, z)$. Other notations for $D_1 f$ are

$$f_1, \quad f_x, \quad \frac{\partial f}{\partial x}, \quad w_x, \quad \frac{\partial w}{\partial x}.$$

If x, y, z are in turn functions of s and t then a classical formulation of the Chain Rule would be

$$\frac{\partial w}{\partial t} = \frac{\partial w}{\partial x} \frac{\partial x}{\partial t} + \frac{\partial w}{\partial y} \frac{\partial y}{\partial t} + \frac{\partial w}{\partial z} \frac{\partial z}{\partial t}. \quad (4.5)$$


This is easily visualized as chasing back along all “dependency chains” from t to w in a diagram where an arrow means “contributes to”:



Unfortunately, for all its mnemonic advantages, the classical notation is a veritable minefield of misinterpretation. Formula (4.5) doesn't indicate where the various partial derivatives are to be evaluated, for one thing. Specifying the variable of differentiation by name rather than by position also becomes confusing when different symbols are substituted for the same variable, especially since the symbols themselves may denote specific values or other variables. For example one can construe many different meanings for the expression

$$\frac{\partial f}{\partial x}(y, x, z).$$

Blurring the distinction between functions and the variables denoting their outputs is even more problematic. If one has, say, $z = f(x, t, u)$, $x = g(t, u)$,



then chasing all paths back to t gives

$$\frac{\partial z}{\partial t} = \frac{\partial z}{\partial x} \frac{\partial x}{\partial t} + \frac{\partial z}{\partial t}$$

with “ $\partial z/\partial t$ ” meaning something different on each side of the equality. While the classical formulas are useful and perhaps simpler to apply in elementary situations, they are not particularly robust until one has a solid understanding of the Chain Rule. On the other hand, the classical formulas work fine in straightforward applications, so several exercises are phrased in the older language to give you practice with it.

Exercises

4.3.1. Explain why in the discussion beginning this section the tangent plane \mathcal{P} consists of all points $(a, b, f(a, b)) + (h, k, T(h, k))$ where $T(h, k) = \varphi'(a)h + \psi'(b)k$.

4.3.2. This exercise shows that all partial derivatives of a function can exist at and about a point without being continuous at the point. Define $f : \mathbf{R}^2 \rightarrow \mathbf{R}$ by

$$f(x, y) = \begin{cases} \frac{2xy}{x^2+y^2} & \text{if } (x, y) \neq (0, 0), \\ 0 & \text{if } (x, y) = (0, 0). \end{cases}$$

- Show that $D_1 f(0, 0) = D_2 f(0, 0) = 0$.
- Show that $D_1 f(a, b)$ and $D_2 f(a, b)$ exist and are continuous at all other $(a, b) \in \mathbf{R}^2$.
- Show that $D_1 f$ and $D_2 f$ are discontinuous at $(0, 0)$.

4.3.3. Define $f : \mathbf{R} \rightarrow \mathbf{R}$ by

$$f(x) = \begin{cases} x^2 \sin \frac{1}{x} & \text{if } x \neq 0, \\ 0 & \text{if } x = 0. \end{cases}$$

Show that $f'(x)$ exists for all x but that f' is discontinuous at 0. Explain how this disproves the converse of Theorem 4.3.3.

4.3.4. Discuss the derivatives of the following mappings at the following points.

(a) $f(x, y) = \frac{x^2 - y}{y + 1}$ on $\{(x, y) \in \mathbf{R}^2 : y \neq -1\}$ at generic (a, b) with $b \neq -1$.

(b) $f(x, y) = \frac{xy^2}{y - 1}$ on $\{(x, y) \in \mathbf{R}^2 : y \neq 1\}$ at generic (a, b) with $b \neq 1$.

(c) $f(x, y) = \begin{cases} \frac{xy}{\sqrt{x^2 + y^2}} & \text{if } (x, y) \neq (0, 0) \\ 0 & \text{if } (x, y) = (0, 0) \end{cases}$ at generic $(a, b) \neq (0, 0)$ and at $(0, 0)$.

4.3.5. As discussed after Theorem 4.3.4, derive the formula for the partial derivatives of a composition without recourse to the Chain Rule by making stronger hypotheses about the partial derivatives of the mappings being composed. (Suggestion: set $n = 2$ and then mimic the proof of Theorem 4.3.3 as described in the text, keeping track of what hypotheses you need as you go.)

For the rest of these exercises, assume as much differentiability as necessary.

4.3.6. For what differentiable mappings $f : A \rightarrow \mathbf{R}^m$ is $f'(a)$ a diagonal matrix for all $a \in A$? (A diagonal matrix is a matrix whose (i, j) th entries for all $i \neq j$ are 0.)

4.3.7. Show that if $z = f(xy)$ then x , y , and z satisfy the differential equation $x \cdot z_x - y \cdot z_y = 0$.

4.3.8. Let $w = F(xz, yz)$. Show that $x \cdot w_x + y \cdot w_y = z \cdot w_z$.

4.3.9. If $z = f(ax + by)$, show that $bz_x = az_y$.

4.3.10. The function $f : \mathbf{R}^2 \rightarrow \mathbf{R}$ is called **homogeneous of degree k** if $f(tx, ty) = t^k f(x, y)$ for all scalars t and vectors (x, y) . Show that such f satisfies the differential equation

$$x f_1(x, y) + y f_2(x, y) = k f(x, y).$$

4.3.11. Let

$$f : \mathbf{R}^2 \rightarrow \mathbf{R}$$

be a function such that for all $(x, y) \in \mathbf{R}^2$, the integral

$$F : \mathbf{R}^2 \rightarrow \mathbf{R}, \quad F(x, y) = \int_{v=0}^y f(x, v) dv$$

exists and is differentiable with respect to x , its partial derivative with respect to x being obtained by passing the x -derivative through the v -integral,

$$\begin{aligned} \frac{\partial F(x, y)}{\partial x} &= \frac{\partial}{\partial x} \int_{v=0}^y f(x, v) dv \\ &= \lim_{h \rightarrow 0} \frac{\int_{v=0}^y f(x+h, v) dv - \int_{v=0}^y f(x, v) dv}{h} \\ &= \lim_{h \rightarrow 0} \int_{v=0}^y \frac{f(x+h, v) - f(x, v)}{h} dv \\ &\stackrel{!}{=} \int_{v=0}^y \lim_{h \rightarrow 0} \frac{f(x+h, v) - f(x, v)}{h} dv \\ &= \int_{v=0}^y \frac{\partial f}{\partial x}(x, v) dv. \end{aligned}$$

(The “!” step requires justification, but under reasonable circumstances it can be carried out.) Define a function

$$G : \mathbf{R} \rightarrow \mathbf{R}, \quad G(x) = \int_{v=0}^x f(x, v) dv.$$

Thus x affects G in two ways: as a parameter for the integrand, and as the upper limit of integration. What is $dG(x)/dx$?

4.4 Higher Order Derivatives

Partial differentiation can be carried out more than once on nice enough functions. For example if $f(x, y) = e^{x \sin y}$ then

$$D_1 f(x, y) = \sin y e^{x \sin y}, \quad D_2 f(x, y) = x \cos y e^{x \sin y}.$$

Taking partial derivatives again yields

$$\begin{aligned} D_1 D_1 f(x, y) &= \sin^2 y e^{x \sin y}, \\ D_1 D_2 f(x, y) &= \cos y e^{x \sin y} + x \sin y \cos y e^{x \sin y}, \\ D_2 D_1 f(x, y) &= \cos y e^{x \sin y} + x \sin y \cos y e^{x \sin y} = D_1 D_2 f(x, y), \\ D_2 D_2 f(x, y) &= -x \sin y e^{x \sin y} + x^2 \cos^2 y e^{x \sin y}, \end{aligned}$$

and some partial derivatives of these in turn are,

$$\begin{aligned}
D_1 D_1 D_2 f(x, y) &= 2 \sin y \cos y e^{x \sin y} + x \sin^2 y \cos y e^{x \sin y}, \\
D_1 D_2 D_1 f(x, y) &= D_1 D_1 D_2 f(x, y), \\
D_2 D_1 D_2 f(x, y) &= -\sin y e^{x \sin y} + 2x \cos^2 y e^{x \sin y} - x \sin^2 y e^{x \sin y} \\
&\quad + x^2 \sin y \cos^2 y e^{x \sin y}, \\
D_2 D_2 D_1 f(x, y) &= D_2 D_1 D_2 f(x, y), \\
D_1 D_2 D_2 f(x, y) &= -\sin y e^{x \sin y} + 2x \cos^2 y e^{x \sin y} - x \sin^2 y e^{x \sin y} \\
&\quad + x^2 \sin y \cos^2 y e^{x \sin y} \\
&= D_2 D_1 D_2 f(x, y), \\
D_2 D_1 D_1 f(x, y) &= 2 \sin y \cos y e^{x \sin y} + x \sin^2 y \cos y e^{x \sin y} \\
&= D_1 D_1 D_2 f(x, y).
\end{aligned}$$

Suspiciously many of these match. The result of two or three partial differentiations seems to depend only on how many were taken with respect to x and how many with respect to y , not on the order in which they were taken.

To analyze the situation, it suffices to consider only two differentiations. Streamline the notation by writing $D_2 D_1 f$ as $D_{12} f$. (The subscripts may look reversed, but reading D_{12} from left to right as “ D -one-two” suggests the appropriate order of differentiating.) The definitions for $D_{11} f$, $D_{21} f$, and $D_{22} f$ are similar. These four functions are called the **second order** partial derivatives of f , and in particular $D_{12} f$ and $D_{21} f$ are the second order **mixed** partial derivatives. More generally, the k th order partial derivatives of a function f are those that come from k partial differentiations. A \mathcal{C}^k -**function** is a function for which all the k th order partial derivatives exist and are continuous. The theorem is that with enough continuity the order of differentiation doesn’t matter. That is, the mixed partial derivatives agree.

Theorem 4.4.1 (Equality of Mixed Partial Derivatives). *Suppose that $f : A \rightarrow \mathbf{R}$ (where $A \subset \mathbf{R}^2$) is a \mathcal{C}^2 function. Then at any point (a, b) of A ,*

$$D_{12} f(a, b) = D_{21} f(a, b).$$

This theorem is similar to Taylor’s Theorem from section 1.3 in that both are stated entirely in terms of derivatives, but they are tricky to prove with derivatives and easier to prove with integrals. We give the easier proof. Extending Theorem 4.4.1 to more variables and to higher derivatives is straightforward provided that one supplies enough continuity. The hypotheses of the theorem can be weakened a bit, in which case a more subtle proof is required, but such technicalities are more distracting than useful.

Proof. Since f is a \mathcal{C}^2 -function on A , every point of A is interior. Take any point $(a, b) \in A$. Then some box $B = [a, a + h] \times [b, b + k]$ lies in A . Compute the nested integral

$$\int_a^{a+h} \int_b^{b+k} dy dx = \int_a^{a+h} k dx = hk.$$

Also, by the Fundamental Theorem of Integral Calculus twice,

$$\begin{aligned} \int_a^{a+h} \int_b^{b+k} D_{12}f(x, y) dy dx &= \int_a^{a+h} (D_1f(x, b+k) - D_1f(x, b)) dx \\ &= f(a+h, b+k) - f(a, b+k) - f(a+h, b) + f(a, b). \end{aligned}$$

Call this quantity $\Delta(h, k)$. Let $m_{h,k}$ be the minimum value of $D_{12}f$ on the box B , and let $M_{h,k}$ be the maximum value. (These exist by Theorem 2.4.15 since B is nonempty compact and $D_{12}f : B \rightarrow \mathbf{R}$ is continuous.) Thus

$$m_{h,k} \leq D_{12}f(x, y) \leq M_{h,k} \quad \text{for all } (x, y) \in B.$$

Integrate this inequality, using the two previous calculations, to get

$$m_{h,k}hk \leq \Delta(h, k) \leq M_{h,k}hk,$$

or

$$m_{h,k} \leq \frac{\Delta(h, k)}{hk} \leq M_{h,k}.$$

As $(h, k) \rightarrow (0^+, 0^+)$, the continuity of $D_{12}f$ at (a, b) forces $m_{h,k}$ and $M_{h,k}$ to $D_{12}f(a, b)$, and hence

$$\frac{\Delta(h, k)}{hk} \rightarrow D_{12}f(a, b) \quad \text{as } (h, k) \rightarrow (0^+, 0^+).$$

But also, reversing the order of the integrations and of the partial derivatives gives the symmetric calculations

$$\int_b^{b+k} \int_a^{a+h} dx dy = hk,$$

and

$$\int_b^{b+k} \int_a^{a+h} D_{21}f(x, y) dx dy = \Delta(h, k),$$

and so the same argument shows that

$$\frac{\Delta(h, k)}{hk} \rightarrow D_{21}f(a, b) \quad \text{as } (h, k) \rightarrow (0^+, 0^+).$$

Since both $D_{12}f(a, b)$ and $D_{21}f(a, b)$ are the limit of $\Delta(h, k)/(hk)$, they are equal. \square

Higher order derivatives are written in many ways. If a function is described by the equation $w = f(x, y, z)$ then $D_{233}f$ is also denoted

$$f_{233}, \quad f_{yzz}, \quad \frac{\partial}{\partial z} \left(\frac{\partial}{\partial z} \left(\frac{\partial f}{\partial y} \right) \right), \quad \frac{\partial^3 f}{\partial z^2 \partial y},$$

$$w_{yzz}, \quad \frac{\partial}{\partial z} \left(\frac{\partial}{\partial z} \left(\frac{\partial w}{\partial y} \right) \right), \quad \frac{\partial^3 w}{\partial z^2 \partial y}.$$

As with one derivative, these combine mnemonic advantages with conceptual dangers.

A calculation using higher order derivatives and the Chain Rule transforms the heat equation of Laplace from cartesian to polar coordinates. The C^2 quantity $u = f(x, y)$ depending on the cartesian variables x and y satisfies **Laplace's equation** if (blurring the distinction between u and f)


$$\frac{\partial^2 u}{\partial x^2} + \frac{\partial^2 u}{\partial y^2} = 0.$$

If instead u is viewed as a function g of the polar variables r and θ , how is Laplace's equation expressed?

The cartesian coordinates in terms of the polar coordinates are

$$x = r \cos \theta, \quad y = r \sin \theta.$$

Thus $u = f(x, y) = f(r \cos \theta, r \sin \theta) = g(r, \theta)$, showing that u depends on r and θ via x and y :



The Chain Rule begins a hieroglyphic calculation,

$$u_r = u_x x_r + u_y y_r,$$

so that by the product rule,

$$u_{rr} = (u_x x_r + u_y y_r)_r$$

$$= u_{xr} x_r + u_x x_{rr} + u_{yr} y_r + u_y y_{rr},$$

and then, since u_x and u_y also depend on r and θ via x and y , the Chain Rule gives expansions of u_{xr} and u_{yr} , and the calculation continues,

$$u_{rr} = (u_{xx} x_r + u_{xy} y_r) x_r + u_x x_{rr} + (u_{yx} x_r + u_{yy} y_r) y_r + u_y y_{rr}$$

$$= u_{xx} x_r^2 + u_{xy} y_r x_r + u_x x_{rr} + u_{yx} x_r y_r + u_{yy} y_r^2 + u_y y_{rr}$$

$$= u_{xx} x_r^2 + 2u_{xy} x_r y_r + u_{yy} y_r^2 + u_x x_{rr} + u_y y_{rr}.$$

Note the use of equality of mixed partial derivatives. The same calculation with θ instead of r gives

$$\begin{aligned}u_\theta &= u_x x_\theta + u_y y_\theta, \\u_{\theta\theta} &= u_{xx} x_\theta^2 + 2u_{xy} x_\theta y_\theta + u_{yy} y_\theta^2 + u_x x_{\theta\theta} + u_y y_{\theta\theta}.\end{aligned}$$

Since $x = r \cos \theta$ and $y = r \sin \theta$ we have the relations

$$x_\theta = -ry_r, \quad y_\theta = rx_r, \quad x_{rr} = y_{rr} = 0.$$

It follows that

$$\begin{aligned}r^2 u_{rr} &= r^2 u_{xx} x_r^2 + 2r^2 u_{xy} x_r y_r + r^2 u_{yy} y_r^2, \\ru_r &= u_x y_\theta - u_y x_\theta, \\u_{\theta\theta} &= r^2 u_{xx} y_r^2 - 2r^2 u_{xy} x_r y_r + r^2 u_{yy} x_r^2 - u_x y_\theta + u_y x_\theta.\end{aligned}$$

Recall that the cartesian form of Laplace's equation is $u_{xx} + u_{yy} = 0$. Now the polar form follows,

$$r^2 u_{rr} + ru_r + u_{\theta\theta} = 0.$$

That is,

$$r^2 \frac{\partial^2 u}{\partial r^2} + r \frac{\partial u}{\partial r} + \frac{\partial^2 u}{\partial \theta^2} = 0.$$

The point of this involved calculation is that having done it once, and only once, we now can check directly whether any function g of the polar variables r and θ satisfies Laplace's equation. We no longer need to transform each g into terms of the cartesian coordinates x and y before checking.

Exercises

4.4.1. This exercise shows that continuity is necessary for the equality of mixed partial derivatives. Let

$$f(x, y) = \begin{cases} \frac{xy(x^2 - y^2)}{x^2 + y^2} & \text{if } (x, y) \neq (0, 0) \\ 0 & \text{if } (x, y) = (0, 0). \end{cases}$$

Show: (a) f , $D_1 f$, and $D_2 f$ are continuous on \mathbf{R}^2 . (b) $D_{12} f$ and $D_{21} f$ exist on \mathbf{R}^2 and are continuous on $\mathbf{R}^2 - \{(0, 0)\}$. (c) $D_{12} f(0, 0) = 1 \neq -1 = D_{21} f(0, 0)$.

4.4.2. Suppose u , as a function of x and y , satisfies the differential equation $u_{xx} - u_{yy} = 0$. Make the change of variables $x = s + t$, $y = s - t$. What corresponding differential equation does u satisfy when viewed as a function of s and t ?

4.4.3. Show that if $u = F(x - ct) + G(x + ct)$ then $c^2 u_{xx} = u_{tt}$.

4.4.4. Show that the substitution $x = e^s$, $y = e^t$ converts the equation

$$x^2 u_{xx} + y^2 u_{yy} + x u_x + y u_y = 0$$

into Laplace's equation $u_{ss} + u_{tt} = 0$.

4.4.5. Show that the substitution $u = x^2 - y^2$, $v = 2xy$ converts Laplace's equation $w_{xx} + w_{yy} = 0$ back into Laplace's equation $w_{uu} + w_{vv} = 0$.

4.5 Extreme Values

In one variable calculus the derivative is used to find maximum and minimum values (extrema) of differentiable functions. Recall the following useful facts.

- (Extreme Value Theorem.) If $f : [\alpha, \beta] \rightarrow \mathbf{R}$ is continuous then it assumes a maximum and a minimum on the interval $[\alpha, \beta]$.
- (Critical Point Theorem.) Suppose that $f : [\alpha, \beta] \rightarrow \mathbf{R}$ is differentiable on (α, β) and that f assumes a maximum or minimum at an interior point a of $[\alpha, \beta]$. Then $f'(a) = 0$.
- (Second Derivative Test.) Suppose that $f : [\alpha, \beta] \rightarrow \mathbf{R}$ is \mathcal{C}^2 on (α, β) and that $f'(a) = 0$ at an interior point a of $[\alpha, \beta]$. If $f''(a) > 0$ then $f(a)$ is a local minimum of f , and if $f''(a) < 0$ then $f(a)$ is a local maximum.

Geometrically the idea is that just as the affine function

$$A(a+h) = f(a) + f'(a)h$$

specifies the tangent line to the graph of f at $(a, f(a))$, the quadratic function

$$P(a+h) = f(a) + f'(a)h + \frac{1}{2}f''(a)h^2$$

determines the best fitting parabola. When $f'(a) = 0$ the tangent line is horizontal and the sign of $f''(a)$ specifies whether the parabola opens upward or downward. When $f'(a) = 0$ and $f''(a) = 0$, the parabola degenerates to the horizontal tangent line, and the second derivative provides no information. (See figure 4.7.)




Figure 4.7. Approximating parabolas

This section generalizes these facts to functions f of n variables. The Extreme Value Theorem has already generalized as Theorem 2.4.15: a continuous function f on a compact subset of \mathbf{R}^n takes maximum and minimum values. The Critical Point Theorem also generalizes easily to say that each extreme value of the function $f : A \rightarrow \mathbf{R}$ that occurs at a point where f is differentiable occurs at **critical point** of f , meaning a point a where Df_a is the zero function.

Theorem 4.5.1 (Multivariable Critical Point Theorem). *Suppose that the function $f : A \rightarrow \mathbf{R}$ (where $A \subset \mathbf{R}^n$) takes an extreme value at the point a of A , and suppose that f is differentiable at a . Then all partial derivatives of f at a are zero.*

Proof. For each $j \in \{1, \dots, n\}$, the value $f(a)$ is an extreme value for the one-variable function φ from definition 4.3.1 of the partial derivative $D_j f(a)$. By the one-variable Critical Point Theorem, $\varphi'(a_j) = 0$. That is, $D_j f(a) = 0$. \square

The generalization of the second derivative test is more elaborate. From now on, all functions are assumed to be of type \mathcal{C}^2 on the interiors of their domains, meaning that all their second order partial derivatives exist and are continuous.

Definition 4.5.2 (Hessian matrix). *Let $f : A \rightarrow \mathbf{R}$ (where $A \subset \mathbf{R}^n$) be a function and let a be an interior point of A . The **Hessian matrix** of f at a is the n -by- n matrix whose (i, j) th entry is the second order partial derivative $D_{ij} f(a)$. Thus*

$$H_f(a) = \begin{bmatrix} D_{11} f(a) & \cdots & D_{1n} f(a) \\ \vdots & \ddots & \vdots \\ D_{n1} f(a) & \cdots & D_{nn} f(a) \end{bmatrix}.$$

By the equality of mixed partial derivatives, the Hessian is a symmetric matrix, i.e., $H_f(a)^t = H_f(a)$. Beware of confusing the Hessian and the Jacobian: the Hessian is a square matrix defined only for scalar valued functions and its entries are second order partial derivatives, while for scalar valued functions the Jacobian is the row vector of first partial derivatives. As an example, if

$$f(x, y) = \sin^2 x + x^2 y + y^2,$$

then for any $(a, b) \in \mathbf{R}^2$,

$$f'(a, b) = [\sin 2a + 2ab \ a^2 + 2b]$$

and

$$H_f(a, b) = \begin{bmatrix} 2 \cos 2a + 2b & 2a \\ 2a & 2 \end{bmatrix}.$$

Any n -by- n matrix M determines a quadratic function

$$Q_M : \mathbf{R}^n \longrightarrow \mathbf{R}, \quad Q_M(h) = h^t M h.$$

Here h is viewed as a column vector. If M has entries m_{ij} and $h = (h_1, \dots, h_n)$ then the rules of matrix multiplication show that

$$Q_M(h) = [h_1 \cdots h_n] \begin{bmatrix} m_{11} & \cdots & m_{1n} \\ \vdots & & \vdots \\ m_{n1} & \cdots & m_{nn} \end{bmatrix} \begin{bmatrix} h_1 \\ \vdots \\ h_n \end{bmatrix} = \sum_{i=1}^n \sum_{j=1}^n m_{ij} h_i h_j.$$

The function Q_M is homogeneous of degree 2, meaning that each of its terms has degree 2 in the entries of h and therefore $Q_M(th) = t^2 Q_M(h)$ for all $t \in \mathbf{R}$ and $h \in \mathbf{R}^n$.

When M is the Hessian of a function f at a point a , the quadratic function Q_M is denoted Qf_a rather than the symbol-heavy $Q_{H_f(a)}$. Just as $f(a) + Df_a(h)$ gives the best affine approximation to $f(a+h)$ for small h , $f(a) + Df_a(h) + \frac{1}{2}Qf_a(h)$ gives the best quadratic approximation.

In the example $f(x, y) = \sin^2 x + x^2 y + y^2$, the Hessian at a point (a, b) defines the function

$$\begin{aligned} Qf_{(a,b)}(h, k) &= [h \ k] \begin{bmatrix} 2 \cos 2a + 2b & 2a \\ 2a & 2 \end{bmatrix} \begin{bmatrix} h \\ k \end{bmatrix} \\ &= 2((\cos 2a + b)h^2 + 2ahk + k^2) \quad \text{for } (h, k) \in \mathbf{R}^2, \end{aligned}$$

and so the best quadratic approximation to f near, for instance, the point $(\pi/2, 1)$ is

$$\begin{aligned} f(\pi/2 + h, 1 + k) &\approx f(\pi/2, 1) + Df_{(\pi/2,1)}(h, k) + \frac{1}{2}Qf_{(\pi/2,1)}(h, k) \\ &= \pi^2/4 + 2 + \pi h + (\pi^2/4 + 2)k + \pi hk + k^2. \end{aligned}$$

Suppose that $f : A \longrightarrow \mathbf{R}$ (where $A \subset \mathbf{R}^2$) has a critical point at (a, b) , i.e., $f'(a, b) = (0, 0)$. The graphs of some quadratic functions on \mathbf{R}^2 with a critical point at $(0, 0)$ are shown in figure 4.8. If the best quadratic approximation of f at (a, b) is a bowl then f should have a minimum at (a, b) . Similarly for an inverted bowl and a maximum. If the best quadratic approximation is a saddle then there should be points (x, y) near (a, b) where $f(x, y) > f(a, b)$ and points (x', y') near (a, b) where $f(x', y') < f(a, b)$. In this case (a, b) is called for obvious reasons a **saddle point** of f .

Returning to the example $f(x, y) = \sin^2 x + x^2 y + y^2$, note that $(0, 0)$ is a critical point of f since $f'(0, 0) = (0, 0)$. The Hessian $H_f(0, 0)$ is $\begin{bmatrix} 2 & 0 \\ 0 & 2 \end{bmatrix}$, and so the quadratic function $\frac{1}{2}Qf_{(0,0)}$ is given by

$$\frac{1}{2}Qf_{(0,0)}(h, k) = \frac{1}{2} [h \ k] \begin{bmatrix} 2 & 0 \\ 0 & 2 \end{bmatrix} \begin{bmatrix} h \\ k \end{bmatrix} = h^2 + k^2.$$

Thus the graph of f looks like a bowl near $(0, 0)$ and $f(0, 0)$ should be a local minimum.




Figure 4.8. Two bowls, two saddles, four halfpipes, and a plane

None of this discussion is rigorous yet. Justifying these ideas and proving the appropriate theorems will occupy the rest of this section. The first task is to study quadratic approximations of C^2 -functions.

Proposition 4.5.3 (Special Case of Taylor's Theorem). *Let I be an open interval in \mathbf{R} containing $[0, 1]$. Let $\varphi : I \rightarrow \mathbf{R}$ be a C^2 -function. Then*

$$\varphi(1) = \varphi(0) + \varphi'(0) + \frac{1}{2}\varphi''(c) \quad \text{for some } c \in [0, 1].$$

This follows from the general Taylor's Theorem in section 1.3 since the first degree Taylor polynomial of φ at 0 is $T_1(t) = \varphi(0) + \varphi'(0)t$, so that in particular, $T_1(1) = \varphi(0) + \varphi'(0)$.

Theorem 4.5.4 (Quadratic Taylor Approximation). *Let $f : A \rightarrow \mathbf{R}$ (where $A \subset \mathbf{R}^n$) be a C^2 -function on the interior points of A . Let a be an interior point of A . Then for all small enough $h \in \mathbf{R}^n$,*

$$f(a+h) = f(a) + Df_a(h) + \frac{1}{2}Q_{f_{a+ch}}(h) \quad \text{for some } c \in [0, 1],$$

or, in matrices, viewing h as a column vector,

$$f(a+h) = f(a) + f'(a)h + \frac{1}{2}h^t H_f(a+ch)h \quad \text{for some } c \in [0, 1].$$

Proof. Define $\varphi : [0, 1] \rightarrow \mathbf{R}$ by $\varphi(t) = f(a + th)$. That is, φ is the restriction of f to the line segment from a to $a + h$. By the Chain Rule,

$$\varphi'(t) = \sum_{i=1}^n D_i f(a + th) h_i = f'(a + th)h = Df_{a+th}(h)$$

and

$$\varphi''(t) = \sum_{i=1}^n \sum_{j=1}^n D_{ij} f(a + th) h_j h_i = h^t H_f(a + th) h = Qf_{a+th}(h).$$

Since $f(a + h) = \varphi(1)$, the special case of Taylor's Theorem says that for some $c \in [0, 1]$,

$$f(a + h) = \varphi(0) + \varphi'(0) + \frac{1}{2}\varphi''(c) = f(a) + Df_a(h) + \frac{1}{2}Qf_{a+ch}(h),$$

giving the result. \square

Thus, to study f near a critical point $a \in \mathbf{R}^n$ where Df_a is zero, we need to look at the sign of $Qf_{a+ch}(h)$ for small vectors h . The next order of business is therefore to discuss the values taken by a homogeneous quadratic function.

Definition 4.5.5 (Positive Definite, Negative Definite, Indefinite Matrix). *The symmetric square n -by- n matrix M is called*

- **positive definite** if $Q_M(h) > 0$ for every nonzero $h \in \mathbf{R}^n$,
- **negative definite** if $Q_M(h) < 0$ for every nonzero $h \in \mathbf{R}^n$,
- **indefinite** if $Q_M(h)$ is positive for some h and negative for others.

The identity matrix I is positive definite since $h^t I h = |h|^2$ for all h . The matrix $\begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix}$ is indefinite. The general question of whether a symmetric n -by- n matrix is positive definite leads to an excursion into linear algebra too lengthy for this course. (See exercise 4.5.10 for the result without proof.) However, in the special case of $n = 2$, basic methods give the answer. Recall that the quadratic polynomial $\alpha h^2 + 2\beta h + \delta$ takes positive and negative values if and only if it has distinct real roots, i.e., $\alpha\delta - \beta^2 < 0$.

Proposition 4.5.6 (Two-by-two Definiteness Test). *Consider a matrix*

$$M = \begin{bmatrix} \alpha & \beta \\ \beta & \delta \end{bmatrix} \in M_2(\mathbf{R}). \text{ Then}$$

- (1) M is positive definite if and only if $\alpha > 0$ and $\alpha\delta - \beta^2 > 0$.
- (2) M is negative definite if and only if $\alpha < 0$ and $\alpha\delta - \beta^2 > 0$.
- (3) M is indefinite if and only if $\alpha\delta - \beta^2 < 0$.

Proof. Since $Q_M(t(h, k)) = t^2 Q_M(h, k)$ for all real t , scaling the input vector (h, k) by nonzero real numbers doesn't affect the sign of the output. The second entry k can therefore be scaled to 0 or 1, and if $k = 0$ then the first entry h can be scaled to 1. Therefore, to show (1), reason that

$$\begin{aligned} M \text{ is positive definite} &\iff Q_M(1, 0) > 0 \text{ and } Q_M(h, 1) > 0 \text{ for all } h \in \mathbf{R} \\ &\iff \alpha > 0 \text{ and } \alpha h^2 + 2\beta h + \delta > 0 \text{ for all } h \in \mathbf{R} \\ &\iff \alpha > 0 \text{ and } \alpha\delta - \beta^2 > 0. \end{aligned}$$

(2) is similar. As for (3),

$$\begin{aligned} M \text{ is indefinite} &\iff \alpha h^2 + 2\beta h + \delta \text{ takes positive and negative values} \\ &\iff \alpha\delta - \beta^2 < 0. \end{aligned}$$

□

The proposition provides no information if $\alpha\delta - \beta^2 = 0$. Geometrically, the proposition gives conditions on M to determine that the graph of Q_M is a bowl, an inverted bowl, or a saddle. The condition $\alpha\delta - \beta^2 = 0$ indicates a degenerate graph: a halfpipe, an inverted halfpipe, or a plane.

The positive definite, negative definite, or indefinite character of a matrix is preserved if the matrix entries vary by small enough amounts. Again we restrict our attention to the 2-by-2 case. Here the result is plausible geometrically, since it says that if the matrix $M(a, b)$ defines a function whose graph is (for example) a bowl, then matrices close to $M(a, b)$ should define functions with similar graphs, which thus should still be bowl-shaped. The same persistence holds for a saddle, but a halfpipe can deform immediately into either a bowl or a saddle, and so can a plane.

Proposition 4.5.7 (Persistence of Definiteness). *Let A be a subset of \mathbf{R}^2 , and let the matrix-valued mapping*

$$M : A \longrightarrow M_2(\mathbf{R}), \quad M(x, y) = \begin{bmatrix} \alpha(x, y) & \beta(x, y) \\ \beta(x, y) & \delta(x, y) \end{bmatrix}$$

be continuous. Let (a, b) be an interior point of A . Suppose that the matrix $M(a, b)$ is positive definite. Then for all (x, y) in some ε -ball about (a, b) , the matrix $M(x, y)$ is also positive definite. Similar statements hold for negative definite and indefinite matrices.

Proof. By the Persistence of Inequality principle (Proposition 2.3.9), the criteria $\alpha > 0$ and $\alpha\delta - \beta^2 > 0$ remain valid if x and y vary by a small enough amount. The other statements follow similarly. □

When a function f has continuous second order partial derivatives, the entries of the Hessian $H_f(a)$ vary continuously with a . The upshot of the last proposition is therefore that instead of studying the nebulous notion of “ Qf_{a+ch} for some c ” it suffices to study the explicit function Qf_a .

Proposition 4.5.8 (Two-variable Max/min Test). *Let $f : A \rightarrow \mathbf{R}$ (where $A \subset \mathbf{R}^2$) be C^2 on its interior points. Let (a, b) be an interior point of A , and suppose that $f'(a, b) = (0, 0)$. Let $H_f(a, b) = \begin{bmatrix} \alpha & \beta \\ \beta & \delta \end{bmatrix}$. Then*

- (1) *If $\alpha > 0$ and $\alpha\delta - \beta^2 > 0$ then $f(a, b)$ is a local minimum.*
- (2) *If $\alpha < 0$ and $\alpha\delta - \beta^2 > 0$ then $f(a, b)$ is a local maximum.*
- (3) *If $\alpha\delta - \beta^2 < 0$ then $f(a, b)$ is a saddle point.*

Proof. This follows from Theorem 4.5.4, Proposition 4.5.6 and Proposition 4.5.7. \square

Again, the test gives no information if $\alpha\delta - \beta^2 = 0$.

Returning once again to the example $f(x, y) = \sin^2 x + x^2 y + y^2$ with its critical point $(0, 0)$ and Hessian $H_f(0, 0) = \begin{bmatrix} 2 & 0 \\ 0 & 2 \end{bmatrix}$, the max/min test shows that f has a local minimum at $(0, 0)$.

Another example is to find the extrema of the function

$$f(x, y) = xy(x + y - 3)$$

on the triangle

$$T = \{(x, y) \in \mathbf{R}^2 : x \geq 0, y \geq 0, x + y \leq 3\}.$$

To solve this, first note that T is compact. Therefore f is guaranteed to take a maximum and a minimum value on T . These are assumed either at interior points of T or along the edge. Examining the signs of x , y and $x + y - 3$ shows that f is zero at all points on the edge of T and negative on the interior of T . Thus f assumes its maximum value—zero—along the boundary of T and must assume its minimum somewhere inside. (See figure 4.9.) To find the extrema of f inside T , find the critical points and try the max/min test. The partial derivatives of f (temporarily viewed as a function on all of \mathbf{R}^2) are

$$f_x(x, y) = y(2x + y - 3) \quad f_y(x, y) = x(x + 2y - 3),$$

and these are both zero if $(x, y) \in \{(0, 0), (0, 3), (3, 0), (1, 1)\}$. (The first three critical points are easy to find; the last is found by assuming that neither x nor y is zero, giving the simultaneous equations $2x + y = 3$, $x + 2y = 3$.) The only critical point interior to T is $(1, 1)$, and therefore $f(1, 1) = -1$ must be the minimum value of f . A quick calculation shows that $H_f(1, 1) = \begin{bmatrix} 2 & 1 \\ 1 & 2 \end{bmatrix}$, and the max/min test confirms the minimum at $(1, 1)$.

Another example is to find the extreme values of the function

$$f : \mathbf{R}^2 \rightarrow \mathbf{R}, \quad f(x, y) = \frac{1}{2}x^2 + xy - 2x - \frac{1}{2}y^2.$$

Since \mathbf{R}^2 is not compact, there is no guarantee that f has any extrema. In fact, for large x , $f(x, 0)$ gets arbitrarily large, and for large y , $f(0, y)$ gets arbitrarily small (where small means negative, not epsilonic). So f has no




Figure 4.9. Zero on the boundary, negative on the interior

global extrema. Nonetheless there may be local ones. Every point of \mathbf{R}^2 is interior, so it suffices to examine the critical points of f . The partial derivatives are

$$f_x(x, y) = x + y - 2 \quad f_y(x, y) = x - y,$$

and the only point where both of them vanish is $(x, y) = (1, 1)$. The Hessian is $H_f(1, 1) = \begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix}$, so the critical point $(1, 1)$ is a saddle point. The function f has no extrema, local or global.

Exercises

4.5.1. Compute the best quadratic approximation to $f(x, y) = e^x \cos y$ at the point $(0, 0)$, $f(h, k) \approx f(0, 0) + Df_{(0,0)}(h, k) + \frac{1}{2}Qf_{(0,0)}(h, k)$.

4.5.2. Compute the best quadratic approximation to $f(x, y) = e^{x+2y}$ at the point $(0, 0)$.

4.5.3. Give a heuristic explanation, making whatever reasonable assumptions seem to be helpful, of why the n -dimensional conceptual analogue of figure 4.8 should have 3^n “pictures.” How does this relate to figure 4.7?

4.5.4. Find the extreme values taken by $f(x, y) = xy(4x^2 + y^2 - 16)$ on the quarter ellipse

$$E = \{(x, y) \in \mathbf{R}^2 : x \geq 0, y \geq 0, 4x^2 + y^2 \leq 16\}.$$

4.5.5. Find the local extrema of the function $f(x, y) = x^2 + xy - 4x + \frac{3}{2}y^2 - 7y$ on \mathbf{R}^2 .

4.5.6. Determine the nature of $f(x, y) = \frac{1}{3}x^3 + \frac{1}{3}y^3 + (x - \frac{3}{2})^2 - (y + 4)^2$ at each of its critical points. Are there global extrema?

4.5.7. Find the critical points. Are they maxima, minima, or saddle points? (The max/min test will not always help.)

$$f(x, y) = x^2y + xy^2 \quad g(x, y) = e^{x+y} \quad h(x, y) = x^5y + xy^5 + xy.$$

4.5.8. Discuss local and global extrema of $f(x, y) = \frac{1}{x-1} - \frac{1}{y-1}$ on the open ball $B((0, 0); 1)$ in \mathbf{R}^2 .

4.5.9. The graph of the function $m(x, y) = 6xy^2 - 2x^3 - 3y^4$ is called a **monkey saddle**. Find the three critical points of m and classify each as a maximum, minimum or saddle. (The max/min test will work on two. Study $m(x, 0)$ and $m(0, y)$ to classify the third.) Explain the name “monkey saddle”—computer graphing software may help.

4.5.10. Linear algebra readily addresses the question of whether an n -by- n matrix is positive definite, negative definite, or indefinite.

Definition 4.5.9 (Characteristic Polynomial). Let M be an n -by- n matrix. Its **characteristic polynomial** is

$$p_M(\lambda) = \det(M - \lambda I).$$

The characteristic polynomial of M is a polynomial of degree n in the scalar variable λ .

While the roots of a polynomial with real coefficients are in general complex, the roots of the characteristic polynomial of a symmetric matrix in $M_n(\mathbf{R})$ are guaranteed to be real. The characterization we want is

Theorem 4.5.10 (Description of Definite/Indefinite Matrices). Let M be a symmetric matrix in $M_n(\mathbf{R})$. Then

- (1) M is positive definite if and only if all the roots of $p_M(\lambda)$ are positive.
- (2) M is negative definite if and only if all the roots of $p_M(\lambda)$ are negative.
- (3) M is indefinite if and only if $p_M(\lambda)$ has positive roots and negative roots.

With this result one can extend the methods in the section to functions of more than two variables.

- (a) Let M be the symmetric matrix $\begin{bmatrix} \alpha & \beta \\ \beta & \delta \end{bmatrix} \in M_2(\mathbf{R})$. Show that

$$p_M(\lambda) = \lambda^2 - (\alpha + \delta)\lambda + (\alpha\delta - \beta^2).$$

- (b) Show that Theorem 4.5.10 is equivalent to Proposition 4.5.6 when $n = 2$.

- (c) Classify the 3-by-3 matrices

$$\begin{bmatrix} 1 & -1 & 0 \\ -1 & 2 & 0 \\ 0 & 0 & 1 \end{bmatrix} \quad \begin{bmatrix} 0 & 1 & 0 \\ 1 & 0 & 1 \\ 0 & 1 & 0 \end{bmatrix}.$$

A generalization of Proposition 4.5.7 also holds, since the roots of a polynomial vary continuously with the polynomial’s coefficients. This leads to

Proposition 4.5.11 (General Max/min Test). *Let $f : A \rightarrow \mathbf{R}$ (where $A \subset \mathbf{R}^n$) be \mathcal{C}^2 on its interior points. Let a be an interior point of A , and suppose that $f'(a) = \mathbf{0}_n$. Let $H_f(a)$, the Hessian matrix of f at a , have characteristic polynomial $p(\lambda)$.*

- (1) *If all roots of $p(\lambda)$ are positive then $f(a)$ is a local minimum.*
- (2) *If all roots of $p(\lambda)$ are negative then $f(a)$ is a local maximum.*
- (3) *If $p(\lambda)$ has positive and negative roots then $f(a)$ is a saddle point.*

4.6 Directional Derivatives and the Gradient

Let f be a scalar-valued function, $f : A \rightarrow \mathbf{R}$ where $A \subset \mathbf{R}^n$, and assume that f is differentiable at a point a of A . While the derivative Df_a is a rather abstract object—the linear mapping that gives the best approximation to $f(a+h) - f(a)$ for small h —the partial derivatives $D_j f(a)$ are easy to understand. The j th partial derivative of f at a ,

$$D_j f(a) = \lim_{t \rightarrow 0} \frac{f(a + te_j) - f(a)}{t},$$

measures the rate of change of f at a as its input varies in the j th direction. Visually, $D_j f(a)$ gives the slope of the j th cross section through a of the graph of f .

Analogous formulas measure the rate of change of f at a as its input varies in a direction that doesn't necessarily parallel a coordinate axis. A direction in \mathbf{R}^n is specified by a *unit* vector d , i.e., a vector d such that $|d| = 1$. As the input to f moves distance t in the d direction, f changes by $f(a+td) - f(a)$. Thus the following definition is natural.

Definition 4.6.1 (Directional Derivative). *Let $f : A \rightarrow \mathbf{R}$ (where $A \subset \mathbf{R}^n$) be a function, let a be an interior point of A , and let $d \in \mathbf{R}^n$ be a unit vector. The **directional derivative of f at a in the d direction** is*

$$D_d f(a) = \lim_{t \rightarrow 0} \frac{f(a + td) - f(a)}{t},$$

if this limit exists.

The directional derivatives of f in the standard basis vector directions are simply the partial derivatives.

When $n = 2$ and f is differentiable at $(a, b) \in \mathbf{R}^2$, its graph has a well-fitting tangent plane through $(a, b, f(a, b))$. The plane is determined by the two slopes $D_1 f(a, b)$ and $D_2 f(a, b)$, and it geometrically determines the rate of increase of f in all other directions. (See figure 4.10.) The geometry suggests that if $f : A \rightarrow \mathbf{R}$ (where $A \subset \mathbf{R}^n$) is differentiable at a then all directional derivatives are expressible in terms of the partial derivatives. This is true and easy to show. A special case of the differentiability property (4.3) is

$$\lim_{t \rightarrow 0} \frac{f(a + td) - f(a) - Df_a(td)}{t} = 0,$$

i.e. (show the steps), $D_d f(a) - f'(a) \cdot d = 0$, or

$$D_d f(a) = \sum_{j=1}^n D_j f(a) d_j$$

as desired.




Figure 4.10. General directional slope determined by axis-directional slopes

The derivative matrix $f'(a)$ of a scalar-valued function f at a is often called the **gradient** of f at a and written $\nabla f(a)$. That is,

$$\nabla f(a) = f'(a) = (D_1 f(a), \dots, D_n f(a)).$$

The previous calculation and this definition lead to

Theorem 4.6.2 (Directional Derivative and Gradient). *Let the function $f : A \rightarrow \mathbf{R}$ (where $A \subset \mathbf{R}^n$) be differentiable at a , and let $d \in \mathbf{R}^n$ be a unit vector. Then the directional derivative of f at a in the d direction exists, and it is equal to*

$$\begin{aligned} D_d f(a) &= \sum_{j=1}^n D_j f(a) d_j \\ &= \langle \nabla f(a), d \rangle \\ &= |\nabla f(a)| \cos \theta_{\nabla f(a), d}. \end{aligned}$$

Therefore:

- The rate of increase of f at a in the d direction varies with d , from $-|\nabla f(a)|$ when d points in the direction opposite to $\nabla f(a)$, to $|\nabla f(a)|$ when d points in the same direction as $\nabla f(a)$.
- In particular, the vector $\nabla f(a)$ points in the direction of greatest increase of f at a , and its modulus $|\nabla f(a)|$ is precisely this greatest rate.
- Also, the directions orthogonal to $\nabla f(a)$ are the directions in which f neither increases or decreases at a .

This theorem gives necessary conditions that arise in consequence of the derivative of f existing at a point a . As in section 4.3, the converse statement, that these conditions are sufficient to make the derivative of f exist at a , is false. Each directional derivative $D_d f(a)$ can exist without the derivative Df_a existing (exercise 4.6.9). Furthermore, each directional derivative can exist at a and satisfy the formula $D_d f(a) = \langle \nabla f(a), d \rangle$ in the theorem, but still without the derivative Df_a existing (exercise 4.6.10). The existence of the multivariable derivative Df_a is a stronger condition than any amount of one-variable cross-sectional derivative data at a .

For an example of using the theorem, if you are skiing on the quadratic mountain $f(x, y) = 9 - x^2 - 2y^2$ at the point $(a, f(a)) = (1, 1, 6)$, then your gradient meter shows

$$\nabla f(1, 1) = (D_1 f(1, 1), D_2 f(1, 1)) = (-2x, -4y)|_{(x,y)=(1,1)} = (-2, -4).$$

Therefore the direction of steepest descent *down* the hillside is the $(2, 4)$ -direction (this could be divided by its modulus $\sqrt{20}$ to make it a unit vector), and the slope of steepest descent is the absolute value $|\nabla f(1, 1)| = \sqrt{20}$. On the other hand, cross-country skiing in the $(2, -1)$ -direction, which is orthogonal to $\nabla f(1, 1)$, neither gains nor loses elevation immediately. (See figure 4.11.) The cross-country skiing trail that neither climbs nor descends has a mathematical name.

Definition 4.6.3 (Level Set). Let $f : A \rightarrow \mathbf{R}$ (where $A \subset \mathbf{R}^n$) be a function. A **level set** of f is the set of points in A that map under f to some fixed value b in \mathbf{R} ,

$$L = \{x \in A : f(x) = b\}.$$

For example, on the mountain

$$f : \mathbf{R}^2 \rightarrow \mathbf{R}, \quad f(x, y) = 9 - x^2 - 2y^2,$$

the level set for $b = 5$ is an ellipse in the plane,

$$L = \{(x, y) \in \mathbf{R}^2 : x^2 + 2y^2 = 4\}.$$

And similarly the level set is an ellipse for any real number b up to 9. In general, plotting the level sets of a function f of two variables gives a topographical




Figure 4.11. Gradient and its orthogonal vector for the parabolic mountain

map description of f . The geometry is different for a function of one variable: each level set is a subset of the line. For example, consider a restriction of the sine function,

$$f : (0, \pi) \longrightarrow \mathbf{R}, \quad f(x) = \sin(x).$$

The level set taken by f to $1/2$ consists of two points,

$$L = \{\pi/6, 5\pi/6\}.$$

As for a function of three variables, each level set is a subset of space. For example, if a , b , and c are positive numbers, and the function is

$$f : \mathbf{R}^3 \longrightarrow \mathbf{R}, \quad f(x, y, z) = (x/a)^2 + (y/b)^2 + (z/c)^2,$$

then its level sets are ellipsoids. Specifically, for every positive r , the level set of points taken by f to r is the ellipsoid of x -radius $a\sqrt{r}$, y -radius $b\sqrt{r}$, and z -radius $c\sqrt{r}$,

$$L = \left\{ (x, y, z) \in \mathbf{R}^3 : \left(\frac{x}{a\sqrt{r}} \right)^2 + \left(\frac{y}{b\sqrt{r}} \right)^2 + \left(\frac{z}{c\sqrt{r}} \right)^2 = 1 \right\}.$$

The third bullet in Theorem 4.6.2 says that *the gradient is normal to the level set*. This may seem surprising, since the gradient is a version of the derivative, and we think of the derivative as describing a tangent object to a graph. The reason that the derivative has become a normal object is that a level set is different from a graph. A level set of f is a subset of the domain of f , whereas the graph of f , which simultaneously shows the domain and the range of f , is a subset of a space that is one dimension larger. For instance, if we think of f as measuring elevation, then the graph of f is terrain in three-dimensional space, while a level set of f is set of points in the plane that lie beneath the terrain at some constant altitude; the level set is typically

a curve. Figure 4.11 illustrates the difference in the case of the mountain function. Note that in the left part of the figure, the gradient is orthogonal to the ellipse on which it starts. Similarly, figure 4.12 illustrates the difference in the case of the restricted sine function from the previous paragraph. In the figure, the x -axis shows the two-point level set from the previous paragraph, and the gradient of f at each of the two points. The fact that one gradient points right indicates that to climb the graph of f over that point, one should move to the right, and the slope to be encountered on the graph will be the length of the gradient on the axis. Similarly, the other gradient points left because to climb the graph over the other point, one should move to the left. Here each gradient is trivially orthogonal to the level set, because the level set consists of isolated points. For the three-variable function from the previous paragraph, we still can see the level sets—they are concentric ellipsoids—but not the graph, which would require four dimensions. Instead, we can conceive of the function as measuring temperature in space, and of the gradient as pointing in the direction to move for greatest rate of temperature-increase, with the length of the gradient being that rate. Figure 4.13 shows a level set for the temperature function, and several gradients, visibly orthogonal to the level set.




Figure 4.12. Level set and gradients for the sine function

Although Theorem 4.6.2 has already stated that the gradient is orthogonal to the level set, we now amplify the argument. Let $f : A \rightarrow \mathbf{R}$ (where $A \subset \mathbf{R}^n$) be given, and assume that it is differentiable. Let a be a point of A , and let $b = f(a)$. Consider the level set of f containing a ,

$$L = \{x \in A : f(x) = b\} \subset \mathbf{R}^n,$$

and consider any smooth curve from some interval into the level set, passing through a ,

$$\gamma : (-\varepsilon, \varepsilon) \rightarrow L, \quad \gamma(0) = a.$$

The composite function

$$f \circ \gamma : (-\varepsilon, \varepsilon) \rightarrow \mathbf{R}$$




Figure 4.13. Level set and gradients for the temperature function

is the constant function b , so that its derivative at 0 is 0. By the chain rule this relation is

$$\nabla f(a) \cdot \gamma'(0) = 0.$$

Every tangent vector to L at a takes the form $\gamma'(0)$ for some γ of the sort that we are considering. Therefore, $\nabla f(a)$ is orthogonal to every tangent vector to L at a , i.e., $\nabla f(a)$ is normal to L at a .

Before continuing to work with the gradient, we pause to remark that level sets and graphs are related. For one thing:

The graph of any function is also the level set of a different function.

To see this, let $n > 1$, let A_0 be a subset of \mathbf{R}^{n-1} , and let $f : A_0 \rightarrow \mathbf{R}$ be any function. Given this information, let $A = A_0 \times \mathbf{R}$ and define a second function $g : A \rightarrow \mathbf{R}$,

$$g(x_1, \dots, x_{n-1}, x_n) = f(x_1, \dots, x_{n-1}) - x_n.$$

Then the graph of f is a level of g , specifically the set of inputs that g takes to 0,

$$\begin{aligned} \text{graph}(f) &= \{x \in A_0 \times \mathbf{R} : x_n = f(x_1, \dots, x_{n-1})\} \\ &= \{x \in A : g(x) = 0\}. \end{aligned}$$

For example, the graph of the mountain function $f(x, y) = 9 - x^2 - 2y^2$ is also a level set of the function $g(x, y, z) = 9 - x^2 - 2y^2 - z$. But in contrast to this quick method defining g explicitly in terms of f to show that every graph is a level set, the converse question is much more subtle:

To what extent is some given level set also a graph?

For example, the level sets of the mountain function f are ellipses (as shown in figure 4.11), but an ellipse is not the graph of y as a function of x or vice versa. The converse question will be addressed by the Implicit Function Theorem in the next chapter.

Returning to the gradient, the geometrical fact that it is normal to the level set makes it easy to find the tangent plane to a two-dimensional surface in \mathbf{R}^3 . For example, consider the surface

$$H = \{(x, y, z) \in \mathbf{R}^3 : x^2 + y^2 - z^2 = 1\}.$$

(This is a hyperboloid of one sheet.) The point $(2\sqrt{2}, 3, 4)$ belongs to H . Note that H as a level set of the function $f(x, y, z) = x^2 + y^2 - z^2$, and compute the gradient

$$\nabla f(2\sqrt{2}, 3, 4) = (4\sqrt{2}, 6, -8).$$

Since this is the normal vector to H at $(2\sqrt{2}, 3, 4)$, the methods of section 3.10 show that the equation of the tangent plane to H at $(2\sqrt{2}, 3, 4)$ is

$$4\sqrt{2}(x - 2\sqrt{2}) + 6(y - 3) - 8(z - 4) = 0.$$

If a function $f : \mathbf{R}^n \rightarrow \mathbf{R}$ has a continuous gradient, then from any starting point $a \in \mathbf{R}^n$ where the gradient $\nabla f(a)$ is nonzero, there is a path of steepest ascent of f (called an **integral curve of ∇f**) starting at a . If $n = 2$ and the graph of f is seen as a surface in 3-space, then the integral curve from the point $(a, b) \in \mathbf{R}^2$ is the shadow of the path followed by a particle climbing the graph, starting at $(a, b, f(a, b))$. If $n = 2$ or $n = 3$ and f is viewed as temperature, then the integral curve is the path followed by a heat-seeking bug.

To find the integral curve, we set up an equation that describes it. The idea is to treat the gradient vector as a divining rod and follow it starting at a . This produces a path in \mathbf{R}^n that describes time-dependent motion, always in the direction of the gradient, and always with speed equal to the modulus of the gradient. Computing the path amounts to finding an interval $I \subset \mathbf{R}$ containing 0 and a mapping

$$\gamma : I \rightarrow \mathbf{R}^n$$

that satisfies the differential equation with initial conditions

$$\gamma'(t) = \nabla f(\gamma(t)), \quad \gamma(0) = a. \quad (4.6)$$

Whether (and how) one can solve this for γ depends on the data f and a .

In the case of the mountain function $f(x, y) = 9 - x^2 - 2y^2$, with gradient $\nabla f(x, y) = (-2x, -4y)$, the path γ has two components γ_1 and γ_2 , and the differential equation and initial conditions (4.6) become

$$(\gamma_1'(t), \gamma_2'(t)) = (-2\gamma_1(t), -4\gamma_2(t)), \quad (\gamma_1(0), \gamma_2(0)) = (a, b),$$

to which the unique solution is

$$(\gamma_1(t), \gamma_2(t)) = (ae^{-2t}, be^{-4t}).$$

Let $x = \gamma_1(t)$ and $y = \gamma_2(t)$. Then the previous display shows that

$$a^2y = bx^2,$$

and so the integral curve lies on a parabola. (The parabola is degenerate if the starting point (a, b) lies on either axis.) Every parabola that forms an integral curve for the mountain function meets orthogonally with every ellipse that forms a level set. (See figure 4.14.) In general, solving the vector differential equation (4.6) to find the integral curves γ of a function f can be difficult.




Figure 4.14. Level sets and integral curves for the parabolic mountain

For another example, suppose the temperature in space is given by $T(x, y, z) = 1/(x^2 + y^2 + z^2)$. (This blows up at the origin, so don't work there.) The level sets of this function are spheres and the integral curves are rays going toward the origin. The level set passing through the point (a, b, c) in space is again orthogonal to the integral curve through the same point.

Exercises

4.6.1. Let $f(x, y, z) = xy^2 + yz$. Find $D_{(\frac{2}{3}, -\frac{1}{3}, \frac{2}{3})}f(1, 1, 2)$.

4.6.2. Let $g(x, y, z) = xyz$, let d be the unit vector in the direction from $(1, 2, 3)$ to $(3, 1, 5)$. Find $D_dg(1, 2, 3)$.

4.6.3. Let f be differentiable at a point a , and let $d = -e_1$, a unit vector. Are the directional derivative $D_d f(a)$ and the partial derivative $D_1 f(a)$ equal? Explain.

4.6.4. Formulate and prove a version of Rolle's theorem for functions of n variables.

4.6.5. Show that if $f : \mathbf{R}^n \rightarrow \mathbf{R}$ and $g : \mathbf{R}^n \rightarrow \mathbf{R}$ are differentiable then so is their product $fg : \mathbf{R}^n \rightarrow \mathbf{R}$ and $\nabla(fg) = f\nabla g + g\nabla f$.

4.6.6. Find the tangent plane to the surface $\{(x, y, z) : x^2 + 2y^2 + 3zx - 10 = 0\}$ in \mathbf{R}^3 at the point $(1, 2, \frac{1}{3})$.

4.6.7. The pheromone concentration in space is given by $P(x, y, z) = e^{-x} + e^{-2y} + e^{-3z}$. In what direction does a moth at $(1, 1, 1)$ fly? At what rate is the pheromone concentration increasing in that direction?

4.6.8. Sketch some level sets and integral curves for the function $f(x, y) = xy$.

4.6.9. Define $f : \mathbf{R}^2 \rightarrow \mathbf{R}$ by

$$f(x, y) = \begin{cases} \frac{x^2 y}{x^2 + y^2} & \text{if } (x, y) \neq 0 \\ 0 & \text{if } (x, y) = 0. \end{cases}$$

- (a) Show that f is continuous at $(0, 0)$.
 (b) Find the partial derivatives $D_1 f(0, 0)$ and $D_2 f(0, 0)$.
 (c) Let d be any unit vector in \mathbf{R}^2 (thus d takes the form $d = (\cos \theta, \sin \theta)$ for some $\theta \in \mathbf{R}$). Show that $D_d f(0, 0)$ exists by finding it.
 (d) Show that in spite of (c), f is not differentiable at $(0, 0)$. (Use your results from parts (b) and (c) to contradict Theorem 4.6.2.) Thus, the existence of every directional derivative at a point is not sufficient for differentiability at the point.

4.6.10. Define $f : \mathbf{R}^2 \rightarrow \mathbf{R}$ by

$$f(x, y) = \begin{cases} 1 & \text{if } y = x^2 \text{ and } (x, y) \neq (0, 0) \\ 0 & \text{otherwise.} \end{cases}$$

- (a) Show that f is discontinuous at $(0, 0)$. It follows that f is not differentiable at $(0, 0)$.
 (b) Let d be any unit vector in \mathbf{R}^2 . Show that $D_d f(0, 0) = 0$. Show that consequently the formula $D_d f(0, 0) = \langle \nabla f(0, 0), d \rangle$ holds for every unit vector d . Thus, the existence of every directional derivative at a point, and the fact that each directional derivative satisfies the formula are still not sufficient for differentiability at the point.

4.6.11. Fix two real numbers a and b satisfying $0 < a < b$. Define a mapping $T = (T_1, T_2, T_3) : \mathbf{R}^2 \rightarrow \mathbf{R}^3$ by

$$T(s, t) = ((b + a \cos s) \cos t, (b + a \cos s) \sin t, a \sin s).$$

- (a) Describe the shape of the set in \mathbf{R}^3 mapped to by T . (The answer will explain the name “ T .”)

(b) Find the points $(s, t) \in \mathbf{R}^2$ such that $\nabla T_1(s, t) = \mathbf{0}_2$. The points map to only four image points p under T . Show that one such p is a maximum of T_1 , another is a minimum, and the remaining two are saddle points.

(c) Find the points $(s, t) \in \mathbf{R}^2$ such that $\nabla T_3(s, t) = \mathbf{0}_2$. To what points q do these (s, t) map under T ? Which such q are maxima of T_3 ? Minima? Saddle points?

4.7 Summary

The multivariable derivative figures in solving various kinds of problems:

- changing variables in a partial differential equation,
- optimization of a scalar-valued function of many variables,
- determining integral curves.

The next chapter will consider a fourth type of problem, optimization with constraints.

Inverse and Implicit Functions

The question of whether a mapping $f : A \rightarrow \mathbf{R}^n$ (where $A \subset \mathbf{R}^n$) is globally invertible is beyond the local techniques of differential calculus. However, a local theorem is finally in reach. The idea sounds plausible: if the derivative of f is invertible at the point a then f itself, being well approximated near a by its derivative, should also be invertible in the small. However, it is by no means a general principle that an approximated object must have the properties of the object approximating it. On the contrary, mathematics often approximates complicated objects by simpler ones. For example, Taylor's Theorem approximates any function that has many derivatives by a polynomial. This does not make the function itself a polynomial as well.

Even in the one-variable case, the Inverse Function Theorem relies on foundational theorems about the real number system, on a property of continuous functions, and on a foundational theorem of differential calculus. We quickly review the ideas. Let $f : A \rightarrow \mathbf{R}$ (where $A \subset \mathbf{R}$) be a function, let a be an interior point of A , and let f be continuously differentiable on some interval about a , meaning that f' exists and is continuous on the interval. Suppose that $f'(a) > 0$. Since f' is continuous about a , the Persistence of Inequality principle (Proposition 2.3.9) says that f' is positive on some closed interval $[a - \delta, a + \delta]$ about a . By an application of the Mean Value Theorem, f is therefore strictly increasing on the interval, and so its restriction to the interval does not take any value twice. By the Intermediate Value Theorem, f takes every value from $f(a - \delta)$ to $f(a + \delta)$ on the interval. Therefore f takes every such value exactly once, making it locally invertible. A slightly subtle point is that the inverse function f^{-1} is continuous at $f(a)$, but then a purely formal calculation with difference quotients will verify that the derivative of f^{-1} exists at $f(a)$ and is $1/f'(a)$. Note how heavily this proof relies on the fact that \mathbf{R} is an ordered field. A proof of the multivariable Inverse Function Theorem must use other methods.

The proof to be given in this chapter has its technical aspects, but the core idea is simple common sense. Let a mapping f be given that takes x -values to y -values, and that in particular takes a to b . Then the local inverse

function must take y -values near b to x -values near a , taking each such y back to the unique x that f took to y in the first place. We need to determine conditions on f that make us believe that a local inverse exists. As explained above, the basic condition is that the derivative of f at a —giving a good approximation of f near a , but easier to understand than f itself—should be invertible. The derivative should be continuous as well, for technical reasons. With these conditions in hand, it is elementary (though a bit painstaking) to show that f is locally injective:

- Given y near b , there is *at most one* x near a that f takes to y .

So the problem reduces to showing that f is locally surjective:

- Given y near b , show that there is *some* x near a that f takes to y .

This problem decomposes to two subproblems. First:

- Given y near b , show that there is some x near a that f takes *closest* to y .

Then:

- Show that f takes this particular x *exactly* to y .

And once the appropriate environment is established, solving each of these is just a matter of applying the main theorems from the previous three chapters.

Not only does the Inverse Function Theorem have a proof that uses so much previous work from this course so nicely, it also has useful consequences. It leads easily to the Implicit Function Theorem, which answers a different question: When does a set of constraining relations among a set of variables make some of the variables dependent on the others? The Implicit Function Theorem in turn justifies the Lagrange multiplier method, a technique for solving optimization problems with constraints. As discussed back in the preface, these problems have no one-variable counterpart, and they can be viewed as the beginning of calculus on curved spaces.

5.1 Preliminaries

The basic elements of topology in \mathbf{R}^n — ε -balls; limit points; closed, bounded, and compact sets—were introduced in section 2.4 to provide the environment for the Extreme Value Theorem. A little more topology is now needed before we proceed to the Inverse Function Theorem. Recall that for any point $a \in \mathbf{R}^n$ and any radius $\varepsilon > 0$, the ε -ball at a is the set

$$B(a, \varepsilon) = \{x \in \mathbf{R}^n : |x - a| < \varepsilon\}.$$

Recall also that a subset of \mathbf{R}^n is called closed if it contains all of its limit points. Not unnaturally, a subset S of \mathbf{R}^n is called **open** if its **complement** $S^c = \mathbf{R}^n - S$ is closed. A set, however, is not a door: it can be neither open or closed, and it can be both open and closed. (Examples?)

Proposition 5.1.1 (ε -balls Are Open). For any $a \in \mathbf{R}^n$ and any $\varepsilon > 0$, the ball $B(a, \varepsilon)$ is open.

Proof. Let x be any point in $B(a, \varepsilon)$, and set $\delta = \varepsilon - |x - a|$, a positive number. The triangle inequality shows that $B(x, \delta) \subset B(a, \varepsilon)$ (exercise 5.1.1), and therefore x is not a limit point of the complement $B(a, \varepsilon)^c$. Consequently all limit points of $B(a, \varepsilon)^c$ are in fact elements of $B(a, \varepsilon)^c$, which is thus closed, making $B(a, \varepsilon)$ itself open. \square

This proof shows that any point $x \in B(a, \varepsilon)$ is an interior point. In fact, an equivalent definition of “open” is that a subset of \mathbf{R}^n is open if each of its points is interior (exercise 5.1.2).

The **closed ε -ball at a** , denoted $\overline{B(a, \varepsilon)}$, consists of the corresponding open ball with its edge added in,

$$\overline{B(a, \varepsilon)} = \{x \in \mathbf{R}^n : |x - a| \leq \varepsilon\}.$$

The **boundary** of the closed ball $\overline{B(a, \varepsilon)}$, denoted $\partial\overline{B(a, \varepsilon)}$, is the points on the edge,

$$\partial\overline{B(a, \varepsilon)} = \{x \in \mathbf{R}^n : |x - a| = \varepsilon\}.$$

(See figure 5.1.) Any closed ball \overline{B} and its boundary $\partial\overline{B}$ are compact sets (exercise 5.1.3).




Figure 5.1. Open ball, closed ball, and boundary

Let $f : A \rightarrow \mathbf{R}^m$ (where $A \subset \mathbf{R}^n$) be continuous, let W be an open subset of \mathbf{R}^m , and let V be the set of all points in A that f maps into W ,

$$V = \{x \in A : f(x) \in W\}.$$

The set V is called the **inverse image of W under f** ; it is often denoted $f^{-1}(W)$, but this is a little misleading since f need not actually have an inverse mapping f^{-1} . For example, if $f : \mathbf{R} \rightarrow \mathbf{R}$ is the squaring function $f(x) = x^2$, then the inverse image of $[4, 9]$ is $[-3, -2] \cup [2, 3]$, and this set is denoted $f^{-1}([4, 9])$ even though f has no inverse. (See figure 5.2, in which f is not the squaring function, but the inverse image $f^{-1}(W)$ also has two components.) The inverse image concept generalizes an idea that we saw in section 4.6: the inverse image of a one-point set under a mapping f is a level set of f , as in Definition 4.6.3.




Figure 5.2. Inverse image with two components

Although the forward image under a continuous function of an open set need not be open (exercise 5.1.4), inverse images behave more nicely. The connection between continuous mappings and open sets is

Theorem 5.1.2 (Inverse Image Characterization of Continuity). *Let $f : A \rightarrow \mathbf{R}^m$ (where A is an open subset of \mathbf{R}^n) be continuous. Let $W \subset \mathbf{R}^m$ be open. Then $f^{-1}(W)$, the inverse image of W under f , is open.*

Proof. Let a be a point of $f^{-1}(W)$. We want to show that it is an interior point. Let $w = f(a)$, a point of W . Since W is open, some ball $B(w, \rho)$ is contained in W . Consider the function

$$g : A \rightarrow \mathbf{R}, \quad g(x) = \rho - |f(x) - w|.$$

This function is continuous and it satisfies $g(a) = \rho > 0$, and so by a slight variant of the Persistence of Inequality principle there exists a ball $B(a, \varepsilon) \subset A$ on which g remains positive. That is,

$$f(x) \in B(w, \rho) \quad \text{for all } x \in B(a, \varepsilon).$$

Since $B(w, \rho) \subset W$, this shows that $B(a, \varepsilon) \subset f^{-1}(W)$, making a an interior point of $f^{-1}(W)$ as desired. \square

The converse to Theorem 5.1.2 is also true and is exercise 5.1.8. We need one last technical result for the proof of the Inverse Function Theorem.

Lemma 5.1.3 (Difference Magnification Lemma). *Let \overline{B} be a closed ball in \mathbf{R}^n and let g be a differentiable mapping from an open set in \mathbf{R}^n containing \overline{B} back to \mathbf{R}^n . Suppose that there is a number c such that $|D_j g_i(x)| \leq c$ for all $i, j \in \{1, \dots, n\}$ and all $x \in \overline{B}$. Then*

$$|g(\tilde{x}) - g(x)| \leq n^2 c |\tilde{x} - x| \quad \text{for all } x, \tilde{x} \in \overline{B}.$$

Proof. Given two points $x, \tilde{x} \in \overline{B}$, make the line segment connecting them the image of a function of one variable,

$$\gamma : [0, 1] \longrightarrow \mathbf{R}^n, \quad \gamma(t) = x + t(\tilde{x} - x).$$

Note that $\gamma(0) = x$, $\gamma(1) = \tilde{x}$, and $\gamma'(t) = \tilde{x} - x$ for all $t \in (0, 1)$. Fix any $i \in \{1, \dots, n\}$ and consider the restriction of g_i to the segment,

$$\varphi : [0, 1] \longrightarrow \mathbf{R}, \quad \varphi(t) = (g_i \circ \gamma)(t).$$

Thus $\varphi(0) = g_i(x)$ and $\varphi(1) = g_i(\tilde{x})$. By the Mean Value Theorem,

$$g_i(\tilde{x}) - g_i(x) = \varphi(1) - \varphi(0) = \varphi'(t) \quad \text{for some } t \in (0, 1),$$

and so since $\varphi = g_i \circ \gamma$ the Chain Rule gives

$$g_i(\tilde{x}) - g_i(x) = (g_i \circ \gamma)'(t) = g_i'(\gamma(t))\gamma'(t) = g_i'(\gamma(t))(\tilde{x} - x).$$

Here $g_i'(\gamma(t))$ is a row vector and $\tilde{x} - x$ is a column vector. Viewing them both as vectors with no reference to rows or columns, we have

$$g_i(\tilde{x}) - g_i(x) = \langle g_i'(\gamma(t)), \tilde{x} - x \rangle,$$

so that by the Cauchy–Schwarz inequality,

$$|g_i(\tilde{x}) - g_i(x)| \leq |g_i'(\gamma(t))| |\tilde{x} - x|.$$

For each j , the j th entry of the vector $g_i'(\gamma(t))$ is $D_j g_i(\gamma(t))$ by Theorem 4.3.2. But we are given that $|D_j g_i(\gamma(t))| \leq c$, so the Size Bounds show that $|g_i'(\gamma(t))| \leq nc$ and therefore

$$|g_i(\tilde{x}) - g_i(x)| \leq nc|\tilde{x} - x|.$$

Again by the Size Bounds, $|g(x) - g(\tilde{x})| \leq \sum_{i=1}^n |g_i(\tilde{x}) - g_i(x)|$, so by the previous display,

$$|g(\tilde{x}) - g(x)| \leq n^2 c |\tilde{x} - x|.$$

This is the desired result. \square

Exercises

5.1.1. Let $x \in B(a; \varepsilon)$ and let $\delta = \varepsilon - |x - a|$. Explain why $\delta > 0$ and why $B(x; \delta) \subset B(a; \varepsilon)$.

5.1.2. Show that a subset of \mathbf{R}^n is open if and only if each of its points is interior.

5.1.3. Prove that any closed ball \overline{B} is indeed a closed set, as is its boundary $\partial \overline{B}$. Show that any closed ball and its boundary are also bounded, hence compact.

5.1.4. Find a continuous function $f : \mathbf{R}^n \rightarrow \mathbf{R}^m$ and an open set $A \subset \mathbf{R}^n$ such that the image $f(A) \subset \mathbf{R}^m$ of A under f is not open. Feel free to choose n and m .

5.1.5. Define $f : \mathbf{R} \rightarrow \mathbf{R}$ by $f(x) = x^3 - 3x$. Compute $f(-1/2)$. Find $f^{-1}((0, 11/8))$, $f^{-1}((0, 2))$, $f^{-1}((-\infty, -11/8) \cup (11/8, \infty))$. Does f^{-1} exist?

5.1.6. Show that for $f : \mathbf{R}^n \rightarrow \mathbf{R}^m$ and $B \subset \mathbf{R}^m$, the inverse image of the complement is the complement of the inverse image,

$$f^{-1}(B^c) = f^{-1}(B)^c.$$

Does the analogous formula hold for forward images?

5.1.7. If $f : \mathbf{R}^n \rightarrow \mathbf{R}^m$ is continuous and $B \subset \mathbf{R}^m$ is closed, show that $f^{-1}(B)$ is closed. What does this say about the level sets of continuous functions?

5.1.8. Prove the converse to Theorem 5.1.2: If $f : A \rightarrow \mathbf{R}^m$ (where $A \subset \mathbf{R}^n$ is open) is such that for any open $W \subset \mathbf{R}^m$ also $f^{-1}(W) \subset A$ is open, then f is continuous.

5.1.9. Let a and b be real numbers with $a < b$. Let $n > 1$, and suppose that the mapping $f : [a, b] \rightarrow \mathbf{R}^n$ is continuous and that f is differentiable on the open interval (a, b) . It is tempting to generalize the Mean Value Theorem (Theorem 1.2.3) to the assertion

$$"f(b) - f(a) = f'(c)(b - a) \quad \text{for some } c \in (a, b)."$$
 (5.1)

This assertion is grammatically meaningful, since it posits an equality between two n -vectors. However, the assertion is false.

(a) Let $f : [0, 2\pi] \rightarrow \mathbf{R}^2$ be $f(t) = (\cos t, \sin t)$. Show that (5.1) fails for this f . Describe the situation geometrically.

(b) Let $f : [0, 2\pi] \rightarrow \mathbf{R}^3$ be $f(t) = (\cos t, \sin t, t)$. Show that (5.1) fails for this f . Describe the situation geometrically.

(c) Here is an attempt to prove (5.1): Let $f = (f_1, \dots, f_n)$. Since each f_i is scalar-valued, we have for $i = 1, \dots, n$ by the Mean Value Theorem,

$$f_i(b) - f_i(a) = f'_i(c)(b - a) \quad \text{for some } c \in (a, b).$$

Assembling the scalar results gives the desired vector result.

What is the error here?

5.2 The Inverse Function Theorem

Theorem 5.2.1 (Inverse Function Theorem). Let $f : A \rightarrow \mathbf{R}^n$ (where $A \subset \mathbf{R}^n$) be a mapping, let a be an interior point of A , and let f be continuously differentiable on some ε -ball about a . (This means first that the derivative mapping Df_x exists for each x in the ball, and second that the entries of

the derivative matrix $f'(x)$, i.e., the partial derivatives $D_j f_i(x)$, are continuous functions of x on the ball.) Suppose that $\det f'(a) \neq 0$. Then there is an open set $V \subset A$ containing a and an open set $W \subset \mathbf{R}^n$ containing $f(a)$ such that $f : V \rightarrow W$ has a continuously differentiable inverse $f^{-1} : W \rightarrow V$. For each $y = f(x) \in W$, the derivative of the inverse is the inverse of the derivative,

$$D(f^{-1})_y = (Df_x)^{-1}.$$

Before the proof, it is worth remarking that the the formula for the derivative of the local inverse, and the fact that the derivative of the local inverse is continuous, are easy to establish once everything else is in place. If the local inverse f^{-1} of f is known to exist and to be differentiable, then for any $x \in V$ the fact that the identity mapping is its own derivative combines with the chain rule to say that

$$\text{id}_n = D(\text{id}_n)_x = D(f^{-1} \circ f)_n = D(f^{-1})_y \circ Df_x \quad \text{where } y = f(x),$$

and similarly $\text{id}_n = Df_x \circ D(f^{-1})_y$, where this time id_n is the identity mapping on y -space. The last formula in the theorem follows. In terms of matrices, the formula is

$$(f^{-1})'(y) = f'(x)^{-1} \quad \text{where } y = f(x).$$

This formula combines with Corollary 3.7.3 (the entries of the inverse matrix are continuous functions of the entries of the matrix) to show that since the mapping is continuously differentiable and the local inverse is differentiable, the local inverse is continuously differentiable. Thus we need to show only that the local inverse exists and is differentiable.

Proof. The proof begins with a simplification. Let $T = Df_a$, a linear mapping from \mathbf{R}^n to \mathbf{R}^n that is invertible because its matrix $f'(a)$ has nonzero determinant. Let

$$\tilde{f} = T^{-1} \circ f.$$

By the chain rule, the derivative of \tilde{f} at a is

$$D\tilde{f}_a = D(T^{-1} \circ f)_a = D(T^{-1})_{f(a)} \circ Df_a = T^{-1} \circ T = \text{id}_n.$$


Also, suppose we have a local inverse \tilde{g} of \tilde{f} , so that

$$\tilde{g} \circ \tilde{f} = \text{id}_n \text{ near } a$$

and

$$\tilde{f} \circ \tilde{g} = \text{id}_n \text{ near } \tilde{f}(a).$$

The situation is shown in the following diagram, in which V is an open set containing a , W is an open set containing $f(a)$, and \tilde{W} is an open set containing $T^{-1}(f(a)) = \tilde{f}(a)$.



The diagram shows that the way to invert f locally, going from W back to V , is to proceed through \widetilde{W} : $g = \tilde{g} \circ T^{-1}$. Indeed, since $f = T \circ \tilde{f}$,

$$g \circ f = (\tilde{g} \circ T^{-1}) \circ (T \circ \tilde{f}) = \text{id}_n \text{ near } a,$$

and, since $T^{-1}(f(a)) = \tilde{f}(a)$,

$$f \circ g = (T \circ \tilde{f}) \circ (\tilde{g} \circ T^{-1}) = \text{id}_n \text{ near } f(a).$$

That is, to invert f , it suffices to invert \tilde{f} . And if \tilde{g} is differentiable then so is $g = \tilde{g} \circ T^{-1}$. The upshot is that we may prove the theorem for \tilde{f} rather than f . Equivalently, we may assume with no loss of generality that $Df_a = \text{id}_n$ and therefore that $f'(a) = I_n$. This normalization will let us carry out a clean, explicit computation in the following paragraph.

Next we find a closed ball \overline{B} around a where the behavior of f is somewhat controlled by the fact that $f'(a) = I_n$. Recall that the (i, j) th entry of I_n is δ_{ij} and that $\det(I_n) = 1$. As x varies continuously near a , the (i, j) th entry $D_j f_i(x)$ of $f'(x)$ varies continuously near δ_{ij} , and so the scalar $\det f'(x)$ varies continuously near 1. Therefore, there is a closed ball \overline{B} small enough that

$$|D_j f_i(x) - \delta_{ij}| < \frac{1}{2n^2} \quad \text{for all } i, j \in \{1, \dots, n\} \text{ and } x \in \overline{B} \quad (5.2)$$

and

$$\det f'(x) \neq 0 \quad \text{for all } x \in \overline{B}. \quad (5.3)$$

Let $g = f - \text{id}_n$, a differentiable mapping near a , whose Jacobian matrix at x , $g'(x) = f'(x) - I_n$, has (i, j) th entry $D_j g_i(x) = D_j f_i(x) - \delta_{ij}$. Equation (5.2) and Lemma 5.1.3 (with $c = 1/(2n^2)$) show that for any two points x and \tilde{x} in \overline{B} ,

$$|g(\tilde{x}) - g(x)| \leq \frac{1}{2} |\tilde{x} - x|,$$

and therefore, since $f = \text{id}_n + g$,

$$\begin{aligned} |f(\tilde{x}) - f(x)| &= |(\tilde{x} - x) + (g(\tilde{x}) - g(x))| \\ &\geq |\tilde{x} - x| - |g(\tilde{x}) - g(x)| \\ &\geq |\tilde{x} - x| - \frac{1}{2} |\tilde{x} - x| \quad (\text{by the previous display}) \\ &= \frac{1}{2} |\tilde{x} - x|. \end{aligned}$$

This implies that f is injective on \overline{B} , i.e., any two distinct points of \overline{B} are taken by f to distinct points of \mathbf{R}^n . For future reference, we note that the result of the previous calculation rearranges as

$$|\tilde{x} - x| \leq 2|f(\tilde{x}) - f(x)| \quad \text{for all } x, \tilde{x} \in \overline{B}. \quad (5.4)$$

The boundary $\partial\overline{B}$ of \overline{B} is compact, and so is the image set $f(\partial\overline{B})$ since f is continuous. Also, $f(a) \notin f(\partial\overline{B})$ since f is injective on \overline{B} . And $f(a)$ is not a limit point of $f(\partial\overline{B})$ since $f(\partial\overline{B})$, being compact, is closed. This means that some open ball $B(f(a), 2\varepsilon)$ contains no point from $f(\partial\overline{B})$. (See figure 5.3.)




Figure 5.3. Ball about $f(a)$ away from $f(\partial\overline{B})$

Let $W = B(f(a), \varepsilon)$, the open ball with radius less than half the distance from $f(a)$ to $f(\partial\overline{B})$. Thus

$$|y - f(a)| < |y - f(x)| \quad \text{for all } y \in W \text{ and } x \in \partial\overline{B}. \quad (5.5)$$

That is, every point y of W is closer to $f(a)$ than it is to any point of $f(\partial\overline{B})$. (See figure 5.4.)




Figure 5.4. Ball closer to $f(a)$ than to $f(\partial\overline{B})$

The goal now is to exhibit a mapping on W that inverts f near a . In other words, the goal is to show that for each $y \in W$, there exists a unique x interior to \overline{B} such that $f(x) = y$. So fix an arbitrary $y \in W$. Define a function $\Delta : \overline{B} \rightarrow \mathbf{R}$ that measures for each x the square of the distance from $f(x)$ to y ,

$$\Delta(x) = |y - f(x)|^2 = \sum_{i=1}^n (y_i - f_i(x))^2.$$

The idea is to show that for one and only one x near a , $\Delta(x) = 0$. Since modulus is always nonnegative, the x we seek must minimize Δ . As mentioned at the beginning of the chapter, this simple idea inside all the technicalities is the heart of the proof: the x to be taken to y by f must be the x that is taken closest to y by f .

The function Δ is continuous and \overline{B} is compact, so the Extreme Value Theorem guarantees that Δ does indeed take a minimum on \overline{B} . Condition (5.5) guarantees that Δ takes no minimum on the boundary $\partial\overline{B}$. Therefore the minimum of Δ must occur at an interior point x of \overline{B} ; this interior point x must be a critical point of Δ , so all partial derivatives of Δ vanish at x . Thus by the Chain Rule,

$$0 = D_j \Delta(x) = -2 \sum_{i=1}^n (y_i - f_i(x)) D_j f_i(x) \quad \text{for } j = 1, \dots, n.$$

This condition is equivalent to the matrix equation

$$\begin{bmatrix} D_1 f_1(x) & \cdots & D_1 f_n(x) \\ \vdots & \ddots & \vdots \\ D_n f_1(x) & \cdots & D_n f_n(x) \end{bmatrix} \begin{bmatrix} y_1 - f_1(x) \\ \vdots \\ y_n - f_n(x) \end{bmatrix} = \begin{bmatrix} 0 \\ \vdots \\ 0 \end{bmatrix}$$

or

$$f'(x)^t (y - f(x)) = \mathbf{0}_n.$$

But $\det f'(x)^t = \det f'(x) \neq 0$ by condition (5.3), so $f'(x)^t$ is invertible and the only solution to the equation is $y - f(x) = \mathbf{0}_n$. This exhibits the desired x interior to \overline{B} such that $y = f(x)$. And there is only one such x because f is injective on \overline{B} . We no longer need the boundary $\partial\overline{B}$, whose role was to make a set compact. In sum, we now know that f is injective on B and that $f(B)$ contains W .

Let $V = f^{-1}(W) \cap B$, the set of all points $x \in B$ such that $f(x) \in W$. (See figure 5.5.) By the inverse image characterization of continuity (Theorem 5.1.2), V is open. We have established that $f : V \rightarrow W$ is inverted by $f^{-1} : W \rightarrow V$. To show that f^{-1} is continuous, substitute the values $\tilde{y} = f(\tilde{x})$, $y = f(x)$, $\tilde{x} = f^{-1}(\tilde{y})$, $x = f^{-1}(y)$ in (5.4) to obtain

$$|f^{-1}(\tilde{y}) - f^{-1}(y)| \leq 2|\tilde{y} - y| \quad \text{for all } y, \tilde{y} \in W.$$

The continuity of f^{-1} is clear from this.

The last thing to prove is that f^{-1} is differentiable on W . Again, reducing the problem makes it easier. By (5.3), the condition $\det f'(x) \neq 0$ is in effect at each $x \in V$. Therefore a is no longer a distinguished point of V , and it suffices to prove that the local inverse f^{-1} is differentiable at $f(a)$. Consider




Figure 5.5. The sets V and W of the Inverse Function Theorem

the mapping \tilde{f} defined by the formula $\tilde{f}(x) = f(x + a) - b$. Since $f(a) = b$ it follows that $\tilde{f}(\mathbf{0}_n) = \mathbf{0}_n$, and since \tilde{f} is f up to prepended and postpended translations, \tilde{f} is locally invertible at $\mathbf{0}_n$ and its derivative there is $D\tilde{f}_{\mathbf{0}} = Df_a = \text{id}_n$. The upshot is that in proving that f^{-1} is differentiable at $f(a)$, there is no loss of generality in normalizing to $a = \mathbf{0}_n$ and $f(a) = \mathbf{0}_n$ while also retaining the normalization that Df_a is the identity mapping.

For any point $k \in W$, let $h = f^{-1}(k)$. By condition (5.4), $|k| \geq \frac{1}{2}|h|$. Since $h \rightarrow \mathbf{0}_n$ exactly as $k \rightarrow \mathbf{0}_n$,

$$\begin{aligned} \lim_{k \rightarrow \mathbf{0}_n} \frac{|f^{-1}(\mathbf{0}_n + k) - f^{-1}(\mathbf{0}_n) - k|}{|k|} &= \lim_{h \rightarrow \mathbf{0}_n} \frac{|h - \mathbf{0}_n - f(h)|}{|k|} \\ &\leq \lim_{h \rightarrow \mathbf{0}_n} \frac{|f(\mathbf{0}_n + h) - f(\mathbf{0}_n) - h|}{\frac{1}{2}|h|}. \end{aligned}$$

This last limit is 0 since $Df_{\mathbf{0}_n}(h) = h$. This shows that f^{-1} is differentiable at $\mathbf{0}_n$ with the identity mapping for its derivative. \square

Note the range of mathematical skills that this proof of the Inverse Function Theorem required. The ideas were motivated and guided by pictures, but the actual argument was symbolic. At the level of fine detail, we normalized the derivative to the identity in order to reduce clutter, we made an adroit choice of quantifier in choosing a small enough \overline{B} to apply the Difference Magnification Lemma with $c = 1/(2n^2)$, and we used the full Triangle Inequality to obtain (5.4). This sufficed to prove that f is locally injective. Since the proof of the Difference Magnification Lemma used the Mean Value Theorem many times, the role of the Mean Value Theorem in the multivariable Inverse Function Theorem is thus similar to its role in the one-variable proof reviewed at the beginning of the chapter. However, while the one-variable proof that f is locally surjective relied on the Intermediate Value Theorem, the multivariable argument was far more elaborate. The idea was that the putative x taken by f to a given y must be the actual x taken by f closest to y . We exploited this idea by working in broad strokes:

- The Extreme Value Theorem from chapter 2 guaranteed that there was such an actual x .

- The Critical Point Theorem and then the Chain Rule from chapter 4 described necessary conditions associated to x .
- And finally, the Linear Invertibility Theorem from chapter 3 showed that indeed $f(x) = y$. Very satisfyingly, the hypothesis that the derivative is invertible sealed the argument that the mapping itself is locally invertible.

Indeed, the proof of local surjectivity used nearly every significant result from chapters 2 through 4 of these notes,

For an example, define $f : \mathbf{R}^2 \rightarrow \mathbf{R}^2$ by $f(x, y) = (x^3 - 2xy^2, x + y)$. Is f locally invertible at $(1, -1)$? If so, what is the best affine approximation to the inverse near $f(1, -1)$? To answer the first question, calculate the Jacobian

$$f'(1, -1) = \begin{bmatrix} 3x^2 - 2y^2 & -4xy \\ 1 & 1 \end{bmatrix} \Big|_{(x,y)=(1,-1)} = \begin{bmatrix} 1 & 4 \\ 1 & 1 \end{bmatrix}.$$

This is invertible with inverse $f'(1, -1)^{-1} = \frac{1}{3} \begin{bmatrix} -1 & 4 \\ 1 & -1 \end{bmatrix}$. Therefore f is locally invertible at $(1, -1)$ and the affine approximation to f^{-1} near $f(1, -1) = (-1, 0)$ is

$$f^{-1}(-1 + h, 0 + k) \approx \begin{bmatrix} 1 \\ -1 \end{bmatrix} + \frac{1}{3} \begin{bmatrix} -1 & 4 \\ 1 & -1 \end{bmatrix} \begin{bmatrix} h \\ k \end{bmatrix} = \left(1 - \frac{1}{3}h + \frac{4}{3}k, -1 + \frac{1}{3}h - \frac{1}{3}k\right).$$

The actual inverse function f^{-1} about $(-1, 0)$ may not be clear, but with the Inverse Function Theorem its affine approximation is easy to find.

Exercises

5.2.1. Define $f : \mathbf{R}^2 \rightarrow \mathbf{R}^2$ by $f(x, y) = (x^3 + 2xy + y^2, x^2 + y)$. Is f locally invertible at $(1, 1)$? If so, what is the best affine approximation to the inverse near $f(1, 1)$?

5.2.2. Same question for $f(x, y) = (x^2 - y^2, 2xy)$ at $(2, 1)$.

5.2.3. Same question for $C(r, \theta) = (r \cos \theta, r \sin \theta)$ at $(1, 0)$.

5.2.4. Same question for $C(\rho, \theta, \phi) = (\rho \cos \theta \sin \phi, \rho \sin \theta \sin \phi, \rho \cos \phi)$ at $(1, 0, \pi/2)$.

5.2.5. At what points $(a, b) \in \mathbf{R}^2$ is each of the following mappings guaranteed to be locally invertible by the Inverse Function Theorem? In each case, find the best affine approximation to the inverse near $f(a, b)$.

(a) $f(x, y) = (x + y, 2xy^2)$.

(b) $f(x, y) = (\sin x \cos y + \cos x \sin y, \cos x \cos y - \sin x \sin y)$.

5.2.6. Define $f : \mathbf{R}^2 \rightarrow \mathbf{R}^2$ by $f(x, y) = (e^x \cos y, e^x \sin y)$. Show that f is locally invertible at each point $(a, b) \in \mathbf{R}^2$, but that f is not globally invertible. Let $(a, b) = (0, \frac{\pi}{3})$; let $(c, d) = f(a, b)$; let g be the local inverse to f near (a, b) . Find an explicit formula for g , compute $g'(c, d)$ and verify that it agrees with $f'(a, b)^{-1}$.

5.2.7. If f and g are functions from \mathbf{R}^3 to \mathbf{R} , show that the mapping $F = (f, g, f+g) : \mathbf{R}^3 \rightarrow \mathbf{R}^3$, does not have a differentiable local inverse anywhere.

5.2.8. Define $f : \mathbf{R} \rightarrow \mathbf{R}$ by

$$f(x) = \begin{cases} x + 2x^2 \sin \frac{1}{x} & \text{if } x \neq 0 \\ 0 & \text{if } x = 0. \end{cases}$$

(a) Show that f is differentiable at $x = 0$ and that $f'(0) \neq 0$. (Since this is a one-dimensional problem you may verify the old definition of derivative rather than the new one.)

(b) Despite the result from (a), show that f is not locally invertible at $x = 0$. Why doesn't this contradict the Inverse Function Theorem?

5.3 The Implicit Function Theorem

Let n and c be positive integers with $c \leq n$, and let $r = n - c$. This section addresses the following question:

When do c conditions on n variables specify c of the variables in terms of the remaining r variables?

The symbols in this question will remain in play throughout the section. That is,

- $n = r + c$ is the total number of variables,
- c is the number of conditions, i.e., the number of constraints on the variables, and therefore the number of variables that might be dependent on the others,
- and r is the number of remaining variables and therefore the number of variables that might be free.

The word *conditions* (or *constraints*) provides a mnemonic for the symbol c , and similarly *remaining* (or *free*) provides a mnemonic for r .

The question can be rephrased:

When is a level set locally a graph?

To understand the rephrasing, we begin by reviewing the idea of a level set, given here in a slightly more general form than in Definition 4.6.3,

Definition 5.3.1 (Level Set). Let $g : A \rightarrow \mathbf{R}^m$ (where $A \subset \mathbf{R}^n$) be a mapping. A **level set** of g is the set of points in A that map under g to some fixed vector w in \mathbf{R}^m ,

$$L = \{v \in A : g(v) = w\}.$$

That is, L is the inverse image under g of the one-point set $\{w\}$.

Also we review the argument in section 4.6 that every graph is a level set. Let A_0 be a subset of \mathbf{R}^r , and let $f : A_0 \rightarrow \mathbf{R}^c$ be any mapping. Let $A = A_0 \times \mathbf{R}^c$ (a subset of \mathbf{R}^n) and define a second mapping $g : A \rightarrow \mathbf{R}^c$,

$$g(x, y) = f(x) - y, \quad (x, y) \in A_0 \times \mathbf{R}^c.$$

Then the graph of f is

$$\begin{aligned} \text{graph}(f) &= \{(x, y) \in A_0 \times \mathbf{R}^c : y = f(x)\} \\ &= \{(x, y) \in A : g(x, y) = \mathbf{0}_c\}, \end{aligned}$$

and this is the set of inputs to g that g takes to $\mathbf{0}_c$, a level set of g as desired.

Now we return to rephrasing the question at the beginning of this section. Let A be an open subset of \mathbf{R}^n , and let $g : A \rightarrow \mathbf{R}^c$ have continuous partial derivatives at every point of A . Points of A can be written (x, y) where $x \in \mathbf{R}^r$ and $y \in \mathbf{R}^c$. Consider the level set

$$L = \{(x, y) \in A : g(x, y) = \mathbf{0}_c\}.$$

The question was whether the c scalar conditions $g(x, y) = \mathbf{0}_c$ on the $n = c + r$ scalar entries of (x, y) define the c scalars of y in terms of the r scalars of x near (a, b) . That is, the question is whether the vector relation $g(x, y) = \mathbf{0}_c$ for (x, y) near (a, b) is equivalent to a vector relation $y = \varphi(x)$ for some mapping φ that takes r -vectors near a to c -vectors near b . This is precisely the question of whether the level set L is locally the graph of such a mapping φ . If the answer is yes, then we would like to understand φ as well as possible by using the techniques of differential calculus. In this context we view the mapping φ as *implicit* in the condition $g = \mathbf{0}_c$, explaining the name of the pending Implicit Function Theorem.

The first phrasing of the question, whether c conditions on n variables specify c of the variables in terms of the remaining r variables, is easy to answer when the conditions are affine. Affine conditions take the matrix form $Pv = w$ where $P \in M_{c,n}(\mathbf{R})$, $v \in \mathbf{R}^n$, and $w \in \mathbf{R}^c$. Partition the matrix P into a left square c -by- r block M and a right c -by- c block N , and partition the vector v into its first r entries h and its last c entries k . Then the relation $Pv = w$ is

$$[M \ N] \begin{bmatrix} h \\ k \end{bmatrix} = w,$$

that is,

$$Mh + Nk = w.$$

Assume that N is invertible. Then subtracting Mh from both sides and then left multiplying by N^{-1} shows that the relation is

$$k = N^{-1}(w - Mh).$$

Thus, when the right c -by- c submatrix of P is invertible, the relation $Pv = w$ explicitly specifies the last c variables of v in terms of the first r variables. A similar statement applies to any invertible c -by- c submatrix of P and the corresponding variables. A special case of this calculation, the linear case, will be used throughout the section: for any $M \in M_{c,r}(\mathbf{R})$, any invertible $N \in M_c(\mathbf{R})$, any $h \in \mathbf{R}^r$, and any $k \in \mathbf{R}^c$,

$$[M \ N] \begin{bmatrix} h \\ k \end{bmatrix} = \mathbf{0}_c \quad \iff \quad k = -N^{-1}Mh. \quad (5.6)$$

When the conditions are nonaffine the situation is not so easy to analyze. However:

- The problem is easy to *linearize*. That is, given a point (a, b) (where $a \in \mathbf{R}^r$ and $b \in \mathbf{R}^c$) on the level set $\{(x, y) : g(x, y) = w\}$, differential calculus tells us how to describe the tangent object to the level set at the point. Depending on the value of r , the tangent object will be a line, or a plane, or higher-dimensional. But regardless of its dimension, it is described by the linear conditions $g'(a, b)v = \mathbf{0}_c$, and these conditions take the form that we have just considered,

$$[M \ N] \begin{bmatrix} h \\ k \end{bmatrix} = \mathbf{0}_c, \quad M \in M_{c,r}(\mathbf{R}), \quad N \in M_c(\mathbf{R}), \quad h \in \mathbf{R}^r, \quad k \in \mathbf{R}^c.$$

Thus if N is invertible then we can solve the linearized problem as in (5.6).

- The Inverse Function Theorem says:
If the linearized inversion problem is solvable, then the nonlinear inversion problem is locally solvable.

With a little work, we can use the Inverse Function Theorem to establish the Implicit Function Theorem:

If the linearized implicit function problem is solvable, then the nonlinear implicit function problem is locally solvable.

And in fact, the Implicit Function Theorem will imply the Inverse Function Theorem as well.

For example, the unit circle C is described by one constraint on two variables,

$$x^2 + y^2 = 1.$$

Globally (in the large), this relation neither specifies x as a function of y nor y as a function of x . It can't: the circle is visibly not the graph of a function of either sort—recall the Vertical Line Test to check whether a curve is the graph of a function $y = \varphi(x)$, and analogously for the Horizontal Line Test. The situation improves, however, if one works locally (in the small) by looking at just part of the circle at a time. Any arc in the bottom half of the circle is described by the function

$$y = \varphi(x) = -\sqrt{1 - x^2}.$$

Similarly, any arc in the right half is described by

$$x = \psi(y) = \sqrt{1 - y^2}.$$

Any arc in the bottom right quarter is described by both functions. (See figure 5.6.) On the other hand, no arc of the circle about the point $(a, b) = (1, 0)$ is described by a function $y = \varphi(x)$, and no arc about $(a, b) = (0, 1)$ is described by a function $x = \psi(y)$. (See figure 5.7.) Thus, about some points (a, b) , the circle relation $x^2 + y^2 = 1$ contains the information to specify each variable as a function of the other. These functions are implicit in the relation. About other points, the relation implicitly defines one variable as a function of the other, but not the second as a function of the first.




Figure 5.6. Arc of a circle




Figure 5.7. Trickier arcs of a circle

To bring differential calculus to bear on the situation, think of the circle as a level set. Consider the function $g(x, y) = x^2 + y^2$. The circle is precisely

$$C = \{(x, y) : g(x, y) = 1\}.$$

Let (a, b) be a point on the circle. The derivative of g at the point is

$$g'(a, b) = [D_1g(a, b) \ D_2g(a, b)] = [2a \ 2b].$$

The tangent line to the circle at (a, b) consists of the points $(a + h, b + k)$ such that (h, k) is orthogonal to $g'(a, b)$,

$$[2a \ 2b] \begin{bmatrix} h \\ k \end{bmatrix} = 0.$$

That is,

$$2ah + 2bk = 0.$$

Thus whenever $b \neq 0$ we have

$$k = -(a/b)h,$$

showing that on the tangent line, the second coordinate is a linear function of the first, and the function has derivative $-a/b$. And so on the circle itself near (a, b) , plausibly the second coordinate is a function of the first as well, provided that $b \neq 0$. Note that indeed this argument excludes the two points $(1, 0)$ and $(-1, 0)$ about which y is not an implicit function of x . But about points $(a, b) \in C$ where $D_2g(a, b) \neq 0$, the circle relation should implicitly define y as a function of x . And at such points, differentiating the circle relation $x^2 + y^2 = 1$ where $y = \varphi(x)$ gives

$$2a + 2b\varphi'(a) = 0,$$

or

$$\varphi'(a) = -a/b.$$

The derivative $-a/b$ is exactly as predicted by solving the linearized problem. The reader may recall from elementary calculus that this technique is called *implicit differentiation*.

It may help the reader visualize the situation if we revisit the idea of the previous paragraph more geometrically. Since C is a level set of g , the gradient $g'(a, b)$ is orthogonal to C at the point (a, b) . When $g'(a, b)$ has a nonzero y -component, C should locally have a big shadow on the x -axis, from which there is a function back to C . (See figure 5.8, in which the arrow drawn is quite a bit shorter than the true gradient, for graphical reasons.)

In the case of the circle, the arguments of the previous two paragraphs are unnecessary because we can explicitly solve for y in terms of x and vice versa. But the purpose of the previous two paragraphs is to illustrate general ideas in a specific context.

Another set defined by a relation is the unit sphere, also specified as a level set. Let

$$g(x, y, z) = x^2 + y^2 + z^2.$$




Figure 5.8. Nonhorizontal gradient and x -shadow

Then the sphere is

$$S = \{(x, y, z) : g(x, y, z) = 1\}.$$

Imposing one condition on three variables should generally leave two of them free (say, the first two) and define the remaining one in terms of the free ones. And indeed, the sphere implicitly describes z as a function $\varphi(x, y)$ about any point $p = (a, b, c) \in S$ off the equator, where $c \neq 0$. (Note: c does not denote the number of constraints in this example.) The equator is precisely the points where $D_3g(p) = 2c$ vanishes. Again geometry makes this plausible. The gradient $g'(p)$ is orthogonal to S at p . When $g'(p)$ has a nonzero z -component, S should locally have a big shadow in the (x, y) -plane from which there is a function back to S and then to the z -axis. (See figure 5.9.)




Figure 5.9. Function from the (x, y) -plane to the z -axis via the sphere

The argument based on calculus and linear algebra to suggest that near points $(a, b, c) \in S$ such that $D_3g(a, b, c) \neq 0$, z is implicitly a function $\varphi(x, y)$ on S is similar to the case of the circle. The derivative of g at the point is

$$g'(a, b, c) = [D_1g(a, b, c) \ D_2g(a, b, c) \ D_3g(a, b, c)] = [2a \ 2b \ 2c].$$

The tangent plane to the sphere at (a, b, c) consists of the points $(a + h, b + k, c + \ell)$ such that (h, k, ℓ) is orthogonal to $g'(a, b, c)$,

$$[2a \ 2b \ 2c] \begin{bmatrix} h \\ k \\ \ell \end{bmatrix} = 0.$$

That is,

$$2ah + 2bk + 2c\ell = 0.$$

Thus whenever $c \neq 0$ we have

$$\ell = -(a/c)h - (b/c)k,$$

showing that on the tangent plane, the third coordinate is a linear function of the first two, and the function has partial derivatives $-a/c$ and $-b/c$. And so on the sphere itself near (a, b, c) , plausibly the third coordinate is a function of the first two as well, provided that $c \neq 0$. This argument excludes points on the equator, about which z is not an implicit function of (x, y) . But about points $(a, b, c) \in S$ where $D_3g(a, b, c) \neq 0$, the sphere relation should implicitly define z as a function of (x, y) . And at such points, differentiating the sphere relation $x^2 + y^2 + z^2 = 1$ with respect to x and y , where $z = \varphi(x, y)$, gives

$$2a + 2cD_1\varphi(a, b) = 0, \quad 2b + 2cD_2\varphi(a, b) = 0,$$

or

$$D_1\varphi(a, b) = -a/c, \quad D_2\varphi(a, b) = -b/c.$$

The partial derivatives are exactly as predicted by solving the linearized problem.

Next consider the intersection of the unit sphere and the 45-degree plane $z = -y$. This is a great circle, again naturally described as a level set. That is, if we consider the mapping

$$g : \mathbf{R}^3 \longrightarrow \mathbf{R}^2, \quad g(x, y, z) = (x^2 + y^2 + z^2, y + z),$$

then the great circle is a level set of g ,

$$GC = \{(x, y, z) : g(x, y, z) = (1, 0)\}.$$

The two conditions on the three variables should generally leave one variable (say, the first one) free and define the other two variables in terms of it. Indeed,

GC is a circle that is orthogonal to the plane of the page, and away from its two points $(\pm 1, 0, 0)$ that are farthest in and out of the page, it does define (y, z) locally as functions of x . (See figure 5.10.) To make the implicit function in the great circle relations explicit, note that near the point $p = (a, b, c)$ in the figure,

$$(y, z) = (\varphi_1(x), \varphi_2(x)) = \left(-\sqrt{\frac{1-x^2}{2}}, \sqrt{\frac{1-x^2}{2}} \right).$$

At p the component functions have derivatives

$$\varphi_1'(a) = \frac{a}{2\sqrt{\frac{1-a^2}{2}}} \quad \text{and} \quad \varphi_2'(a) = \frac{-a}{2\sqrt{\frac{1-a^2}{2}}}.$$

But $1 - a^2 = 2b^2 = 2c^2$, and $\sqrt{b^2} = -b$ since $b < 0$ while $\sqrt{c^2} = c$ since $c > 0$, so the derivatives are

$$\varphi_1'(a) = -\frac{a}{2b} \quad \text{and} \quad \varphi_2'(a) = -\frac{a}{2c}. \quad (5.7)$$




Figure 5.10. y and z locally as functions of x on a great circle

Now we show that linearizing the problem reproduces the results without explicitly finding φ_1 and φ_2 . The derivative matrix of g at p is

$$g'(a, b, c) = \begin{bmatrix} 2a & 2b & 2c \\ 0 & 1 & 1 \end{bmatrix}.$$

The level set GC is defined by the condition that g remain constant as (x, y, z) varies. Thus the tangent line to GC at a point (a, b, c) consists of points $(a+h, b+k, c+\ell)$ such that neither component function of g is instantaneously changing in the (h, k, ℓ) -direction,

$$\begin{bmatrix} 2a & 2b & 2c \\ 0 & 1 & 1 \end{bmatrix} \begin{bmatrix} h \\ k \\ \ell \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \end{bmatrix}.$$

The right 2-by-2 submatrix of $g'(a, b, c)$ has nonzero determinant whenever $b \neq c$, that is, at all points of GC except the two aforementioned extreme points $(\pm 1, 0, 0)$. Assuming that $b \neq c$, let M denote the first column of $g'(a, b, c)$ and let N denote the right 2-by-2 submatrix. Then by (5.6), the linearized problem has solution

$$\begin{bmatrix} k \\ \ell \end{bmatrix} = -N^{-1}Mh = \frac{1}{2(c-b)} \begin{bmatrix} 1 & -2c \\ -1 & 2b \end{bmatrix} \begin{bmatrix} 2a \\ 0 \end{bmatrix} h = \begin{bmatrix} -\frac{a}{2b} \\ -\frac{a}{2c} \end{bmatrix} h. \quad (5.8)$$

(The condition $c = -b$ was used in the last step.) That is, for all points $(a+h, b+k, c+\ell)$ on the tangent line to GC at (a, b, c) , the last two coordinate-offsets k and ℓ are specified in terms of the first coordinate offset h via (5.8), and the component functions have partial derivatives $-a/(2b)$ and $-a/(2c)$. Predictably enough, these are the previously-calculated component derivatives of the true mapping φ defining the last two coordinates y and z in terms of the first coordinate x for points on GC itself near p , shown in (5.7).

In the examples of the circle, the sphere, and the great circle, the functions implicit in the defining relations could in fact be found explicitly. But in general, relations may snarl the variables so badly that we expressing some as functions of the others is beyond our algebraic capacity. For instance, do the simultaneous conditions

$$y^2 = e^z \cos(y + x^2) \quad \text{and} \quad y^2 + z^2 = x^2 \quad (5.9)$$

define y and z implicitly as functions $y = \varphi_1(x)$, $z = \varphi_2(x)$ near the point $(1, -1, 0)$? (This point meets both conditions.) Answering this directly by solving for y and z is manifestly unappealing. But linearizing the problem is easy. At our point $(1, -1, 0)$, the mapping

$$g(x, y, z) = (y^2 - e^z \cos(y + x^2), y^2 + z^2 - x^2)$$

has derivative matrix

$$\begin{aligned} g'(1, -1, 0) &= \begin{bmatrix} 2xe^z \sin(y + x^2) & 2y + e^z \sin(y + x^2) & -e^z \cos(y + x^2) \\ -2x & 2y & 2z \end{bmatrix} \Big|_{(1, -1, 0)} \\ &= \begin{bmatrix} 0 & -2 & -1 \\ -2 & -2 & 0 \end{bmatrix}. \end{aligned}$$

Since the right 2-by-2 determinant is nonzero, we expect that indeed y and z are implicit functions $\varphi_1(x)$ and $\varphi_2(x)$ near $(1, -1, 0)$. Furthermore, solving the linearized problem as in the previous example with M and N similarly defined suggests that if $(y, z) = \varphi(x) = (\varphi_1(x), \varphi_2(x))$ then

$$\varphi'(1) = -N^{-1}M = - \begin{bmatrix} -2 & -1 \\ -2 & 0 \end{bmatrix}^{-1} \begin{bmatrix} 0 \\ -2 \end{bmatrix} = \frac{1}{2} \begin{bmatrix} 0 & 1 \\ 2 & -2 \end{bmatrix} \begin{bmatrix} 0 \\ -2 \end{bmatrix} = \begin{bmatrix} -1 \\ 2 \end{bmatrix}.$$

Thus for a point $(x, y, z) = (1 + h, -1 + k, 0 + l)$ near $(1, -1, 0)$ satisfying conditions (5.9), we expect that $(k, l) \approx (-h, 2h)$, i.e.,

$$y \approx -1 - h \quad \text{and} \quad z \approx 0 + 2h \quad \text{where } x = 1 + h.$$

The Implicit Function Theorem fulfills these expectations.

Theorem 5.3.2 (Implicit Function Theorem). *Let c and n be positive integers with $n > c$, and let $r = n - c$. Let A be an open subset of \mathbf{R}^n , and let $g : A \rightarrow \mathbf{R}^c$ have continuous partial derivatives at every point of A . Consider the level set*

$$L = \{v \in A : g(v) = \mathbf{0}_c\}.$$

Let p be a point of L , i.e., let $g(p) = \mathbf{0}_c$. Let $p = (a, b)$ where $a \in \mathbf{R}^r$ and $b \in \mathbf{R}^c$, and let $g'(p) = [M \ N]$ where M is the left c -by- r submatrix and N is the remaining right square c -by- c submatrix.

If $\det N \neq 0$ then the level set L is locally a graph near p . That is, the condition $g(x, y) = \mathbf{0}_c$ for (x, y) near (a, b) implicitly defines y as a function $y = \varphi(x)$ where φ takes r -vectors near a to c -vectors near b , and in particular $\varphi(a) = b$. The function φ is differentiable at a with derivative matrix $\varphi'(a) = -N^{-1}M$. Hence φ is well approximated near a by its affine approximation,

$$\varphi(a + h) \approx b - N^{-1}Mh.$$

We make three remarks before the proof.

- The condition $g(x, y) = \mathbf{0}_c$ could just as easily be $g(x, y) = w$ for any fixed point $w \in \mathbf{R}^c$, as in our earlier examples. Normalizing to $w = \mathbf{0}_c$ amounts to replacing g by $g - w$, which we do to tidy up the statement of the theorem.
- The Implicit Function Theorem gives no information when $\det N = 0$. In this case, the condition $g(x, y) = \mathbf{0}_c$ may or may not define y in terms of x .
- While the theorem strictly addresses only whether the last c of n variables subject to c conditions depend on the first r variables, it can be suitably modified to address whether any c variables depend on the remaining ones. This is a matter of reindexing or permuting the variables, not worth the cumbersome notation of discussing formally, but feel free to use the modified version.

Proof. Examining the derivative has already shown the theorem’s plausibility in specific instances. Shoring up these considerations into a proof is easy with a little trick and the Inverse Function Theorem, which produces a mapping given an invertible derivative. The trick is to define

$$G : A \longrightarrow \mathbf{R}^n$$

as follows: for all $x \in \mathbf{R}^r$ and $y \in \mathbf{R}^c$ such that $(x, y) \in A$,

$$G(x, y) = (x, g(x, y)).$$

Note that G incorporates g , but unlike g it is a map between spaces of the same dimension n . Note also that the augmentation that changes g into G is highly invertible, being the identity mapping on the x -coordinates. That is, it is easy to recover g from G . The mapping G affects only y -coordinates, and it is designed to take the level set $L = \{(x, y) \in A : g(x, y) = \mathbf{0}_c\}$ to the x -axis. (See figure 5.11, in which the inputs and the outputs of G are shown in the same copy of \mathbf{R}^n .)




Figure 5.11. Mapping A to \mathbf{R}^n and the constrained set to x -space

The mapping G is differentiable at the point $p = (a, b)$ with derivative matrix

$$G'(a, b) = \begin{bmatrix} I_r & \mathbf{0}_{r \times c} \\ M & N \end{bmatrix} \in M_n(\mathbf{R}).$$

This matrix has determinant $\det G'(a, b) = \det N \neq 0$, and so by the Inverse Function Theorem G has a local inverse mapping Φ defined near the point $G(a, b) = (a, \mathbf{0}_c)$. (See figure 5.12.) Since the first r components of G are the identity mapping, the same holds for the inverse. That is, the inverse takes the form

$$\Phi(x, y) = (x, \phi(x, y)),$$

where ϕ maps n -vectors near $(a, \mathbf{0}_c)$ to c -vectors near b . The inversion criterion is that for all (x, y) near (a, b) and all (x, \tilde{y}) near $(a, \mathbf{0}_c)$,

$$G(x, y) = (x, \tilde{y}) \iff (x, y) = \Phi(x, \tilde{y}).$$

Equivalently, since neither G nor Φ affects x -coordinates, for all x near a , y near b , and \tilde{y} near $\mathbf{0}_c$,

$$g(x, y) = \tilde{y} \iff y = \phi(x, \tilde{y}). \tag{5.10}$$

Also by the Inverse Function Theorem and a short calculation.

$$\Phi'(a, \mathbf{0}_c) = G'(a, b)^{-1} = \begin{bmatrix} I_r & \mathbf{0}_{r \times c} \\ -N^{-1}M & N^{-1} \end{bmatrix}$$




Figure 5.12. Local inverse of G

Now we can exhibit the desired mapping implicit in the original g . Define a mapping

$$\varphi(x) = \phi(x, \mathbf{0}_c) \text{ for } x \text{ near } a. \tag{5.11}$$

The idea is that locally this lifts the x -axis to the level set L where $g(x, y) = \mathbf{0}_c$ and then projects horizontally to the y -axis. (See figure 5.13.) For any (x, y) near (a, b) , a specialization of condition (5.10) combines with the definition (5.11) of φ to give

$$g(x, y) = \mathbf{0}_c \iff y = \varphi(x).$$

This exhibits y as a local function of x on the level set of g , as desired. And since by definition (5.11), φ is the last c component functions of Φ restricted to the first r inputs to Φ , the derivative $\varphi'(a)$ is exactly the lower left c -by- r block of $\Phi'(a, \mathbf{0}_c)$, which is $-N^{-1}M$. This completes the proof. \square




Figure 5.13. The implicit mapping from x -space to y -space via the level set

Thus the Implicit Function Theorem follows easily from the Inverse Function Theorem. The converse implication is even easier. Imagine a scenario where somehow we know the Implicit Function Theorem but not the Inverse Function Theorem. Let $f : A \rightarrow \mathbf{R}^n$ (where $A \subset \mathbf{R}^n$) be a mapping that satisfies the hypotheses for the Inverse Function Theorem at a point $a \in A$. That is, f is continuously differentiable in an open set containing a , and $\det f'(a) \neq 0$. Define a mapping

$$g : A \times \mathbf{R}^n \rightarrow \mathbf{R}^n, \quad g(x, y) = f(x) - y.$$

(This should look familiar from the beginning of the section.) Let $b = f(a)$. Then $g(a, b) = 0$, and the derivative matrix of g at (a, b) is

$$g'(a, b) = [f'(a) \quad -I_n].$$

Since $f'(a)$ is invertible, we may apply the Implicit Function Theorem, with the roles of c , r , and n in theorem taken by the values n , n , and $2n$ here, and with the theorem modified as in the third remark before its proof so that we are checking whether the first n variables depend on the last n values. The theorem supplies us with a differentiable mapping φ defined for values of y near b such that for all (x, y) near (a, b) ,

$$g(x, y) = 0 \iff x = \varphi(y).$$

But by the definition of g , this equivalence is

$$y = f(x) \iff x = \varphi(y).$$

That is, φ inverts f . Also by the Implicit Function Theorem, the derivative of φ at b is

$$\varphi'(b) = -f'(a)^{-1}(-I_n) = f'(a)^{-1},$$

and we have recovered the Inverse Function Theorem. However, this is not really impressive since proving the Implicit Function Theorem without citing the Inverse Function Theorem would be just as hard as the route we took of proving the Inverse Function Theorem first. The point is that the two theorems have essentially the same content. (See figure 5.14.)




Figure 5.14. The Inverse Function Theorem from the Implicit Function Theorem

Exercises

5.3.1. Does the relation $x^2 + y + \sin(xy) = 0$ implicitly define y as a function of x near the origin? If so, what is its best affine approximation? How about x as a function of y and its affine approximation?

5.3.2. Does the relation $xy - z \log y + e^{xz} = 1$ implicitly define z as a function of (x, y) near $(0, 1, 1)$? How about y as a function of (x, z) ? When possible, give the affine approximation to the function.

5.3.3. Do the simultaneous conditions $x^2(y^2 + z^2) = 5$ and $(x - z)^2 + y^2 = 2$ implicitly define (y, z) as a function of x near $(1, -1, 2)$? If so, then what is the function's affine approximation?

5.3.4. Same question for the conditions $x^2 + y^2 = 4$ and $2x^2 + y^2 - 8z^2 = 8$ near $(2, 0, 0)$.

5.3.5. Do the simultaneous conditions $xy + 2yz = 3xz$ and $xyz + x - y = 1$ implicitly define (x, y) as a function of z near $(1, 1, 1)$? How about (x, z) as a function of y ? How about (y, z) as a function of x ? Give affine approximations when possible.

5.3.6. Do the conditions $xy^2 + xzu + yv^2 = 3$ and $u^3yz + 2xv - u^2v^2 = 2$ implicitly define (u, v) in terms of (x, y, z) near the point $(1, 1, 1, 1, 1)$? If so, what is the derivative matrix of the implicitly defined mapping at $(1, 1, 1)$?

5.3.7. Do the conditions $x^2 + yu + xv + w = 0$ and $x + y + uvw = -1$ implicitly define (x, y) in terms of (u, v, w) near $(x, y, u, v, w) = (1, -1, 1, 1, -1)$? If so, what is the best affine approximation to the implicitly defined mapping?

5.3.8. Do the conditions

$$\begin{aligned}2x + y + 2z + u - v &= 1 \\xy + z - u + 2v &= 1 \\yz + xz + u^2 + v &= 0\end{aligned}$$

define the first three variables (x, y, z) as a function $g(u, v)$ near the point $(x, y, z, u, v) = (1, 1, -1, 1, 1)$? If so, find the derivative matrix $g'(1, 1)$.

5.3.9. Define $g : \mathbf{R}^2 \rightarrow \mathbf{R}$ by $g(x, y) = 2x^3 - 3x^2 + 2y^3 + 3y^2$ and let L be the level set $\{(x, y) : g(x, y) = 0\}$. Find those points of L about which y need not be defined implicitly as a function of x , and find the points about which x need not be defined implicitly as a function of y . Describe L precisely—the result should explain the points you found.

5.4 Lagrange Multipliers: Geometric Motivation and Specific Examples

How close does the intersection of the planes $x + y + z = 1$ and $x - y + 2z = -1$ in \mathbf{R}^3 come to the origin? This is an example of an *optimization problem with constraints*. The goal in such problems is to maximize or minimize some function, but with relations imposed on its variables. Equivalently, the problem is to optimize some function whose domain is a level set.

A geometric solution to the sample problem just given is that the planes intersect in a line through the point $p = (0, 1, 0)$ in direction $d = (1, 1, 1) \times (1, -1, 2)$, so the point-to-line distance formula from exercise 3.10.11 answers the question. This is easy and efficient.

A more generic method of solution is via substitution. The equations of the constraining planes are $x + y = 1 - z$ and $x - y = -1 - 2z$; adding gives $x = -3z/2$, subtracting gives $y = 1 + z/2$. (These are easily checked.) To finish the problem, minimize the function $d^2(z) = (-3z/2)^2 + (1 + z/2)^2 + z^2$, where d^2 denotes distance squared from the origin. Minimizing d^2 rather than d avoids square roots.

Not all constrained problems yield readily to either of these methods. The more irregular the conditions, the less amenable they are to geometry, and the more tangled the variables, the less readily they distill. Merely adding more

variables to the previous problem produces a nuisance: How close does the intersection of the planes $v + w + x + y + z = 1$ and $v - w + 2x - y + z = -1$ in \mathbf{R}^5 come to the origin? Now no geometric procedure lies conveniently at hand. As for substitution, linear algebra shows that

$$\begin{bmatrix} 1 & 1 & 1 & 1 & 1 \\ 1 & -1 & 2 & -1 & 1 \end{bmatrix} \begin{bmatrix} v \\ w \\ x \\ y \\ z \end{bmatrix} = \begin{bmatrix} 1 \\ -1 \end{bmatrix}$$

implies

$$\begin{bmatrix} v \\ w \end{bmatrix} = \begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix}^{-1} \left(\begin{bmatrix} 1 \\ -1 \end{bmatrix} - \begin{bmatrix} 1 & 1 & 1 \\ 2 & -1 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ z \end{bmatrix} \right) = \begin{bmatrix} -3x/2 - z \\ 1 + x/2 - y \end{bmatrix}.$$

Since the resulting function $d^2(x, y, z) = (-3x/2 - z)^2 + (1 + x/2 - y)^2 + x^2 + y^2 + z^2$ is quadratic, partial differentiation and more linear algebra will find its critical points. But the process is getting tedious.

Let's step back from specifics (but we will return to the currently unresolved example soon) and consider in general the necessary nature of a critical point in a constrained problem. The discussion will take place in two stages: first we consider the domain of the problem, and then we consider the critical point.

The domain of the problem is the points in n -space that satisfy a set of c constraints. To satisfy the constraints is to meet a condition

$$g(x) = \mathbf{0}_c$$

where $g : A \rightarrow \mathbf{R}^c$ is a \mathcal{C}^1 -mapping, with $A \subset \mathbf{R}^n$ an open set. That is, the constrained set forming the domain in the problem is a level set L , the intersection of the level sets of the component functions g_i of g . (See figures 5.15 and 5.16. The first figure shows two individual level sets for scalar-valued functions on \mathbf{R}^3 , and the second figure shows them together and then shows their intersection, the level set for a vector-valued mapping.)

At any point $p \in L$, the set L must be locally orthogonal to each gradient $\nabla g_i(p)$. (See figures 5.17 and 5.18. The first figure shows the level sets for the component functions of the constraint mapping, and the gradients of the component functions at p , while the second figure shows the tangent line and the normal plane to the level set at p . In the first figure, neither gradient is tangent to the other surface, and so in the second figure the two gradients are not normal to one another.) Therefore:

- L is orthogonal at p to every linear combination of the gradients,

$$\sum_{i=1}^c \lambda_i \nabla g_i(p) \quad \text{where } \lambda_1, \dots, \lambda_c \text{ are scalars.}$$




Figure 5.15. Level sets for two scalar-valued functions on \mathbf{R}^3




Figure 5.16. The intersection is a level set for a vector-valued mapping on \mathbf{R}^3

Equivalently:

- Every such linear combination of gradients is orthogonal to L at p .

But we want to turn this idea around and assert the converse, that:

- Every vector that is orthogonal to L at p is such a linear combination.

This does not always follow. Intuitively, the argument is that if the gradients $\nabla g_1(p), \dots, \nabla g_c(p)$ are linearly independent (i.e., they point in c nonredundant directions) then the Implicit Function Theorem should say that the level set L therefore looks $(n - c)$ -dimensional near p , so the space of vectors orthogonal to L at p is c -dimensional, and so any such vector is indeed a linear combination of the gradients. This is not a proof, but for now it is a good heuristic.

Proceeding to the second stage of the discussion, now suppose that p is a critical point of the restriction to L of some \mathcal{C}^1 -function $f : A \rightarrow \mathbf{R}$. (Thus f has the same domain $A \subset \mathbf{R}^n$ as g .) Then for any unit vector d




Figure 5.17. Gradients to the level sets at a point of intersection




Figure 5.18. Tangent line and normal plane to the intersection

describing a direction in L at p , the directional derivative $D_d f(p)$ must be 0. But $D_d f(p) = \langle \nabla f(p), d \rangle$, so this means that:

- $\nabla f(p)$ must be orthogonal to L at p .

This observation combines with our description of the most general vector orthogonal to L at p , in the third bullet above, to give Lagrange's condition:

Let p be a critical point of the function f restricted to the level set $L = \{x : g(x) = 0_c\}$ of g . If the gradients $\nabla g_i(p)$ are linearly independent, then for some scalars $\lambda_1, \dots, \lambda_c$,

$$\nabla f(p) = \sum_{i=1}^c \lambda_i \nabla g_i(p),$$

and since p is in the level set, also

$$g(p) = \mathbf{0}_c.$$

Approaching a constrained problem by setting up these conditions and then working with the new variables $\lambda_1, \dots, \lambda_c$ is sometimes easier than the other methods. The λ_i are *useful but irrelevant constants*. The next section will prove the Lagrange criterion carefully and then give some general examples. The remainder of this section is dedicated to specific examples,

Returning to the unresolved second example at the beginning of the section, the functions in question are

$$\begin{aligned} f(v, w, x, y, z) &= v^2 + w^2 + x^2 + y^2 + z^2 \\ g_1(v, w, x, y, z) &= v + w + x + y + z - 1 \\ g_2(v, w, x, y, z) &= v - w + 2x - y + z + 1 \end{aligned}$$

and the corresponding Lagrange condition and constraints are (after absorbing a 2 into the λ 's, whose particular values are irrelevant anyway)

$$\begin{aligned} (v, w, x, y, z) &= \lambda_1(1, 1, 1, 1, 1) + \lambda_2(1, -1, 2, -1, 1) \\ &= (\lambda_1 + \lambda_2, \lambda_1 - \lambda_2, \lambda_1 + 2\lambda_2, \lambda_1 - \lambda_2, \lambda_1 + \lambda_2) \\ v + w + x + y + z &= 1 \\ v - w + 2x - y + z &= -1. \end{aligned}$$

Substitute the expressions from the Lagrange condition into the constraints to get $5\lambda_1 + 2\lambda_2 = 1$ and $2\lambda_1 + 8\lambda_2 = -1$. So

$$\begin{bmatrix} \lambda_1 \\ \lambda_2 \end{bmatrix} = \begin{bmatrix} 5 & 2 \\ 2 & 8 \end{bmatrix}^{-1} \begin{bmatrix} 1 \\ -1 \end{bmatrix} = \begin{bmatrix} 10/36 \\ -7/36 \end{bmatrix}.$$

Note how much more convenient the two λ 's are to work with than the five original variables. Their values are auxiliary to the original problem, but substituting back now gives that the nearest point to the origin is

$$(v, w, x, y, z) = \frac{1}{36}(3, 17, -4, 17, 3)$$

and its distance is $\sqrt{612}/36$. This example is just one instance of a general problem of finding the nearest point to the origin in \mathbf{R}^n subject to c affine constraints. We will solve the general problem in the next section.

An example from geometry is Euclid's Least Area Problem. Given an angle ABC and a point P interior to the angle as shown in figure 5.19, what line through P cuts off from the angle the triangle of least area?

Draw the line L through P parallel to AB and let D be its intersection with AC . Let a denote the distance AD and let h denote the altitude from AC to P . Both a and h are constants. Given any other line L' through P , let x denote its intersection with AC and H denote the altitude from AC to




Figure 5.19. Setup for Euclid's Least Area Problem




Figure 5.20. Construction for Euclid's Least Area Problem

the intersection of L' with AB . (See figure 5.20.) The shaded triangle and its subtriangle in the figure are similar, giving the relation $x/H = (x - a)/h$.

The problem is now to minimize the function $f(x, H) = \frac{1}{2}xH$ subject to the constraint $g(x, H) = 0$ where $g(x, H) = (x - a)H - xh = 0$. Lagrange's condition $\nabla f(x, H) = \lambda \nabla g(x, H)$ and the constraint $g(x, H) = 0$ become, after absorbing a 2 into λ ,

$$\begin{aligned} (H, x) &= \lambda(H - h, x - a), \\ (x - a)H &= xh. \end{aligned}$$

The first relation quickly yields $(x - a)H = x(H - h)$. Combining this with the second shows that $H - h = h$, that is, $H = 2h$. The solution to Euclid's problem is, therefore, to take the segment that is *bisected* by P between the two sides of the angle. (See figure 5.21.)




Figure 5.21. Solution to Euclid's Least Area Problem

An example from optics is Snell's Law. A particle travels through medium 1 at speed v , and through medium 2 at speed w . If the particle travels from point A to point B as shown in the least possible amount of time, what is the relation between angles α and β ? (See figure 5.22.)




Figure 5.22. Geometry of Snell's Law

Since time is distance over speed, a little trigonometry shows that this problem is equivalent to minimizing $f(\alpha, \beta) = a \sec \alpha / v + b \sec \beta / w$ subject to the constraint $g(\alpha, \beta) = a \tan \alpha + b \tan \beta = d$. (g measures lateral distance traveled.) The Lagrange condition $\nabla f(\alpha, \beta) = \lambda \nabla g(\alpha, \beta)$ is

$$\left(\frac{a}{v} \sin \alpha \sec^2 \alpha, \frac{b}{w} \sin \beta \sec^2 \beta \right) = \lambda (a \sec^2 \alpha, b \sec^2 \beta).$$

Therefore $\lambda = \sin \alpha / v = \sin \beta / w$, giving Snell's famous relation,

$$\frac{\sin \alpha}{\sin \beta} = \frac{v}{w}.$$

For an example from analytic geometry, let the function f measure the square of the distance between the points $x = (x_1, x_2)$ and $y = (y_1, y_2)$ in the plane,

$$f(x_1, x_2, y_1, y_2) = (x_1 - y_1)^2 + (x_2 - y_2)^2.$$

Fix points $a = (a_1, a_2)$ and $b = (b_1, b_2)$ in the plane, and fix positive numbers r and s . Define

$$\begin{aligned} g_1(x_1, x_2) &= (x_1 - a_1)^2 + (x_2 - a_2)^2 - r^2, \\ g_2(y_1, y_2) &= (y_1 - b_1)^2 + (y_2 - b_2)^2 - s^2 \\ g(x_1, x_2, y_1, y_2) &= (g_1(x_1, x_2), g_2(y_1, y_2)). \end{aligned}$$

Then the set of four-tuples (x_1, x_2, y_1, y_2) such that

$$g(x_1, x_2, y_1, y_2) = (0, 0)$$

can be viewed as the set of pairs of points x and y that lie respectively on the circles centered at a and b with radii r and s . Thus, to optimize the function f subject to the constraint $g = \mathbf{0}$ is to optimize the distance between pairs of points on the circles. The rows of the 2-by-4 matrix

$$g'(x, y) = 2 \begin{bmatrix} x_1 - a_1 & x_2 - a_2 & 0 & 0 \\ 0 & 0 & y_1 - b_1 & y_2 - b_2 \end{bmatrix}$$

are linearly independent because $x \neq a$ and $y \neq b$. The Lagrange condition works out to

$$\begin{aligned} (x_1 - y_1, x_2 - y_2, y_1 - x_1, y_2 - x_2) &= \lambda_1(x_1 - a_1, x_2 - a_2, 0, 0) \\ &\quad - \lambda_2(0, 0, y_1 - b_1, y_2 - b_2), \end{aligned}$$

or

$$(x - y, y - x) = \lambda_1(x - a, \mathbf{0}_2) - \lambda_2(\mathbf{0}_2, y - b).$$

The second half of the vector on the left is the additive inverse of the first, so the condition rewrites as

$$x - y = \lambda_1(x - a) = \lambda_2(y - b).$$

If $\lambda_1 = 0$ or $\lambda_2 = 0$ then $x = y$ and both λ_i are 0. Otherwise λ_1 and λ_2 are nonzero. This makes x and y distinct points such that

$$x - y \parallel x - a \parallel y - b,$$

and so the points x , y , a , and b are collinear. Granted, these results are obvious geometrically, but it is pleasing to see them follow so easily from the Lagrange

multiplier condition. On the other hand, not all points x and y such that x , y , a , and b are collinear are solutions to the problem. For example, if both circles are bisected by the x -axis and neither circle sits inside the other, then x and y could be the leftmost points of the circles, neither the closest nor the farthest pair.

The last example of this section begins by maximizing the geometric mean of n positive numbers,

$$f(x_1, \dots, x_n) = (x_1 \cdots x_n)^{1/n},$$

subject to the constraint that their arithmetic mean is 1,

$$\frac{x_1 + \cdots + x_n}{n} = 1.$$

Thus the constraining function is $g(x_1, \dots, x_n) = (x_1 + \cdots + x_n)/n$, and the gradients are

$$\begin{aligned}\nabla f(x_1, \dots, x_n) &= \frac{f(x_1, \dots, x_n)}{n} \left(\frac{1}{x_1}, \dots, \frac{1}{x_n} \right) \\ \nabla g(x_1, \dots, x_n) &= \frac{1}{n} (1, \dots, 1).\end{aligned}$$

The Lagrange condition quickly shows that all x_i are equal, and the constraint forces their value to be 1. Therefore, an extremum of the geometric mean when the arithmetic mean is 1 is the value

$$f(1, \dots, 1) = (1 \cdots 1)^{1/n} = 1.$$

On the other hand, let ε be a small positive number and let

$$x_1 = n - \varepsilon, \quad x_2 = \cdots = x_n = \varepsilon/(n - 1).$$

These x -values also satisfy the constraint that their arithmetic mean is 1, and their geometric mean is close to 0. Therefore the extremum that we found a moment ago is a maximum. That is, $(x_1 \cdots x_n)^{1/n} \leq 1$ for all positive x_1, \dots, x_n such that $(x_1 + \cdots + x_n)/n = 1$. This provides most of the proof of

Theorem 5.4.1 (Arithmetic–Geometric Mean Inequality). *The geometric mean of n positive numbers is at most their arithmetic mean:*

$$(a_1 \cdots a_n)^{1/n} \leq \frac{a_1 + \cdots + a_n}{n} \quad \text{for all positive } a_1, \dots, a_n.$$

Proof. Given positive numbers a_1, \dots, a_n , let $a = (a_1 + \cdots + a_n)/n$ and let $x_i = a_i/a$ for $i = 1, \dots, n$. Then $(x_1 + \cdots + x_n)/n = 1$ and therefore

$$(a_1 \cdots a_n)^{1/n} = a(x_1 \cdots x_n)^{1/n} \leq a = \frac{a_1 + \cdots + a_n}{n}.$$

□

Despite these pleasing examples, Lagrange multipliers are in general no computational panacea. Some problems of optimization with constraint are solved at least as easily by geometry or substitution. Nonetheless, Lagrange's method provides a unifying idea that addresses many different types of optimization problem without reference to geometry or physical considerations. In the following exercises, use whatever methods you find convenient.

Exercises

5.4.1. Find the nearest point to the origin on the intersection of the hyperplanes $x + y + z - 2w = 1$ and $x - y + z + w = 2$ in \mathbf{R}^4 .

5.4.2. Find the nearest point on the ellipse $x^2 + 2y^2 = 1$ to the line $x + y = 4$.

5.4.3. Minimize $f(x, y, z) = z$ subject to the constraints $2x + 4y = 5$, $x^2 + z^2 = 2y$.

5.4.4. Maximize $f(x, y, z) = xy + yz$ subject to the constraints $x^2 + y^2 = 2$, $yz = 2$.

5.4.5. Find the extrema of $f(x, y, z) = xy + z$ subject to the constraints $x \geq 0$, $y \geq 0$, $xz + y = 4$, $yz + x = 5$.

5.4.6. Find the largest rectangular box with sides parallel to the coordinate axes that can be inscribed in the ellipsoid $\left(\frac{x}{a}\right)^2 + \left(\frac{y}{b}\right)^2 + \left(\frac{z}{c}\right)^2 = 1$.

5.4.7. The lengths of the twelve edges of a rectangular block sum to a , and the areas of the six faces sum to $\frac{a^2}{25}$. Find the lengths of the edges when the excess of the block's volume over that of a cube with edge equal to the least edge of the block is greatest.

5.4.8. A cylindrical can (with top and bottom) has volume V . Subject to this constraint, what dimensions give it the least surface area?

5.4.9. Find the distance in the plane from the point $(0, b)$ to the parabola $y = ax^2$ assuming $0 < \frac{1}{2a} < b$.

5.4.10. This exercise extends the Arithmetic–Geometric Mean Inequality. Let e_1, \dots, e_n be positive numbers with $\sum_{i=1}^n e_i = 1$. Maximize the function $f(x_1, \dots, x_n) = x_1^{e_1} \cdots x_n^{e_n}$ (where all $x_i > 0$) subject to the constraint $\sum_{i=1}^n e_i x_i = 1$. Use your result to derive the generalized Arithmetic–Geometric Mean inequality,

$$a_1^{e_1} \cdots a_n^{e_n} \leq e_1 a_1 + \cdots + e_n a_n \quad \text{for all positive } a_1, \dots, a_n.$$

What values of e_1, \dots, e_n reduce this to the first Arithmetic–Geometric Mean Inequality?

5.4.11. Let p and q be positive numbers satisfying the equation $\frac{1}{p} + \frac{1}{q} = 1$. Maximize the function of $2n$ variables $f(x_1, \dots, x_n, y_1, \dots, y_n) = \sum_{i=1}^n x_i y_i$ subject to the constraints $\sum_{i=1}^n x_i^p = 1$ and $\sum_{i=1}^n y_i^q = 1$. Derive **Hölder's Inequality**: For all nonnegative $a_1, \dots, a_n, b_1, \dots, b_n$,

$$\sum_{i=1}^n a_i b_i \leq \left(\sum_{i=1}^n a_i^p \right)^{1/p} \left(\sum_{i=1}^n b_i^q \right)^{1/q}.$$

5.5 Lagrange Multipliers: Analytic Proof and General Examples

Recall that the environment for optimization with constraints consists of

- an open set $A \subset \mathbf{R}^n$,
- a constraining \mathcal{C}^1 -mapping $g : A \rightarrow \mathbf{R}^c$,
- the corresponding level set $L = \{v \in A : g(v) = \mathbf{0}_c\}$,
- and a \mathcal{C}^1 -function $f : A \rightarrow \mathbf{R}$ to optimize on L .

We have argued geometrically, and not fully rigorously, that if f on L is optimized at a point $p \in L$ then the gradient $f'(p)$ is orthogonal to L at p . Also, every linear combination of the gradients of the component functions of g is orthogonal to L at p . We want to assert the converse, that every vector that is orthogonal to L at p is such a linear combination. The desired converse assertion does not always hold, but when it does it gives the Lagrange condition,

$$\nabla f(p) = \sum_{i=1}^c \lambda_i \nabla g_i(p).$$

Here is the rigorous analytic justification that the Lagrange multiplier method usually works. The Implicit Function Theorem will do the heavy lifting, and it will reaffirm that the method is guaranteed only where the gradients of the component functions of g are linearly independent.

Theorem 5.5.1 (Lagrange Multiplier Condition). *Let n and c be positive integers with $n > c$. Let $g : A \rightarrow \mathbf{R}^c$ (where $A \subset \mathbf{R}^n$) be a mapping that is continuously differentiable at each interior point of A . Consider the level set*

$$L = \{x \in A : g(x) = \mathbf{0}_c\}.$$

Let $f : A \rightarrow \mathbf{R}$ be a function. Suppose that the restriction of f to L has an extreme value at a point $p \in L$ that is an interior point of A . Suppose that f is differentiable at p , and suppose that the c -by- n derivative matrix $g'(p)$ contains a c -by- c block that is invertible. Then the following conditions hold:

$$\begin{aligned} \nabla f(p) &= \lambda g'(p) \quad \text{for some row vector } \lambda \in \mathbf{R}^c, \\ g(p) &= \mathbf{0}_c. \end{aligned}$$

Proof. The second condition holds since p is a point in L . The first condition needs to be proved. Let $r = n - c$, the number of variables that should remain free under the constraint $g(x) = \mathbf{0}_c$, and notate the point in p as $p = (a, b)$, where $a \in \mathbf{R}^r$ and $b \in \mathbf{R}^c$. Using this notation, we have $g(a, b) = \mathbf{0}_c$ and $g'(a, b) = [M \ N]$ where M is c -by- r and N is c -by- c and invertible. (We may assume that N is the invertible block in the hypotheses to the theorem because we may freely permute the variables.)

The Implicit Function Theorem gives a mapping $\varphi : A_0 \rightarrow \mathbf{R}^c$ (where $A_0 \subset \mathbf{R}^r$ and a is an interior point of A_0) with $\varphi(a) = b$, $\varphi'(a) = -N^{-1}M$, and for all points $(x, y) \in A$ near (a, b) , $g(x, y) = \mathbf{0}_c$ if and only if $y = \varphi(x)$. Make f depend only on the free variables by defining

$$\tilde{f} : A_0 \rightarrow \mathbf{R}^c \quad \text{by} \quad \tilde{f}(x) = f(x, \varphi(x)).$$

(See figure 5.23.) Since the domain of \tilde{f} doesn't curve around in some larger space, the optimization techniques from chapter 4 apply, i.e., \tilde{f} has a critical point at a . Since \tilde{f} is a composition, the Chain Rule shows that the condition $\nabla \tilde{f}(a) = 0$ is

$$\nabla f(a, \varphi(a)) \begin{bmatrix} I_r \\ \varphi'(a) \end{bmatrix} = 0,$$

and since $\varphi(a) = b$ and $\varphi'(a) = -N^{-1}M$, this condition is

$$\nabla f(a, b) \begin{bmatrix} I_r \\ -N^{-1}M \end{bmatrix} = 0.$$

Let $\nabla f(a, b) = (u, v)$ where $u \in \mathbf{R}^r$ and $v \in \mathbf{R}^c$ are row vectors. Then the previous display becomes

$$u = vN^{-1}M,$$

which rearranges as

$$[u \ v] = vN^{-1} [M \ N].$$

Set $\lambda = vN^{-1} \in \mathbf{R}^c$ and we have the desired condition,

$$\nabla f(a, b) = \lambda g'(a, b).$$

□

We have seen that the Lagrange Multiplier Condition is necessary but not sufficient for an extreme value. That is, it can report a false positive, as in the two-circle problem in the previous section. This is not a serious problem since inspecting all the points that meet the Lagrange condition will determine which of them give the true extrema of f . A false negative would be a worse situation, giving us no indication that an extreme value might exist, much less how to find it. The following example shows that the false negative scenario can arise without the invertible c -by- c block required in Theorem 5.5.1.

Let the temperature in the plane be given by




Figure 5.23. The Lagrange Multiplier Criterion from the Implicit Function Theorem

$$f(x, y) = x,$$

and consider a plane set defined by one constraint on two variables,

$$L = \{(x, y) \in \mathbf{R}^2 : y^2 = x^3\}.$$

(See figure 5.24.) Since temperature increases as we move to the right, the coldest point of L is its leftmost point, the cusp at $(0, 0)$. However, the Lagrange condition does not find this point. Indeed, the constraining function is $g(x, y) = x^3 - y^2$ (which does have continuous derivatives, notwithstanding that its level set has a cusp: the graph of a smooth function is smooth, but the level set of a smooth function need not be smooth—this is exactly the issue addressed by the Implicit Function Theorem). Therefore the Lagrange condition and the constraint are

$$\begin{aligned} (1, 0) &= \lambda(3x^2, -2y), \\ x^3 &= y^2. \end{aligned}$$

These equations have no solution. The problem is that the gradient at the cusp is $\nabla g(0, 0) = (0, 0)$, and neither of its 1-by-1 subblocks is invertible. In general, the Lagrange Multiplier Condition will not report a false negative so long as we remember that it only claims to check for extrema at the *nonsingular* points of L , the points p such that $g'(p)$ has an invertible c -by- c subblock.

The previous section gave specific examples of the Lagrange multiplier method. This section now gives some general families of examples.

Recall that the previous section discussed the problem of optimizing the distance between two points in the plane, each point lying on an associated circle. Now, as the first general example of the Lagrange multiplier method, let $(x, y) \in \mathbf{R}^n \times \mathbf{R}^n$ denote a pair of points each from \mathbf{R}^n , and let the function f measure the square of distance between such a pair,




Figure 5.24. Curve with cusp

$$f : \mathbf{R}^n \times \mathbf{R}^n \longrightarrow \mathbf{R}, \quad f(x, y) = |x - y|^2.$$

Note that $\nabla f(x, y) = [x - y \ y - x]$, viewing x and y as row vectors. Given two mappings $g_1 : \mathbf{R}^n \longrightarrow \mathbf{R}^{c_1}$ and $g_2 : \mathbf{R}^n \longrightarrow \mathbf{R}^{c_2}$, define

$$g : \mathbf{R}^n \times \mathbf{R}^n \longrightarrow \mathbf{R}^{c_1+c_2}, \quad g(x, y) = (g_1(x), g_2(y)).$$

To optimize the function f subject to the constraint $g(x, y) = (\mathbf{0}_{c_1}, \mathbf{0}_{c_2})$ is to optimize the distance between pairs of points x and y on the respective level sets cut out of \mathbf{R}^n by the c_1 conditions $g_1(x) = \mathbf{0}_{c_1}$ and the c_2 conditions $g_2(y) = \mathbf{0}_{c_2}$. Assuming that the Lagrange condition holds for the optimizing pair, it is

$$\begin{aligned} [x - y \ y - x] &= \lambda g'(x, y) = [\lambda_1 \ -\lambda_2] \begin{bmatrix} g'_1(x) & \mathbf{0}_{c_2 \times n} \\ \mathbf{0}_{c_1 \times n} & g'_2(y) \end{bmatrix} \\ &= \lambda_1(g'_1(x), \mathbf{0}_{c_2 \times n}) - \lambda_2(\mathbf{0}_{c_1 \times n}, g'_2(y)), \end{aligned}$$

where $\lambda_1 \in \mathbf{R}^{c_1}$ and $\lambda_2 \in \mathbf{R}^{c_2}$ are row vectors. The symmetry of ∇f reduces this equality of $2n$ -vectors to an equality of n -vectors,

$$x - y = \lambda_1 g'_1(x) = \lambda_2 g'_2(y).$$

That is, either $x = y$ or the line through x and y is normal to the first level set at x and normal to the second level set at y , generalizing the result from the two-circle problem. With this result in mind, you may want to revisit exercise 0.0.1 from the preface to these notes.

The second general example is the problem of optimizing a linear function subject to affine constraints. The data are (viewing vectors as columns)

$$\begin{aligned} f : \mathbf{R}^n &\longrightarrow \mathbf{R}, & f(x) &= a^t x, \\ g : \mathbf{R}^n &\longrightarrow \mathbf{R}^c, & g(x) &= Mx \text{ where } M \in M_{c,n}(\mathbf{R}). \end{aligned}$$

The Lagrange condition and the constraints are

$$\begin{aligned} a^t &= \lambda^t M, \\ Mx &= b. \end{aligned}$$

The Lagrange condition is usually not solvable for λ since we are tacitly assuming that $c < n$, i.e., the condition is an inhomogeneous linear system with more equations than variables. But when the condition is solvable, it immediately gives the optimal value of f ,

$$f(x) = a^t x = \lambda^t Mx = \lambda^t b.$$

Note that we don't need to find an optimizing x in order to find the optimal value $f(x)$. In fact, when the Lagrange condition is solvable, the function f subject to the constraint $g = b$ is constant. Indeed, the condition says that a^t is a linear combination of the rows of M , making a orthogonal to all vectors \tilde{x} such that $M\tilde{x} = \mathbf{0}_c$. A result from linear algebra says that the vectors x satisfying the constraint $Mx = b$ are the vectors

$$x = x_0 + \tilde{x},$$

where x_0 is a fixed linear combination of the rows of M such that $Mx_0 = b$, and \tilde{x} is any vector such that $M\tilde{x} = \mathbf{0}_c$. Thus for all such x ,

$$f(x) = a^t(x_0 + \tilde{x}) = \lambda^t M(x_0 + \tilde{x}) = \lambda^t b.$$

That is, f subject to the constraint $g = b$ is constant, as claimed.

For geometric insight into what the calculation is telling us, envision the space of vectors \tilde{x} such that $M\tilde{x} = \mathbf{0}_c$ as an axis, and envision the space of linear combinations of the rows of M as a plane orthogonal to the axis. Then the first condition is that a lies in the plane, and the second is that x lies on an axis parallel to the \tilde{x} -axis. The constant value of f is $a^t x$ for any x on the axis. In particular, the value is $a^t x_0$ where x_0 is the point where the axis meets the plane.

The third general example is to optimize a quadratic function subject to affine constraints. Here the data are

$$\begin{aligned} f : \mathbf{R}^n &\longrightarrow \mathbf{R}, & f(x) &= x^t A x \text{ where } A \in M_n(\mathbf{R}) \text{ is symmetric,} \\ g : \mathbf{R}^n &\longrightarrow \mathbf{R}^c, & g(x) &= Mx \text{ where } M \in M_{c,n}(\mathbf{R}). \end{aligned}$$

To set up the Lagrange condition, we need to differentiate the quadratic function f . Compute that

$$f(x+h) - f(x) = (x+h)^t A(x+h) - x^t A x = 2x^t A h + h^t A h,$$

and so the best linear approximation of this difference is $T(h) = 2x^t A h$. It follows that

$$\nabla f(x) = 2x^t A.$$

Returning to the optimization problem, the Lagrange condition and the constraints are

$$\begin{aligned}x^t A &= \lambda^t M, \\ Mx &= b.\end{aligned}$$

So the possible optimal values of f take the form

$$f(x) = x^t Ax = \lambda^t Mx = \lambda^t b,$$

which we will know as soon as we find the possible values of λ , without needing to find x . Assume that A is invertible. Transpose the Lagrange condition to get $Ax = M^t \lambda$, whence $x = A^{-1} M^t \lambda$ and thus $b = Mx = MA^{-1} M^t \lambda$, so that (also assuming that the c -by- c matrix $MA^{-1} M^t$ is invertible) $\lambda = (MA^{-1} M^t)^{-1} b$. That is, the optimal value $\lambda^t b$ of f is

$$f(x) = b^t (MA^{-1} M^t)^{-1} b.$$

Also, the x -value at which f is optimized is

$$x = A^{-1} M^t (MA^{-1} M^t)^{-1} b.$$

In particular, letting $A = I$, the closest point x to the origin such that $Mx = b$ is

$$x = M^t (MM^t)^{-1} b,$$

and its distance-squared from the origin is

$$|x|^2 = b^t (MM^t)^{-1} b.$$

The fourth general example is to optimize a linear function subject to a quadratic constraint. The data are

$$\begin{aligned}f : \mathbf{R}^n &\longrightarrow \mathbf{R}, & f(x) &= a^t x, \\ g : \mathbf{R}^n &\longrightarrow \mathbf{R}, & g(x) &= x^t Mx \text{ where } M \in M_n(\mathbf{R}) \text{ is symmetric.}\end{aligned}$$

The Lagrange condition and the constraint are

$$\begin{aligned}a^t &= \lambda x^t M, \\ x^t Mx &= b.\end{aligned}$$

Therefore the optimized value of f is

$$f(x) = a^t x = \lambda x^t Mx = \lambda b,$$

and so to find this value it suffices to find λ . Assuming that M is invertible, the Lagrange condition gives $x^t = \lambda^{-1} a^t M^{-1}$, hence $x = \lambda^{-1} M^{-1} a$. Substitute for x^t and x in the constraint to get $\lambda^{-2} a^t M^{-1} a = b$, and thus (assuming that $a^t M^{-1} a / b > 0$)

$$\lambda = \pm \sqrt{a^t M^{-1} a / b}.$$

The final general example is to optimize a quadratic function subject to a quadratic constraint. The data are

$$\begin{aligned} f : \mathbf{R}^n &\longrightarrow \mathbf{R}, & f(x) &= x^t A x \text{ where } A \in M_n(\mathbf{R}) \text{ is symmetric,} \\ g : \mathbf{R}^n &\longrightarrow \mathbf{R}, & g(x) &= x^t M x \text{ where } M \in M_n(\mathbf{R}) \text{ is symmetric.} \end{aligned}$$

The Lagrange condition and the constraint are

$$\begin{aligned} x^t A &= \lambda x^t M, \\ x^t M x &= b. \end{aligned}$$

Note that as a special case, if $b = 0$ then these conditions hold for $x = \mathbf{0}$, giving $f(\mathbf{0}) = 0$ as a candidate optimal value of f . From now on, we assume that $x \neq \mathbf{0}$, a necessary assumption when $b \neq 0$. By the Lagrange condition and the constraint, the possible optimal values of f take the form

$$f(x) = x^t A x = \lambda x^t M x = \lambda b,$$

which we will know as soon as we find the possible values of λ , without needing to find x . Assuming that M is invertible, the Lagrange condition gives

$$M^{-1} A x = \lambda x.$$

In other words, x must satisfy the condition that *multiplying x by $M^{-1}A$ gives a scalar multiple of x* . Any nonzero vector x that satisfies this condition is called an *eigenvector* of $M^{-1}A$. The scalar multiple factor λ is the corresponding *eigenvalue*.

The eigenvalues of any square matrix B are found by a systematic procedure. The first step is to observe that the condition $Bx = \lambda x$ is

$$(B - \lambda I)x = \mathbf{0}.$$

Since any eigenvector x is nonzero by definition, $B - \lambda I$ is not invertible, i.e.,

$$\det(B - \lambda I) = 0.$$

Conversely, for every $\lambda \in \mathbf{R}$ satisfying this equation there is at least one eigenvector x of B because the equation $(B - \lambda I)x = \mathbf{0}$ has nonzero solutions. And so the eigenvalues are the real roots of the polynomial

$$p_B(\lambda) = \det(B - \lambda I).$$

This is the *characteristic polynomial* of B , already discussed in exercise 4.5.10. For example, part (a) of that exercise covered the case $n = 2$, showing that if $B = \begin{bmatrix} a & b \\ b & d \end{bmatrix}$ then

$$p_B(\lambda) = \lambda^2 - (a + d)\lambda + (ad - b^2).$$

The discriminant of this quadratic polynomial is

$$\Delta = (a + d)^2 - 4(ad - b^2) = (a - d)^2 + 4b^2.$$

Since Δ is nonnegative, all roots of characteristic polynomial are real. And a result of linear algebra says that for any positive n , all roots of the characteristic polynomial of a symmetric n -by- n matrix B are real as well. However, returning to our example, even though the square matrices A and M are assumed to be symmetric, the product $M^{-1}A$ need not be.

As a particular case of this last general example, if $A = I$ then finding the eigenvectors of M encompasses finding the points of a quadric surface that are closest to the origin or farthest from the origin. For instance, if $n = 2$ and $M = \begin{bmatrix} a & b \\ b & d \end{bmatrix}$ then we are optimizing on the set of points $(x_1, x_2) \in \mathbf{R}^2$ such that, say,

$$ax_1^2 + 2bx_1x_2 + dx_2^2 = 1.$$

This is the equation of a conic section. When $b = 0$ we have an unrotated ellipse or hyperbola, and the only possible optimal points will be the scalar multiples of e_1 and e_2 that lie on the curve. For an ellipse, a pair of points on one axis is closest to the origin, and a pair on the other axis is farthest; for a hyperbola, a pair on one axis is closest and there are no points on the other axis. In the case of a circle, the matrix M is a scalar multiple of the identity matrix, and so all vectors are eigenvectors compatibly with the geometry that all points are equidistant from the origin. Similarly if $n = 3$ then L is a surface such as an ellipsoid or a hyperboloid.

Exercises

5.5.1. Let $f(x, y) = y$ and let $g(x, y) = y^3 - x^4$. Graph the level set $L = \{(x, y) : g(x, y) = 0\}$. Show that the Lagrange multiplier criterion does not find any candidate points where f is optimized on L . Optimize f on L nonetheless.

5.5.2. Consider the linear mapping

$$g(x, y, z) = (x + 2y + 3z, 4x + 5y + 6z).$$

(a) Use a general method given in this section to optimize the linear function $f(x, y, z) = 6x + 9y + 12z$ subject to the affine constraints $g(x, y, z) = (7, 8)$.

(b) Verify without using the Lagrange multiplier method that the function f subject to the constraints $g = (7, 8)$ (with f and g from part (a)) is constant, always taking the value that you found in part (a).

(c) Show that the function $f(x, y, z) = 5x + 7y + z$ can not be optimized subject to any constraint $g(x, y, z) = b$.

5.5.3. (a) Use a general method given in this section to minimize the quadratic function $f(x, y) = x^2 + y^2$ subject to the affine constraint $3x + 5y = 8$.

(b) Use the method to find the extrema of $f(x, y, z) = 2xy + z^2$ subject to the constraints $x + y + z = 1$, $x + y - z = 0$.

(c) Use the method to find the nearest point to the origin on the intersection of the hyperplanes $x + y + z - 2w = 1$ and $x - y + z + w = 2$ in \mathbf{R}^4 , reproducing your answer to exercise 5.4.1.

5.5.4. (a) Use a general method to maximize $f(x, y, z) = x - 2y + 2z$ on the sphere of radius 3.

(b) Use the method to optimize the function $f(x, y, z, w) = x + y - z - w$ subject to the constraint $g(x, y, z, w) = 1$, $g(x, y, z, w) = x^2/2 - y^2 + z^2 - w^2$.

5.5.5. (a) Use a general method to optimize the function $f(x, y) = 2xy$ subject to the constraint $g(x, y) = 1$ where $g(x, y) = x^2 + 2y^2$.

(b) Use the method to optimize the function $f(x, y, z) = 2(xy + yz + zx)$ subject to the constraint $g(x, y, z) = 1$ where $g(x, y, z) = x^2 + y^2 - z^2$.

5.6 Summary

The Inverse Function Theorem implies the Implicit Function Theorem, and conversely. The Implicit Function Theorem implies the Lagrange multiplier condition, a systematic approach to problems of optimization with constraints. The Lagrange multiplier method can also be viewed as optimization on a level set inside a larger-dimensional space.

Integration

The integral of a scalar-valued function of many variables, taken over a box of its inputs, is defined in sections 6.1 and 6.2. Intuitively, the integral can be understood as representing mass or volume, but the definition is purely mathematical: the integral is a limit of sums, as in one-variable calculus. Multivariable integration has many familiar properties—for example, the integral of a sum is the sum of the integrals. Section 6.3 shows that continuous functions can be integrated over boxes. However, we want to carry out multivariable integration over more generally-shaped regions. That is, the theory has geometric aspects not present in the one-dimensional case, where integration is carried out over intervals. After a quick review of the one-variable theory in section 6.4, section 6.5 shows that continuous functions can also be integrated over nonboxes that have manageable shapes. The main tools for evaluating multivariable integrals are Fubini’s Theorem (section 6.6), which reduces an n -dimensional integral to an n -fold nesting of one-dimensional integrals, and the Change of Variable Theorem (section 6.7), which replaces one multivariable integral by another that may be easier to evaluate. Section 6.8 provides some preliminaries for the proof of the Change of Variable Theorem, and then section 6.9 gives the proof.

6.1 Machinery: Boxes, Partitions, and Sums

The integral represents fairly clear ideas, but defining it properly requires some care. Here is some terminology that is standard from the calculus of one variable, perhaps other than *compact* (meaning *closed and bounded*) from section 2.4 of these notes.

Definition 6.1.1 (Compact Interval, Length, Partition, Subinterval).
A nonempty compact interval in \mathbf{R} is a set

$$I = [a, b] = \{x \in \mathbf{R} : a \leq x \leq b\},$$

where a and b are real numbers with $a \leq b$. The **length** of the interval is

$$\text{length}(I) = b - a.$$

A **partition** of I is a set of real numbers

$$P = \{t_0, t_1, \dots, t_k\}$$

satisfying

$$a = t_0 < t_1 < \dots < t_k = b.$$

Such a partition divides I into k **subintervals** J_1, \dots, J_k where

$$J_j = [t_{j-1}, t_j], \quad j = 1, \dots, k.$$

A generic nonempty compact subinterval of I is denoted J . (See figure 6.1.) Since the only intervals that we are interested in are nonempty and compact, either or both of these properties will often be tacit from now on, rather than stated again and again. As a special case, Definition 6.1.1 says that any length-zero interval $[a, a]$ has only one partition, $P = \{a\}$, which divides it into no subintervals.




Figure 6.1. Interval and subintervals

Definition 6.1.2 (Bounded Function). Let A be a subset of \mathbf{R} , and let $f : A \rightarrow \mathbf{R}$ be a function. Then f is **bounded** if its range, $\{f(x) : x \in A\}$, is bounded as a set in \mathbf{R} , as in Definition 2.4.6. That is, f is bounded if there exists some $R > 0$ such that $|f(x)| < R$ for all $x \in A$.

Visually, a function is bounded if its graph is contained inside a horizontal strip. On the other hand, the graph of a bounded function needn't be contained in a vertical strip. This is because the domain (and therefore the graph) need not be bounded. For example, these functions are bounded:

$$f(x) = \sin x, \quad f(x) = 1/(1 + x^2), \quad f(x) = \arctan x,$$

but these functions are not:

$$f(x) = e^x, \quad f(x) = 1/x \text{ for } x \neq 0.$$

Definition 6.1.3 (One-dimensional Lower Sum and Upper Sum). Let I be a nonempty compact interval in \mathbf{R} , and let $f : I \rightarrow \mathbf{R}$ be a bounded function. For any subinterval J of I , the greatest lower bound of the values taken by f on J is denoted $m_J(f)$,

$$m_J(f) = \inf \{f(x) : x \in J\},$$

and similarly the least upper bound is denoted $M_J(f)$,

$$M_J(f) = \sup \{f(x) : x \in J\}.$$

The **lower sum** of f over P is

$$L(f, P) = \sum_J m_J(f) \text{length}(J),$$

and the **upper sum** of f over P is

$$U(f, P) = \sum_J M_J(f) \text{length}(J).$$

If the interval I in Definition 6.1.3 has length zero, then the lower and upper sums are empty and so they are assigned the value 0 by convention.

The function f in Definition 6.1.3 is not required to be differentiable or even continuous, only bounded. Even so, the values $m_J(f)$ and $M_J(f)$ in the previous definition exist by the set-bound phrasing of the principle that the real number system is complete. To review this idea, see Theorem 1.1.5. When f is in fact continuous, the Extreme Value Theorem (Theorem 2.4.15) justifies substituting *min* and *max* for *inf* and *sup* in the definitions of $m_J(f)$ and $M_J(f)$, since each subinterval J is nonempty and compact. It may be easiest at first to understand $m_J(f)$ and $M_J(f)$ by imagining f to be continuous and mentally substituting appropriately. But we will need to integrate discontinuous functions f . Such functions may take no minimum or maximum on J , and so we may run into a situation like the one pictured in figure 6.2, in which the values $m_J(f)$ and $M_J(f)$ are not actual outputs of f . Thus the definition must be as given to make sense.

The technical properties of *inf* and *sup* will figure in Lemmas 6.1.6, 6.1.8, and 6.2.2. To see them in isolation first, we rehearse them now. So, let S and T be nonempty sets of real numbers, both bounded. In the context of integration, S and T will be sets of outputs of a bounded function f . This is irrelevant for the moment, but it may help you to see later how these ideas are used in context if you now imagine S and T on a vertical axis, as in figure 6.2, rather than on a horizontal one. In any case, the necessary results are as follows.

- $\inf(S) \leq \sup(S)$. In fact *any* lower bound of S is at most as big as *any* upper bound, because any element of S lies between them. In particular, this argument applies to the greatest lower bound $\inf(S)$ and the least upper bound $\sup(S)$, giving the stated inequality.




Figure 6.2. Sup and inf but no max or min

- If $S \subset T$ then $\inf(T) \leq \inf(S) \leq \sup(S) \leq \sup(T)$. We already have the middle inequality. To establish the others, the idea is that since S is a subset, the bounds on S are innately at least as tight as those on T . More specifically, since $\inf(T)$ is a lower bound of T , it is a lower bound of the subset S , and since $\inf(S)$ is the *greatest* lower bound of S the first inequality follows. The third inequality is similar.
- If $s \leq t$ for all $s \in S$ and $t \in T$ then $\sup(S) \leq \inf(T)$. Imprecisely, the idea is that S is entirely below T on the vertical axis, and so the smallest number that traps S from above is still below the largest number that traps T from below. This will be proved more carefully in the next section.

Graphing f over I in the usual fashion and interpreting the lower and upper sum as sums of rectangle-areas shows that they are respectively too small and too big to be the area under the graph. (See figure 6.3.) Alternatively, thinking of f as the density function of a wire stretched over the interval I shows that the lower and upper sum are too small and too big to be the mass of the wire. The hope is that the lower and upper sums are trapping a yet-unknown quantity (possibly to be imagined as area or mass) from each side, and that as the partition P becomes finer, the lower and upper sums will actually converge to this value.




Figure 6.3. Too small and too big

All the terminology so far generalizes easily from one dimension to many, i.e., from \mathbf{R} to \mathbf{R}^n . Recall that if S_1, S_2, \dots, S_n are subsets of \mathbf{R} then their **cartesian product** is a subset of \mathbf{R}^n ,

$$S_1 \times S_2 \times \cdots \times S_n = \{(s_1, s_2, \dots, s_n) : s_1 \in S_1, s_2 \in S_2, \dots, s_n \in S_n\}.$$

(See figure 6.4, in which $n = 2$, and S_1 has two components, and S_2 has one component, so that the cartesian product $S_1 \times S_2$ has two components.)




Figure 6.4. Cartesian product

Definition 6.1.4 (Compact Box, Box Volume, Partition, Subbox). A nonempty compact box in \mathbf{R}^n is a cartesian product

$$B = I_1 \times I_2 \times \cdots \times I_n$$

of nonempty compact intervals I_j for $j = 1, \dots, n$. The **volume** of the box is the product of the lengths of its sides,

$$\text{vol}(B) = \prod_{j=1}^n \text{length}(I_j).$$

A **partition** of B is a cartesian product of partitions P_j of I_j for $j = 1, \dots, n$,

$$P = P_1 \times P_2 \times \cdots \times P_n.$$

Such a partition divides B into **subboxes** J , each such subbox being a cartesian product of subintervals. By a slight abuse of language, these are called the **subboxes of** P .

(See figure 6.5, and imagine its three-dimensional Rubik's cube counterpart.) Every nonempty compact box in \mathbf{R}^n has partitions, even such boxes with

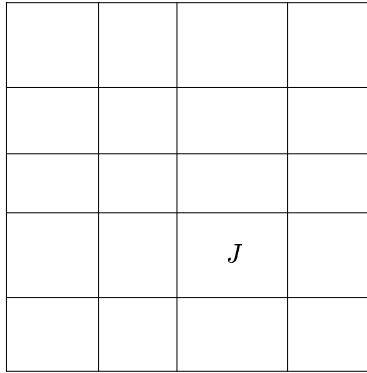


Figure 6.5. Box and subboxes

some length-zero sides. This point will arise at the very beginning of the next section.

The definition of a bounded function $f : A \rightarrow \mathbf{R}$, where now A is a subset of \mathbf{R}^n , is virtually the same as earlier in the section: again the criterion is that its range must be bounded as a set. (In fact the definition extends just as easily to mappings $f : A \rightarrow \mathbf{R}^m$, but we need only scalar-valued functions here.)

Definition 6.1.5 (n -dimensional Lower Sum and Upper Sum). *Let B be a nonempty compact box in \mathbf{R}^n , and let $f : B \rightarrow \mathbf{R}$ be a bounded function. For any subbox J of B , define $m_J(f)$ and $M_J(f)$ analogously to before,*

$$m_J(f) = \inf \{f(x) : x \in J\} \quad \text{and} \quad M_J(f) = \sup \{f(x) : x \in J\}.$$

The lower sum and upper sum of f over P are similarly

$$L(f, P) = \sum_J m_J(f) \text{vol}(J) \quad \text{and} \quad U(f, P) = \sum_J M_J(f) \text{vol}(J).$$

With minor grammatical modifications, this terminology includes the previous definition as a special case when $n = 1$ (e.g., *volume* reverts to *length*, as it should revert to *area* when $n = 2$), so from now on we work in \mathbf{R}^n . However, keeping the cases $n = 1$ and $n = 2$ in mind should help to make the pan-dimensional ideas of multivariable integration geometrically intuitive. If the box B in Definition 6.1.5 has any sides of length zero then the upper and lower sums are 0.

Graphing f over B in the usual fashion when $n = 2$ and interpreting the lower and upper sum as sums of box-volumes shows that they are respectively too small and too big to be the volume under the graph. (See figure 6.6.) Alternatively, if $n = 2$ or $n = 3$, then thinking of f as the density of a plate

or a block occupying the box B shows that the lower and upper sum are too small and too big to be the object's mass. Again, the hope is that as the partitions become finer, the lower and upper sums will converge to a common value that they are trapping from either side.




Figure 6.6. Too small and too big

The first result supports this intuition.

Lemma 6.1.6. *For any box B , any partition P of B , and any bounded function $f : B \rightarrow \mathbf{R}$,*

$$L(f, P) \leq U(f, P).$$

Proof. For any subbox J of P , $m_J(f) \leq M_J(f)$ by definition, while also $\text{vol}(J) \geq 0$, and therefore $m_J(f) \text{vol}(J) \leq M_J(f) \text{vol}(J)$. Sum this relation over all subboxes J to get the result. \square

The next thing to do is express the notion of taking a finer partition.

Definition 6.1.7 (Refinement). *Let P and P' be partitions of B . Then P' is a refinement of P if $P' \supset P$.*

Figure 6.7 illustrates the fact that if P' refines P then every subbox of P' is contained in a subbox of P . The literal manifestation in the figure of the containment $P' \supset P$ is that the set of points where a horizontal line segment and a vertical line segment meet in the right side of the figure subsumes the set of such points in the left side.

Refining a partition brings the lower and upper sums nearer each other:

Lemma 6.1.8. *Suppose that P' refines P as a partition of the box B . Then*

$$L(f, P) \leq L(f, P') \quad \text{and} \quad U(f, P') \leq U(f, P).$$

See figure 6.8 for a picture-proof for lower sums when $n = 1$, thinking of the sums in terms of area. The formal proof is just a symbolic rendition of the figure's features.




Figure 6.7. Refinement




Figure 6.8. Lower sum increasing under refinement

Proof. Every subbox J of P divides further under the refinement P' into subboxes J' . Since each $J' \subset J$, f has less opportunity to be small on J' than on J , and so $m_{J'}(f) \geq m_J(f)$. Thus

$$\begin{aligned} \sum_{J' \subset J} m_{J'}(f) \operatorname{vol}(J') &\geq \sum_{J' \subset J} m_J(f) \operatorname{vol}(J') \\ &= m_J(f) \sum_{J' \subset J} \operatorname{vol}(J') = m_J(f) \operatorname{vol}(J). \end{aligned}$$

Sum the relation $\sum_{J' \subset J} m_{J'}(f) \operatorname{vol}(J') \geq m_J(f) \operatorname{vol}(J)$ over all subboxes J of P to get $L(f, P') \geq L(f, P)$. The argument is similar for upper sums. \square

The proof uncritically uses the fact that the volumes of a box's subboxes sum to the volume of the box. This is true, and left as an exercise. The emphasis here isn't on boxes (which are straightforward), but on defining the integral of a function f whose domain is a box. The next result helps investigate whether the lower and upper sums indeed trap some value from both sides. First we need a definition.

Definition 6.1.9 (Common Refinement). Given two partitions of B ,

$$P = P_1 \times P_2 \times \cdots \times P_n \quad \text{and} \quad P' = P'_1 \times P'_2 \times \cdots \times P'_n,$$

their **common refinement** is the partition

$$P'' = (P_1 \cup P'_1) \times (P_2 \cup P'_2) \times \cdots \times (P_n \cup P'_n).$$

(See figure 6.9.) The common refinement of two partitions P and P' is certainly a partition that refines both P and P' , and it is the smallest such partition. The union $P \cup P'$ is not taken as the definition of the common refinement because it need not be a partition at all. The common refinement does all the work for the next result.




Figure 6.9. Common refinement

Proposition 6.1.10 (Lower Sums Are At Most Upper Sums). Let P and P' be partitions of the box B , and let $f : B \rightarrow \mathbf{R}$ be any bounded function. Then

$$L(f, P) \leq U(f, P').$$

Proof. Let P'' be the common refinement of P and P' . By the two lemmas,

$$L(f, P) \leq L(f, P'') \leq U(f, P'') \leq U(f, P'),$$

proving the result. \square

Exercises

6.1.1. (a) Let $I = [0, 1]$, let $P = \{0, 1/2, 1\}$, let $P' = \{0, 3/8, 5/8, 1\}$, and let P'' be the common refinement of P and P' . What are the subintervals of P , and what are their lengths? Same question for P' . Same question for P'' .

(b) Let $B = I \times I$, let $Q = P \times \{0, 1/2, 1\}$, let $Q' = P' \times \{0, 1/2, 1\}$, and let Q'' be the common refinement of Q and Q' . What are the subboxes of Q and what are their areas? Same question for Q' . Same question for Q'' .

6.1.2. Show that the lengths of the subintervals of any partition of $[a, b]$ sum to the length of $[a, b]$. Same for the areas of the subboxes of $[a, b] \times [c, d]$. Generalize to \mathbf{R}^n .

6.1.3. Let $J = [0, 1]$. Compute $m_J(f)$ and $M_J(f)$ for each of the following functions $f : J \rightarrow \mathbf{R}$.

(a) $f(x) = x(1 - x)$,

(b) $f(x) = \begin{cases} 1 & \text{if } x \text{ is irrational} \\ 1/m & \text{if } x = n/m \text{ in lowest terms, } n, m \in \mathbf{Z} \text{ and } m > 0, \end{cases}$

(c) $f(x) = \begin{cases} (1 - x) \sin(1/x) & \text{if } x \neq 0 \\ 0 & \text{if } x = 0. \end{cases}$

6.1.4. (a) Let I, P, P' and P'' be as in exercise 6.1.1(a), and let $f(x) = x^2$ on I . Compute the lower sums $L(f, P), L(f, P'), L(f, P'')$ and the corresponding upper sums, and check that they conform to Lemma 6.1.6, Lemma 6.1.8, and Proposition 6.1.10.

(b) Let B, Q, Q' and Q'' be as in exercise 6.1.1(b), and define $f : B \rightarrow \mathbf{R}$ by

$$f(x, y) = \begin{cases} 0 & \text{if } 0 \leq x < 1/2 \\ 1 & \text{if } 1/2 \leq x \leq 1. \end{cases}$$

Compute $L(f, Q), L(f, Q'), L(f, Q'')$ and the corresponding upper sums, and check that they conform to Lemma 6.1.6, Lemma 6.1.8, and Proposition 6.1.10.

6.1.5. Draw the cartesian product $([a_1, b_1] \cup [c_1, d_1]) \times ([a_2, b_2] \cup [c_2, d_2]) \subset \mathbf{R}^2$ where $a_1 < b_1 < c_1 < d_1$ and similarly for the other subscript.

6.1.6. When is a cartesian product empty?

6.1.7. Show that the union of partitions of a box B need not be a partition of B .

6.1.8. Draw a picture illustrating the proof of Proposition 6.1.10 when $n = 1$.

6.2 Definition of the Integral

Fix a nonempty compact box B and a bounded function $f : B \rightarrow \mathbf{R}$. The set of lower sums of f over all partitions P of B ,

$$\{L(f, P) : P \text{ is a partition of } B\},$$

is nonempty because such partitions exist (as observed in the previous section), and similarly for the set of upper sums. Proposition 6.1.10 shows that the set of lower sums is bounded above by any upper sum, and similarly the set of upper sums is bounded below. Thus the next definition is natural.

Definition 6.2.1 (Lower Integral, Upper Integral, Integrability, Integral). The lower integral of f over B is the least upper bound of the lower sums of f over all partitions P ,

$$L \int_B f = \sup \{L(f, P) : P \text{ is a partition of } B\}.$$

Similarly, the upper integral of f over B is the greatest lower bound of the upper sums of f over all partitions P ,

$$U \int_B f = \inf \{U(f, P) : P \text{ is a partition of } B\}.$$

The function f is called **integrable over B** if the lower and upper integrals are equal, i.e., if $L \int_B f = U \int_B f$. In this case, their shared value is called the **integral of f over B** and written $\int_B f$.

So we have a quantitative definition that seems appropriate. The integral, if it exists, is at least as big as any lower sum and at least as small as any upper sum; and it is specified as the common value that is approached from below by lower sums and from above by upper sums. Less formally, if quantities that we view as respectively too small and too big approach a common value, then that value must be what we're after.

The following lemma shows that $L \int_B f \leq U \int_B f$. Its proof provides an example of how to work with lower and upper bounds. Note that the argument does not require a contradiction or an ε , but rather it goes directly to the point.

Lemma 6.2.2 (Persistence of Order). Let \mathcal{L} and \mathcal{U} be nonempty sets of real numbers such that

$$\ell \leq u \quad \text{for all } \ell \in \mathcal{L} \text{ and } u \in \mathcal{U}. \quad (6.1)$$

Then $\sup(\mathcal{L})$ and $\inf(\mathcal{U})$ exist, and they satisfy

$$\sup(\mathcal{L}) \leq \inf(\mathcal{U}).$$

Proof. The given condition (6.1) says that every element ℓ of \mathcal{L} is a lower bound of \mathcal{U} . Since \mathcal{U} is nonempty and has lower bounds, it has a greatest lower bound $\inf(\mathcal{U})$. Since each $\ell \in \mathcal{L}$ is a lower bound and $\inf(\mathcal{U})$ is the greatest lower bound,

$$\ell \leq \inf(\mathcal{U}) \quad \text{for all } \ell \in \mathcal{L}.$$

This shows that $\inf(\mathcal{U})$ is an upper bound of \mathcal{L} . Since \mathcal{L} is nonempty and has an upper bound, it has a least upper bound $\sup(\mathcal{L})$. Since $\sup(\mathcal{L})$ is the least upper bound and $\inf(\mathcal{U})$ is an upper bound,

$$\sup(\mathcal{L}) \leq \inf(\mathcal{U}).$$

This is the desired result. \square

Since always $L \int_B f \leq U \int_B f$, to show that $\int_B f$ exists it suffices to show that also the reverse inequality holds, $L \int_B f \geq U \int_B f$.

Not all bounded functions $f : B \rightarrow \mathbf{R}$ are integrable. The standard counterexample is the interval $B = [0, 1]$ and the function

$$f : B \rightarrow \mathbf{R}, \quad f(x) = \begin{cases} 0 & \text{if } x \text{ is irrational,} \\ 1 & \text{if } x \text{ is rational.} \end{cases}$$

Chasing through the definitions shows that for this B and f , any lower sum is $L(f, P) = 0$, so the lower integral is $L \int_B f = \sup\{0\} = 0$. Similarly, $U \int_B f = 1$. Since the upper and lower integrals don't agree, $\int_B f$ does not exist.

So the questions are: what functions *are* integrable, or at least, what are some general classes of integrable functions, and how does one evaluate their integrals? Working from the definitions, as in the last example, is a good exercise in simple cases to get familiar with the machinery, but as a general procedure it is hopelessly unwieldy. Here is one result that will help us in the next section to show that continuous functions are integrable.

Proposition 6.2.3 (Integrability Criterion). *Let B be a box, and let $f : B \rightarrow \mathbf{R}$ be a bounded function. Then f is integrable over B if and only if for every $\varepsilon > 0$, there exists a partition P of B such that $U(f, P) - L(f, P) < \varepsilon$.*

Proof. (\implies) Let f be integrable over B and let $\varepsilon > 0$ be given. There exist partitions P and P' of B such that

$$L(f, P) > \int_B f - \varepsilon/2 \quad \text{and} \quad U(f, P') < \int_B f + \varepsilon/2.$$

Let P'' be the common refinement of P and P' . Then since refining increases lower sums and decreases upper sums, also

$$L(f, P'') > \int_B f - \varepsilon/2 \quad \text{and} \quad U(f, P'') < \int_B f + \varepsilon/2.$$

This shows that $U(f, P'') - L(f, P'') < \varepsilon$, as required.

(\impliedby) We need to show that $L \int_B f = U \int_B f$. To do so, use the cute principle that to prove that two numbers are equal, it suffices to show that they are within arbitrary positive ε of each other. So for any given $\varepsilon > 0$, we must show that

$$U \int_B f - L \int_B f < \varepsilon.$$

But by assumption there exists a partition P such that

$$U(f, P) - L(f, P) < \varepsilon,$$

and by the definition of upper and lower integral, also

$$L(f, P) \leq L \int_B f \leq U \int_B f \leq U(f, P).$$

The last two displays combine to give exactly what we need. \square

Here is an example of how to use the Integrability Criterion.

Proposition 6.2.4. *Let B be a box, let $f : B \rightarrow \mathbf{R}$ be a bounded function, and let P be a partition of B . If f is integrable over B then f is integrable over each subbox J of P , in which case*

$$\sum_J \int_J f = \int_B f.$$

Proof. Consider any partition P' of B that refines P . For each subbox J of P , let $P'_J = P' \cap J$, a partition of J . Let the symbol J' denote subboxes of P' , and compute that

$$L(f, P') = \sum_{J'} m_{J'}(f) \operatorname{vol}(J') = \sum_J \sum_{J' \subset J} m_{J'}(f) \operatorname{vol}(J') = \sum_J L(f, P'_J).$$

Similarly, $U(f, P') = \sum_J U(f, P'_J)$.

Suppose that f is integrable over B . Let an arbitrary $\varepsilon > 0$ be given. By “ \implies ” of the Integrability Criterion, there exists a partition P' of B such that

$$U(f, P') - L(f, P') < \varepsilon.$$

Since refining a partition cannot increase the difference between the upper and lower sums, we may replace P' by its common refinement with P and thus assume that P' refines P . Therefore the formulas from the previous paragraph show that

$$\sum_J (U(f, P'_J) - L(f, P'_J)) < \varepsilon,$$

and so

$$U(f, P'_J) - L(f, P'_J) < \varepsilon \quad \text{for each subbox } J \text{ of } B.$$

Therefore f is integrable over each subbox J of B by “ \impliedby ” of the Integrability Criterion.

Now assume that f is integrable over B and hence over each subbox J . Still letting P' be any partition of B that refines P , the integral over B lies between the lower and upper sums,

$$L(f, P') \leq \int_B f \leq U(f, P'),$$

and this same fact about the integral over each subbox J combines with the formulas from the first paragraph of the proof to give

$$L(f, P') = \sum_J L(f, P'_J) \leq \sum_J \int_J f \leq \sum_J U(f, P'_J) = U(f, P').$$

Since both $\int_B f$ and $\sum_J \int_J f$ lie between $L(f, P')$ and $U(f, P')$, they are within ε of each other. Since $\varepsilon > 0$ is arbitrary, they are equal by the cute principle enunciated in the proof of the Integrability Criterion. \square

Similar techniques show that the converse of the proposition holds as well, so that given B , f , and P , f is integrable over B if and only if f is integrable over each subbox J , but we do not need this full result.

The symbol B denotes a box in the next set of exercises.

Exercises

6.2.1. Let $f : B \rightarrow \mathbf{R}$ be a bounded function. Explain how Lemma 6.2.2 shows that $L \int_B f \leq U \int_B f$.

6.2.2. Let U and L be real numbers satisfying $U \geq L$. Show that $U = L$ if and only if for all $\varepsilon > 0$, $U - L < \varepsilon$.

6.2.3. Let $f : B \rightarrow \mathbf{R}$ be the constant function $f(x) = k$ for all $x \in B$. Show that f is integrable over B and $\int_B f = k \cdot \text{vol}(B)$.

6.2.4. Fill in the details in the argument that the function $f : [0, 1] \rightarrow \mathbf{R}$ with $f(x) = 0$ for irrational x and $f(x) = 1$ for rational x is not integrable over $[0, 1]$.

6.2.5. Let $B = [0, 1] \times [0, 1] \subset \mathbf{R}^2$. Define a function $f : B \rightarrow \mathbf{R}$ by

$$f(x, y) = \begin{cases} 0 & \text{if } 0 \leq x < 1/2, \\ 1 & \text{if } 1/2 \leq x \leq 1. \end{cases}$$

Show that f is integrable and $\int_B f = 1/2$.

6.2.6. This exercise shows that integration is linear. Let $f : B \rightarrow \mathbf{R}$ and $g : B \rightarrow \mathbf{R}$ be integrable.

(a) Let P be a partition of B and let J be some subbox of P . Show that

$$m_J(f) + m_J(g) \leq m_J(f + g) \leq M_J(f + g) \leq M_J(f) + M_J(g).$$

Show that consequently,

$$L(f, P) + L(g, P) \leq L(f + g, P) \leq U(f + g, P) \leq U(f, P) + U(g, P).$$

(b) Part (a) of this exercise obtained comparisons between lower and upper sums, analogously to the first paragraph of the proof of Proposition 6.2.4. Argue analogously to the rest of the proof to show $\int_B (f + g)$ exists and equals

$\int_B f + \int_B g$. (One way to begin is by using the Integrability Criterion twice and then a common refinement to show that there exists a partition P of B such that $U(f, P) - L(f, P) < \varepsilon/2$ and $U(g, P) - L(g, P) < \varepsilon/2$.)

(c) Let $c \geq 0$ be any constant. Let P be any partition of B . Show that for any subbox J of P ,

$$m_J(cf) = c m_J(f) \quad \text{and} \quad M_J(cf) = c M_J(f).$$

Explain why consequently

$$L(cf, P) = c L(f, P) \quad \text{and} \quad U(cf, P) = c U(f, P).$$

Explain why consequently

$$L \int_B cf = c L \int_B f \quad \text{and} \quad U \int_B cf = c U \int_B f.$$

Explain why consequently $\int_B cf$ exists and

$$\int_B cf = c \int_B f.$$

(d) Let P be any partition of B . Show that for any subbox J of P ,

$$m_J(-f) = -M_J(f) \quad \text{and} \quad M_J(-f) = -m_J(f).$$

Explain why consequently

$$L(-f, P) = -U(f, P) \quad \text{and} \quad U(-f, P) = -L(f, P).$$

Explain why consequently

$$L \int_B (-f) = -U \int_B f \quad \text{and} \quad U \int_B (-f) = -L \int_B f.$$

Explain why consequently $\int_B (-f)$ exists and

$$\int_B (-f) = - \int_B f.$$

Explain why the work so far here in part (d) combines with part (c) to show that for any $c \in \mathbf{R}$ (positive, zero, or negative), $\int_B cf$ exists and

$$\int_B cf = c \int_B f.$$

6.2.7. This exercise shows that integration preserves order. Let $f : B \rightarrow \mathbf{R}$ and $g : B \rightarrow \mathbf{R}$ both be integrable, and suppose that $f \leq g$, meaning that $f(x) \leq g(x)$ for all $x \in B$. Show that $\int_B f \leq \int_B g$. (It might be tidiest to begin by explaining why the previous exercise lets us assume that $f = 0$.)

6.2.8. Suppose that $f : B \rightarrow \mathbf{R}$ is integrable, and that so is $|f|$. Show that $|\int_B f| \leq \int_B |f|$.

6.2.9. Prove the converse to Proposition 6.2.4: Let B be a box, let $f : B \rightarrow \mathbf{R}$ be a bounded function, and let P be a partition of B . If f is integrable over each subbox J of P then f is integrable over B . (You may quote the formulas from the first paragraph of the proof in the text, since that paragraph makes no assumptions of integrability. It may help to let b denote the number of subboxes J , so that this quantity has a handy name.)

6.3 Continuity and Integrability

Although the Integrability Criterion gives a test for the integrability of any specific function f , it is cumbersome to apply case by case. But handily, it will provide the punchline of the proof of the next theorem, which says that a natural class of functions is integrable.

Theorem 6.3.1 (Continuity Implies Integrability). *Let B be a box, and let $f : B \rightarrow \mathbf{R}$ be a continuous function. Then f is integrable over B .*

To prove this, as we will at the end of the section, we first need to sharpen our understanding of continuity on boxes. The version of continuity that we're familiar with isn't strong enough to prove certain theorems, this one in particular. Formulating the stronger version of continuity requires first revising the grammar of the familiar brand.

Definition 6.3.2 (Sequential Continuity). *Let $S \subset \mathbf{R}^n$ be a set, and let $f : S \rightarrow \mathbf{R}^m$ be a mapping. For any $x \in S$, f is **sequentially continuous at x** if for every sequence $\{x_\nu\}$ in S converging to x , the sequence $\{f(x_\nu)\}$ converges to $f(x)$. The mapping f is **sequentially continuous on S** if f is sequentially continuous at each point x in S .*

Definition 6.3.3 (ε - δ Continuity). *Let $S \subset \mathbf{R}^n$ be a set, and let $f : S \rightarrow \mathbf{R}^m$ be a mapping. For any $x \in S$, f is **ε - δ continuous at x** if for every $\varepsilon > 0$ there exists some $\delta > 0$ such that*

$$\text{if } \tilde{x} \in S \text{ and } |\tilde{x} - x| < \delta \text{ then } |f(\tilde{x}) - f(x)| < \varepsilon.$$

*The mapping f is **ε - δ continuous on S** if f is ε - δ continuous at each point x in S .*

Both definitions of continuity at a point x capture the idea that as inputs to f approach x , the corresponding outputs from f should approach $f(x)$. This is exactly the substance of sequential continuity. (See figure 6.10.)

For ε - δ continuity at x , imagine that someone has drawn a ball of radius ε (over which you have no control, and it's probably quite small) about the




Figure 6.10. Sequential continuity

point $f(x)$ in \mathbf{R}^m . The idea is that in response, you can draw a ball of some radius—this is the δ in the definition—about the point x in S , so that every point in the δ -ball about x gets taken by f into the ε -ball about $f(x)$. (See figure 6.11.)




Figure 6.11. ε - δ continuity

For example, the function $f : \mathbf{R}^n \rightarrow \mathbf{R}$ given by $f(x) = 2|x|$ is ε - δ continuous on \mathbf{R}^n . To show this, let $\varepsilon > 0$ be given. Set $\delta = \varepsilon/2$. Then whenever $|\tilde{x} - x| < \delta$, a calculation that uses the generalized triangle inequality at the third step shows that

$$|f(\tilde{x}) - f(x)| = |2|\tilde{x}| - 2|x|| = 2||\tilde{x}| - |x|| \leq 2|\tilde{x} - x| < 2\delta = \varepsilon,$$

as needed.

For another example, to prove that the function $f : \mathbf{R} \rightarrow \mathbf{R}$ given by $f(x) = x^2$ is ε - δ continuous on \mathbf{R} , pick any $x \in \mathbf{R}$ and let $\varepsilon > 0$ be given. This time set

$$\delta = \min\{1, \varepsilon/(2|x| + 1)\}.$$

This may look strange, but its first virtue is that since $\delta \leq 1$, for any $\tilde{x} \in \mathbf{R}$ with $|\tilde{x} - x| < \delta$, we have $|\tilde{x}| < |x| + 1$ and therefore $|\tilde{x}| + |x| < 2|x| + 1$; and its second virtue is that also $\delta \leq \varepsilon/(2|x| + 1)$. These conditions fit perfectly into the following calculation,

$$\begin{aligned}
|f(\tilde{x}) - f(x)| &= |\tilde{x}^2 - x^2| \\
&= |\tilde{x} + x| |\tilde{x} - x| \\
&\leq (|\tilde{x}| + |x|)\delta && \text{by the triangle inequality} \\
&< (2|x| + 1)\frac{\varepsilon}{2|x| + 1} && \text{by the two virtues of } \delta \\
&= \varepsilon.
\end{aligned}$$

And this is exactly what we needed to show that f is continuous at x .

The tricky part of writing this sort of proof is finding the right δ . Doing so generally requires some preliminary fiddling around on scratch paper. For the proof just given, the key scratch calculation would be

$$|f(\tilde{x}) - f(x)| = |\tilde{x} + x| |\tilde{x} - x|,$$

exhibiting the quantity that we need to bound by ε as a product of two terms, the second bounded directly by whatever δ we choose. The idea is to make the first term reasonably small (by stipulating that δ be at most 1), and then to constrain δ further to make the second term small enough that the product is less than ε . Hence the choice of δ in the proof.

In fact there is no need to continue proving that specific functions already known to be sequentially continuous are also ε - δ continuous. Each type of continuity implies the other.

Proposition 6.3.4 (Sequential Continuity and ε - δ Continuity are Equivalent). *For any set $S \subset \mathbf{R}^n$ and any mapping $f : S \rightarrow \mathbf{R}^m$, f is sequentially continuous on S if and only if f is ε - δ continuous on S .*

Proof. Let x be any point of S .

(\Leftarrow) Suppose that f is ε - δ continuous at x . We need to show that f is sequentially continuous at x . So, let $\{x_\nu\}$ be a sequence in S converging to x . To show that $\{f(x_\nu)\}$ converges to $f(x)$ means that given an arbitrary $\varepsilon > 0$, we need to exhibit a starting index N such that

$$\text{for all } \nu > N, |f(x_\nu) - f(x)| < \varepsilon.$$

The definition of ε - δ continuity gives a δ such that

$$\text{if } \tilde{x} \in S \text{ and } |\tilde{x} - x| < \delta \text{ then } |f(\tilde{x}) - f(x)| < \varepsilon.$$

And since $\{x_\nu\}$ converges in S to x , there is some starting index N such that

$$\text{for all } \nu > N, |x_\nu - x| < \delta.$$

The last two displays combine to imply the first display, showing that f is sequentially continuous at x .

(\Rightarrow) Now suppose that f is sequentially continuous at x . This time we need to show that f is ε - δ continuous at x . The proof is by contradiction.

So, let $\varepsilon > 0$ be given. Suppose that no $\delta > 0$ satisfies the condition for ε - δ continuity. Then in particular, $\delta = 1/\nu$ doesn't work for $\nu = 1, 2, 3, \dots$. So there is a sequence $\{x_\nu\}$ in S such that

$$|x_\nu - x| < 1/\nu \quad \text{and} \quad |f(x_\nu) - f(x)| \geq \varepsilon, \quad \nu = 1, 2, 3, \dots$$

This condition contradicts the sequential continuity of f at x . So, given $\varepsilon > 0$, some $\delta > 0$ must satisfy the condition for ε - δ continuity at x .

Since the two types of continuity imply one another at each point x of S , they imply one another on S . \square

The fact that the second half of this proof has to proceed by contradiction, whereas the first half is straightforward, shows that ε - δ continuity is a little more powerful than sequential continuity on the face of it, until we do the work of showing that they are equivalent. Also, the very definition of ε - δ continuity seems harder for students than the definition of sequential continuity, which is why these notes have used sequential continuity up to now. However, the exceptionally alert reader may have recognized that the second half of this proof is essentially identical to the proof of the Persistence of Inequality Principle (Proposition 2.3.9). Thus, the occasional arguments in these notes that cited Persistence of Inequality were tacitly using ε - δ continuity already, because sequential continuity was not transparently strong enough for their purposes. The reader who dislikes redundancy is encouraged to rewrite the second half of this proof to quote Persistence of Inequality rather than re-prove it.

The reason that we bother with this new ε - δ type of continuity, despite its equivalence to sequential continuity meaning that it is nothing new, is that its grammar generalizes to describe the more powerful continuity that we need. The two examples above of ε - δ continuity differed: in the example $f(x) = x^2$, the choice of $\delta = \min\{1, \varepsilon/(2|x| + 1)\}$ for any given x and ε to satisfy the definition of ε - δ continuity at x depended not only on ε but on x as well. In the example $f(x) = 2|x|$, the choice of $\delta = \varepsilon/2$ for any given x and ε depended only on ε , i.e., it was independent of x . Here, one value of δ works simultaneously at all values of x once ε is specified. This technicality has enormous consequences.

Definition 6.3.5 (Uniform Continuity). *Let $S \subset \mathbf{R}^n$ be a set, and let $f : S \rightarrow \mathbf{R}^m$ be a mapping. Then f is **uniformly continuous on S** if for every $\varepsilon > 0$ there exists some $\delta > 0$ such that*

$$\text{if } x, \bar{x} \in S \text{ and } |\bar{x} - x| < \delta \text{ then } |f(\bar{x}) - f(x)| < \varepsilon.$$

The nomenclature *uniformly continuous on S* is meant to emphasize that given $\varepsilon > 0$, a single, *uniform* value of δ works in the definition of ε - δ continuity simultaneously for all points $x \in S$. The scope of its effectiveness is large-scale. Uniform continuity depends on both the mapping f and the set S .

A visual image may help distinguish between the old notion of continuity (henceforth called **pointwise continuity**) and the new, stronger notion of

uniform continuity. Imagine the graph of a function $f : S \rightarrow \mathbf{R}$ (where $S \subset \mathbf{R}$), and take some input point x . Then f is pointwise continuous at x if for any $\varepsilon > 0$, one can draw a rectangle of height 2ε centered at the point $(x, f(x))$ that is narrow enough that the graph of f protrudes only from the sides of the rectangle, not the top or bottom. The base of the rectangle is 2δ , where δ comes from ε - δ continuity. Note that for a given ε , one may need rectangles of various widths at different points. A rectangle that works at x may not be narrow enough to work again at some other point \tilde{x} . (See figure 6.12, where ever-narrower rectangles are required as we move to the left on the graph.) On the other hand, the function f is uniformly continuous if given $\varepsilon > 0$, there is a single 2ε -by- 2δ rectangle that can slide along the entire graph of f with its centerpoint on the graph, and the graph never protruding from the top or bottom. (See figure 6.13. A tacit assumption here is that the graph of f either doesn't extend beyond the picture frame, or it continues to rise and fall tamely if it does.) By contrast, no single rectangle will work in figure 6.12.




Figure 6.12. One ε can require different values δ at different points x




Figure 6.13. Or one δ can work uniformly for ε at all x

The domain of the nonuniformly continuous function $f(x) = \sin(1/x)$ in figure 6.12 is not compact, not being closed at its left endpoint. We are about to prove that on a compact domain, uniform continuity follows for free from pointwise continuity. In conjunction with the compactness of the boxes B over which we integrate, this is the crucial ingredient for proving Theorem 6.3.1 (continuous functions are integrable over boxes), the goal of this section.

Theorem 6.3.6 (Continuity on Compact Sets is Uniform). *Let $K \subset \mathbf{R}^n$ be compact, and let $f : K \rightarrow \mathbf{R}^m$ be pointwise continuous on K . Then f is uniformly continuous on K .*

Proof. The proof proceeds by contradiction, so suppose that we have such K and f , but f is not uniformly continuous on K . Then for some $\varepsilon > 0$ there exists no uniform δ , so in particular no reciprocal positive integer $1/\nu$ will serve as δ in the definition of uniform continuity. Thus for each $\nu \in \mathbf{Z}^+$ there exist points x_ν and y_ν in K such that

$$|x_\nu - y_\nu| < 1/\nu \quad \text{and} \quad |f(x_\nu) - f(y_\nu)| \geq \varepsilon. \quad (6.2)$$

Consider the sequences $\{x_\nu\}$ and $\{y_\nu\}$ in K . By the sequential characterization of compactness (Theorem 2.4.13), $\{x_\nu\}$ has a convergent subsequence converging in K ; call it $\{x_{\nu_k}\}$. Throw away the rest of the x_ν 's and throw away the y_ν 's of corresponding index, reindex the remaining terms of the two sequences, and now $\{x_\nu\}$ converges to some $p \in K$. Since $|x_\nu - y_\nu| < 1/\nu$ for each ν (this remains true after the reindexing), $\{y_\nu\}$ converges to p as well. So

$$\lim x_\nu = p = \lim y_\nu,$$

and now the continuity of f gives

$$\lim f(x_\nu) = f(p) = \lim f(y_\nu).$$

This violates the second condition in (6.2) even though the first condition holds, and so the proof by contradiction is complete. \square

Recall the main result that we want: If B is a box in \mathbf{R}^n and $f : B \rightarrow \mathbf{R}$ is continuous then $\int_B f$ exists. This is easy to prove now. The crucial line of the proof is the opener.

Proof (of Theorem 6.3.1). The continuity of f on B is uniform. Thus, given $\varepsilon > 0$, there exists $\delta > 0$ such that

$$\text{if } x, \tilde{x} \in B \text{ and } |\tilde{x} - x| < \delta \text{ then } |f(\tilde{x}) - f(x)| < \frac{\varepsilon}{\text{vol}(B)}.$$

(We may take $\text{vol}(B) > 0$, making the volume safe to divide by, since otherwise all lower sums and upper sums are 0, making the integral 0 as well, and there is nothing to prove.) Take a partition P of B whose subboxes J have sides

of length less than δ/n . By the Size Bounds (Proposition 2.2.7), any points x and \tilde{x} in any subbox J satisfy $|\tilde{x} - x| < \delta$, so

$$\text{if } x, \tilde{x} \in J \text{ then } |f(\tilde{x}) - f(x)| < \frac{\varepsilon}{\text{vol}(B)}.$$

Let x and \tilde{x} vary over J , and cite the Extreme Value Theorem (Theorem 2.4.15) to show that

$$M_J(f) - m_J(f) < \frac{\varepsilon}{\text{vol}(B)}.$$

Multiply by $\text{vol}(J)$ to get

$$M_J(f)\text{vol}(J) - m_J(f)\text{vol}(J) < \frac{\varepsilon \text{vol}(J)}{\text{vol}(B)},$$

and sum this relation over subboxes J to get

$$U(f, P) - L(f, P) < \varepsilon.$$

The Integrability Criterion now shows that $\int_B f$ exists. \square

Integration synthesizes local data at each point of a domain into one whole. The idea of this section is that integrating a continuous function over a box is more than a purely local process: it requires the uniform continuity of the function all through the box, a large-scale simultaneous estimate that holds in consequence of the box being compact.

Exercises

6.3.1. Reread the proof that sequential and ε - δ continuity are equivalent, then redo the proof with the book closed.

6.3.2. Let $f : \mathbf{R} \rightarrow \mathbf{R}$ be the cubing function $f(x) = x^3$. Give a direct proof that f is ε - δ continuous on \mathbf{R} . (Hint: $A^3 - B^3 = (A - B)(A^2 + AB + B^2)$.)

6.3.3. Here is a proof that the squaring function $f(x) = x^2$ is not uniformly continuous on \mathbf{R} . Suppose that some $\delta > 0$ satisfies the definition of uniform continuity for $\varepsilon = 1$. Set $x = 1/\delta$ and $\tilde{x} = 1/\delta + \delta/2$. Then certainly $|\tilde{x} - x| < \delta$, but

$$|f(\tilde{x}) - f(x)| = \left| \left(\frac{1}{\delta} + \frac{\delta}{2} \right)^2 - \frac{1}{\delta^2} \right| = \left| \frac{1}{\delta^2} + 1 + \frac{\delta^2}{4} - \frac{1}{\delta^2} \right| = 1 + \frac{\delta^2}{4} > \varepsilon.$$

This contradicts uniform continuity.

Is the cubing function of the previous exercise uniformly continuous on \mathbf{R} ? On $[0, 500]$?

6.3.4. (a) Show: If $I \subset \mathbf{R}$ is an interval (possibly all of \mathbf{R}), $f : I \rightarrow \mathbf{R}$ is differentiable, and there exists a positive constant R such that $|f'(x)| \leq R$ for all $x \in I$ then f is uniformly continuous on I .

(b) Prove that sine and cosine are uniformly continuous on \mathbf{R} .

6.3.5. Let $f : [0, +\infty)$ be the square root function $f(x) = \sqrt{x}$. You may take for granted that f is ε - δ continuous on $[0, +\infty)$.

(a) What does part (a) of the previous problem say about the uniform continuity of f ?

(b) Is f uniformly continuous?

6.3.6. Let J be a box in \mathbf{R}^n with sides of length less than δ/n . Show that any points x and \tilde{x} in J satisfy $|\tilde{x} - x| < \delta$.

6.3.7. For $\int_B f$ to exist, it is sufficient that $f : B \rightarrow \mathbf{R}$ be continuous, but it is not necessary. What preceding exercise provides an example of this? Here is another example. Let $B = [0, 1]$ and let $f : B \rightarrow \mathbf{R}$ be monotonic increasing, meaning that if $x_1 < x_2$ in B then $f(x_1) \leq f(x_2)$. Show that such a function is bounded, though it need not be continuous. Use the Integrability Criterion to show that $\int_B f$ exists.

6.4 Integration of Functions of One Variable

In a first calculus course one learns to do computations such as: to evaluate

$$\int_{x=1}^e \frac{(\log x)^2}{x} dx,$$

let $u = \log x$; then $du = dx/x$, and as x goes from 1 to e , u goes from 0 to 1, so the integral equals

$$\int_{u=0}^1 u^2 du = \frac{1}{3}u^3 \Big|_0^1 = \frac{1}{3}.$$

Or such as: to evaluate

$$\int_0^9 \frac{dx}{\sqrt{1+\sqrt{x}}}.$$

let $u = \sqrt{1+\sqrt{x}}$. Then some algebra shows that $x = (u^2 - 1)^2$, and so $dx = 4(u^2 - 1)u du$. Also, when $x = 0$, $u = 1$, and when $x = 9$, $u = 2$. Therefore the integral is

$$\begin{aligned} \int_0^9 \frac{dx}{\sqrt{1+\sqrt{x}}} &= 4 \int_1^2 \frac{(u^2 - 1)u}{u} du = 4 \int_1^2 (u^2 - 1) du \\ &= 4 \left(\frac{1}{3}u^3 - u \right) \Big|_1^2 = \frac{16}{3}. \end{aligned}$$

Although both of these examples use substitution, they differ from one another in a way that a first calculus course may not explain. The first substitution involved picking an x -dependent u (i.e., $u = \ln x$) where $u'(x)$ (i.e., $1/x$) was present in the integral and got absorbed by the substitution. The second substitution was of an opposite form to the first: this time the x -dependent u was inverted to produce a u -dependent x , and the factor $u'(x)$ was introduced into the integral rather than eliminated from it. Somehow, two different things are going on under the guise of “ u -substitution.”

In this section we specialize our theory of multivariable integration to $n = 1$ and review two tools for evaluating one-dimensional integrals, the Fundamental Theorem of Integral Calculus (FTIC) and the Change of Variable Theorem. Writing these down precisely will clarify the examples we just worked. More importantly, generalizing these results appropriately to n dimensions is the subject of the remainder of these notes.

The multivariable integral notation of this chapter, specialized to one dimension, is $\int_{[a,b]} f$. For familiarity, replace this by the usual notation,

$$\int_a^b f = \int_{[a,b]} f \quad \text{for } a \leq b.$$

As matters stand, the redefined notation $\int_a^b f$ makes sense only when $a \leq b$, so extend its definition to

$$\int_a^b f = - \int_b^a f \quad \text{for } a > b.$$

Once this is done, the same relation between “signed” integrals holds regardless of which (if either) of a and b is larger,

$$\int_a^b f = - \int_b^a f \quad \text{for all } a \text{ and } b.$$

Something nontrivial is happening here: when the multivariable integration of this chapter is specialized to one dimension, it can be extended to incorporate a sign convention to represent the order on \mathbf{R} . If $a < b$ then \int_a^b describes positive traversal along the real line from a up to b , while \int_b^a describes negative traversal from b down to a . This sort of thing does not obviously generalize to higher dimensions since \mathbf{R}^n is not ordered.

Casewise inspection shows that for any three points $a, b, c \in \mathbf{R}$ in any order, and for any integrable function $f : [\min\{a, b, c\}, \max\{a, b, c\}] \rightarrow \mathbf{R}$,

$$\int_a^c f = \int_a^b f + \int_b^c f.$$

Also, if $f : [\min\{a, b\}, \max\{a, b\}] \rightarrow \mathbf{R}$ takes the constant value k then

$$\int_a^b f = k(b-a),$$

again regardless of which of a and b is larger. These facts generalize Proposition 6.2.4 and exercise 6.2.3 to signed one-variable integration.

Each of the next two theorems describes a sense in which one-variable differentiation and integration are inverse operations. Both are called the Fundamental Theorem of Integral Calculus, but the second is more deserving of the title because of how far it generalizes.

Theorem 6.4.1. *Let the function $f : [a, b] \rightarrow \mathbf{R}$ be continuous. Define a function*

$$F : [a, b] \rightarrow \mathbf{R}, \quad F(x) = \int_a^x f.$$

Then F is differentiable on $[a, b]$ and $F' = f$.

Proof. Let x and $x+h$ lie in $[a, b]$ with $h \neq 0$. Study the difference quotient

$$\frac{F(x+h) - F(x)}{h} = \frac{\int_a^{x+h} f - \int_a^x f}{h} = \frac{\int_x^{x+h} f}{h}.$$

If $h > 0$ then $m_{[x, x+h]}(f) \cdot h \leq \int_x^{x+h} f \leq M_{[x, x+h]}(f) \cdot h$, and dividing through by h shows that the difference quotient lies between $m_{[x, x+h]}(f)$ and $M_{[x, x+h]}(f)$. This forces it to $f(x)$ as h goes to 0, since f is continuous. Similar analysis applies when $h < 0$. \square

The alert reader will recall the convention in these notes that a mapping can be differentiable only at an interior point of its domain. In particular, the derivative of a function $F : [a, b] \rightarrow \mathbf{R}$ is undefined at a and b . Hence the statement of Theorem 6.4.1 is inconsistent with our usage, and strictly speaking the theorem should conclude that F is continuous on $[a, b]$ and differentiable on (a, b) with derivative $F' = f$. The given proof does show this, since the existence of the one-sided derivative of F at each endpoint makes F continuous there.

However, we prohibited derivatives at endpoints only to tidy up our statements. An alternative would have been to make the definition that for any compact, connected set $K \subset \mathbf{R}^n$ (both of these terms were discussed in section 2.4), a mapping $f : K \rightarrow \mathbf{R}^m$ is differentiable on K if there exists an open set $A \subset \mathbf{R}^n$ containing K , and an extension of f to a differentiable mapping $f : A \rightarrow \mathbf{R}^m$. Here the word “extension” means that the new function f on A has the same behavior on K as the old f . One reason that we avoided this slightly more general definition is that it is tortuous to track through the material in chapter 4, especially for the student who is seeing the ideas for the first time. Also, this definition requires that the Critical Point Theorem (stating that the extrema a function occur at points where its derivative is 0) be fussily rephrased to say that this criterion applies only to the extrema that

occur at the interior points of the domain. From the same preference for tidy statements over fussy ones, now we do allow the more general definition of the derivative.

Proving the FTIC from Theorem 6.4.1 requires the observation that if two functions $F_1, F_2 : [a, b] \rightarrow \mathbf{R}$ are differentiable, and $F_1' = F_2'$, then $F_1 = F_2 + c$ for some constant c . This follows from the Mean Value Theorem and is an exercise.

Theorem 6.4.2 (Fundamental Theorem of Integral Calculus). *Suppose that the function $F : [a, b] \rightarrow \mathbf{R}$ is differentiable and F' is continuous. Then*

$$\int_a^b F' = F(b) - F(a).$$

Proof. Define $F_2 : [a, b] \rightarrow \mathbf{R}$ by $F_2(x) = \int_a^x F'$. Then $F_2' = F'$ by the preceding theorem, so there exists a constant c such that for all $x \in [a, b]$,

$$F_2(x) = F(x) + c. \quad (6.3)$$

Plug $x = a$ into (6.3) to get $0 = F(a) + c$, so $c = -F(a)$. Next plug in $x = b$ to get $F_2(b) = F(b) - F(a)$. Since $F_2(b) = \int_a^b F'$ by definition, the proof is complete. \square

In one-variable calculus one learns various techniques to find antiderivatives; i.e., given continuous f , one finds F such that $F' = f$. Once this is done, evaluating $\int_a^b f$ is mere plug-in to the FTIC. But since not all continuous functions have antiderivatives that are readily found, or even possible to write in an elementary form (for example, try $f(x) = e^{-x^2}$ or $f(x) = \sin(x^2)$), the FTIC has its limitations.

Another tool for evaluating one-dimensional integrals is the Change of Variable Theorem. The idea is to transform one integral to another that may be better suited to the FTIC.

Theorem 6.4.3 (Change of Variable Theorem; Forward Substitution Formula). *Let $\phi : [a, b] \rightarrow \mathbf{R}$ be differentiable with continuous derivative and let $f : \phi[a, b] \rightarrow \mathbf{R}$ be continuous. Then*

$$\int_a^b (f \circ \phi) \cdot \phi' = \int_{\phi(a)}^{\phi(b)} f. \quad (6.4)$$

Proof. Use Theorem 6.4.1 to define $F : \phi[a, b] \rightarrow \mathbf{R}$ such that $F' = f$. By the chain rule, $F \circ \phi$ has derivative $(F \circ \phi)' = (F' \circ \phi) \cdot \phi' = (f \circ \phi) \cdot \phi'$, which is continuous on $[a, b]$. Thus by the FTIC twice,

$$\begin{aligned} \int_a^b (f \circ \phi) \cdot \phi' &= \int_a^b (F \circ \phi)' = (F \circ \phi)(b) - (F \circ \phi)(a) \\ &= F(\phi(b)) - F(\phi(a)) = \int_{\phi(a)}^{\phi(b)} F' = \int_{\phi(a)}^{\phi(b)} f. \end{aligned}$$

\square

One way to apply the Change of Variable Theorem to an integral $\int_a^b g$ is to recognize that the integrand takes the form $g = (f \circ \phi) \cdot \phi'$, giving the left side of (6.4) for suitable f and ϕ such that the right side $\int_{\phi(a)}^{\phi(b)} f$ is easier to evaluate. This method is called **integration by forward substitution**. For instance, for the first integral $\int_{x=1}^e ((\log x)^2)/x dx$ at the beginning of the section, take

$$g : \mathbf{R}^+ \longrightarrow \mathbf{R}, \quad g(x) = (\log x)^2/x.$$

To evaluate $\int_1^e g$, define

$$\phi : \mathbf{R}^+ \longrightarrow \mathbf{R}, \quad \phi(x) = \log x$$

and

$$f : \mathbf{R} \longrightarrow \mathbf{R}, \quad f(u) = u^2.$$

Then $g = (f \circ \phi) \cdot \phi'$, and $\phi(1) = 0$, $\phi(e) = 1$, so by the Change of Variable Theorem,

$$\int_1^e g = \int_1^e (f \circ \phi) \cdot \phi' = \int_{\phi(1)}^{\phi(e)} f = \int_0^1 f.$$

Since f has antiderivative F where $F(u) = u^3/3$, the last integral equals $F(1) - F(0) = 1/3$ by the FTIC.

The second integral at the beginning of the section was evaluated not by the Change of Variable Theorem as given, but by a consequence of it:

Corollary 6.4.4 (Inverse Substitution Formula). *Let $\phi : [a, b] \longrightarrow \mathbf{R}$ be differentiable with continuous derivative and let $f : \phi[a, b] \longrightarrow \mathbf{R}$ be continuous. Suppose further that ϕ is invertible and that ϕ^{-1} is differentiable. Then*

$$\int_a^b (f \circ \phi) = \int_{\phi(a)}^{\phi(b)} f \cdot (\phi^{-1})'.$$

This is the formula for **integration by inverse substitution**. To obtain it from (6.4), replace ϕ by ϕ^{-1} , f by $f \circ \phi$, a by $\phi(a)$, and b by $\phi(b)$, and then exchange the two sides of the equality.

To apply this formula to an integral $\int_a^b g$, write the integrand as $g = f \circ \phi$, giving the left side, and then invert ϕ and differentiate the inverse to see if the right side is easier to evaluate. For instance, for the second integral $\int_0^9 dx/\sqrt{1+\sqrt{x}}$ at the beginning of the section, let id denote the identity function, and take

$$g : \mathbf{R}_{\geq 0} \longrightarrow \mathbf{R}, \quad g = 1/\sqrt{1+\sqrt{id}}.$$

To evaluate $\int_0^9 g$, define

$$\phi : \mathbf{R}_{\geq 0} \longrightarrow \mathbf{R}, \quad \phi = \sqrt{1+\sqrt{id}}$$

and

$$f : \mathbf{R}^+ \longrightarrow \mathbf{R}, \quad f(u) = 1/\text{id}.$$

Thus $g = f \circ \phi$ as desired. From the definition of ϕ , it follows by algebra that $(\phi^2 - 1)^2 = \text{id}$. Letting s denote the squaring function, this relation rewrites as $s \circ (s \circ \phi - 1) = \text{id}$. Consequently,

$$\phi^{-1} = s \circ (s - 1).$$

Since $s' = 2 \text{id}$, it follows by the chain rule that

$$(\phi^{-1})' = (2 \text{id} \circ (s - 1)) \cdot 2 \text{id} = 4(s - 1) \cdot \text{id}.$$

Also, $\phi(0) = 1$ and $\phi(9) = 2$. By the inverse substitution formula, the integral is now

$$\int_0^9 (f \circ \phi) = \int_{\phi(0)}^{\phi(9)} f \cdot (\phi^{-1})' = 4 \int_1^2 \frac{(s - 1) \cdot \text{id}}{\text{id}} = 4 \int_1^2 (s - 1).$$

Since $s - 1$ has antiderivative $c/3 - \text{id}$ where c is the cubing function, the integral equals $4(c(2)/3 - 2 - c(1)/3 + 1) = 16/3$ as before.

The variable-based notation used to work the two integrals at the beginning of the section, with x and u and dx and du , is much easier mnemonically than the function-based notation used to rework them with the Change of Variable Theorem and its corollary. But a purist would object to it on two counts. First, expressions such as $(\log x)^2/x$ and u^2 are not functions, they are the outputs of functions, so strictly speaking we can't integrate them. This is not a serious problem, it is mere pedantry: we simply need to loosen our notation to let $\int_{x=a}^b f(x)$ be synonymous with $\int_a^b f$, at the cost of an unnecessary new symbol x . This x is called a "dummy variable" since another symbol would do just as well: $\int_{y=a}^b f(y)$ and $\int_{\heartsuit=a}^b f(\heartsuit)$ also denote $\int_a^b f$. At the theoretical level, where we deal with functions *qua* functions, this extra notation is useless and cumbersome, but in any down-to-earth example it is in fact a convenience since describing functions by formulas is easier and more direct than introducing new symbols to name them.

The second, more serious objection to the variable-based notation is to the dx , the du , and mysterious relations such as $du = dx/x$ between them. What kind of objects are dx and du ? In a first calculus course they are typically described as infinitesimally small changes in x and u , but our theory of integration is not based on such hazy notions; in fact it was created in the 19th century to answer objections to their validity. (Though infinitesimals were revived and put on a firm footing in the 1960s, we have no business with them here.) An alternative is to view dx and du as formal symbols that serve, along with the integral sign \int , as bookends around the expression for the function being integrated. This leaves notation such as $du = dx/x$ still meaningless in its own right. In a first calculus course it may be taught as

a procedure with no real justification, whereas by contrast, the revisited versions of the two integral-calculations of this section are visibly applications of results that have been proved. However, the classical method is probably easier for most of us, its notational conventions dovetailing with the Change of Variable Theorem and its corollary so well. So feel free to continue using it. (And remember to switch the limits of integration when you do.)

However, to underscore that dx is an unnecessary, meaningless symbol, it will not be used in these notes until it is defined next chapter, as something called a *differential form*.

Exercises

6.4.1. (a) Show that for three points $a, b, c \in \mathbf{R}$ in any order, and any integrable function $f : [\min\{a, b, c\}, \max\{a, b, c\}] \rightarrow \mathbf{R}$, $\int_a^c f = \int_a^b f + \int_b^c f$.

(b) Show that if $f : [\min\{a, b\}, \max\{a, b\}] \rightarrow \mathbf{R}$ takes the constant value k then $\int_a^b f = k(b - a)$, regardless of which of a and b is larger.

6.4.2. Complete the proof of Theorem 6.4.1 by analyzing the case $h < 0$.

6.4.3. Show that if $F_1, F_2 : [a, b] \rightarrow \mathbf{R}$ are differentiable and $F_1' = F_2'$, then $F_1 = F_2 + C$ for some constant C .

6.4.4. (a) Suppose that $0 \leq a \leq b$ and $f : [a^2, b^2] \rightarrow \mathbf{R}$ is continuous. Define $F : [a, b] \rightarrow \mathbf{R}$ by $F(x) = \int_{a^2}^{x^2} f$. Does F' exist, and if so then what is it?

(b) More generally, suppose $f : \mathbf{R} \rightarrow \mathbf{R}$ is continuous, and $\alpha, \beta : \mathbf{R} \rightarrow \mathbf{R}$ are differentiable. Define $F : \mathbf{R} \rightarrow \mathbf{R}$ by $F(x) = \int_{\alpha(x)}^{\beta(x)} f$. Does F' exist, and if so then what is it?

6.4.5. Let $f : [0, 1] \rightarrow \mathbf{R}$ be continuous and suppose that for all $x \in [0, 1]$, $\int_0^x f = \int_x^1 f$. What is f ?

6.4.6. Find all differentiable functions $f : \mathbf{R}_{\geq 0} \rightarrow \mathbf{R}$ such that for all $x \in \mathbf{R}_{\geq 0}$, $(f(x))^2 = \int_0^x f$.

6.4.7. Define $f : \mathbf{R}^+ \rightarrow \mathbf{R}$ by $f(u) = e^{(u+1/u)}/u$ and $F : \mathbf{R}^+ \rightarrow \mathbf{R}$ by $F(x) = \int_1^x f$. Show that F behaves somewhat like a logarithm in that $F(1/x) = -F(x)$ for all $x \in \mathbf{R}^+$. Interpret this property of F as a statement about area under the graph of f . (Hint: define $\phi : \mathbf{R}^+ \rightarrow \mathbf{R}^+$ by $\phi(u) = 1/u$, and show that $(f \circ \phi) \cdot \phi' = -f$.)

6.5 Integration Over Nonboxes

So far we know that $\int_B f$ exists if B is a box and $f : B \rightarrow \mathbf{R}$ is continuous (Theorem 6.3.1). With some more work, the theorem can be refined to relax

these requirements. The basic idea is that $\int_B f$ still exists if f is discontinuous on a small enough subset of B . This isn't hard conceptually, but it requires some bookkeeping. Once it is established, integration over compact sets K other than boxes is easy to define provided that their boundaries are suitably small.

To quantify the notion of small, and more generally the notion of set size, let a set $S \subset \mathbf{R}^n$ be given. The **characteristic function of S** is

$$\chi_S : \mathbf{R}^n \longrightarrow \mathbf{R}, \quad \chi_S(x) = \begin{cases} 1 & \text{if } x \in S \\ 0 & \text{otherwise.} \end{cases}$$

Suppose that S is bounded, meaning that S sits in some box B .

Definition 6.5.1 (Volume of a Set). *The volume of a bounded set $S \subset \mathbf{R}^n$ is*

$$\text{vol}(S) = \int_B \chi_S \quad \text{where } B \text{ is any box containing } S,$$

if this integral exists.

This definition requires several comments. At first glance it seems ill-posed. Conceivably, $\int_B \chi_S$ could exist for some boxes B containing S but not others, and it could take different values for the various B where it exists. In fact, some technique shows that if $\int_B \chi_S$ exists for *some* box B containing S then it exists for *any* such box and always takes the same value, so the definition makes sense after all. See the exercises. Also, an exercise shows that the volume of a box B is the same under Definition 6.5.1 as under Definition 6.1.4, as it must be for grammatical consistency. Finally, note that not all sets have volume, only those whose characteristic functions are integrable.

Sets of volume zero are small enough that they don't interfere with integration. To prove such a result explicitly, we first translate the definition of volume zero into statements about the machinery of the integral. Let $S \subset \mathbf{R}^n$ sit in a box B , and let P be a partition of B . The subboxes J of P consist of two types:

$$\text{type I : } J \text{ such that } J \cap S \neq \emptyset$$

and

$$\text{type II : } J \text{ such that } J \cap S = \emptyset.$$

Thus S sits in the union of subboxes J of type I and the sum of their volumes gives an upper sum for $\int_B \chi_S$.

For example, figure 6.14 shows a circle S inside a box B , and a partition P of B , where the type I subboxes of the partition are shaded. The shaded boxes visibly have a small total area. Similarly, figure 6.15 shows a smooth piece of surface in \mathbf{R}^3 , then shows it inside a partitioned box, and figure 6.16 shows some of the type I subboxes of the partition. Figure 6.16 also shows a smooth arc in \mathbf{R}^3 and some of the type I rectangles that cover it, with the ambient

box and the rest of the partition now tacit. Figure 6.16 is meant to show that all the type I boxes, which combine to cover the surface or the arc, have a small total volume.




Figure 6.14. Circle, box, partition, and type I subboxes




Figure 6.15. A two dimensional set in \mathbf{R}^3 ; the set inside a partitioned box

The following fact is convenient.

Proposition 6.5.2 (Volume Zero Criterion). *A set S contained in the box B has volume zero if and only if for every $\varepsilon > 0$ there exists a partition P of B such that*




Figure 6.16. Some type I subboxes of the partition, and for an arc in \mathbf{R}^3

$$\sum_{J:\text{type I}} \text{vol}(J) < \varepsilon.$$

The proof is an exercise. This criterion makes it plausible that any bounded smooth arc in \mathbf{R}^2 has volume zero, and similarly for a bounded smooth arc or smooth piece of surface in \mathbf{R}^3 . The next result uses the criterion to provide a general class of volume zero sets. Recall that for any set $S \subset \mathbf{R}^n$ and any mapping $\varphi : S \rightarrow \mathbf{R}^m$, the **graph** of φ is a subset of $\mathbf{R}^n \times \mathbf{R}^m$,

$$\text{graph}(\varphi) = \{(x, \varphi(x)) : x \in S\}.$$

Proposition 6.5.3 (Graphs Have Volume Zero). *Let B be a box in \mathbf{R}^n , and let $\varphi : B \rightarrow \mathbf{R}^m$ be continuous. Then $\text{graph}(\varphi)$ has volume zero.*

Proof. The continuity of φ on B is uniform, and the image of φ , being compact, sits in some interval I .

Let $\varepsilon > 0$ be given. Set ε' equal to any positive number less than $\varepsilon/(2\text{vol}(B))$ such that $\text{length}(I)/\varepsilon'$ is an integer. There exists a partition Q of I whose subintervals K have length ε' , and a $\delta > 0$ such that for all $x, \tilde{x} \in B$,

$$|\tilde{x} - x| < \delta \quad \implies \quad |\varphi(\tilde{x}) - \varphi(x)| < \varepsilon'. \quad (6.5)$$

Now take a partition P of B whose subboxes J have sides of length less than δ/n , so that if two points are in a common subbox J then the distance between them is less than δ . Consider the partition $P \times Q$ of $B \times I$. For each subbox J of P there exist at most two subboxes $J \times K$ of $P \times Q$ over J that intersect the graph of φ , i.e., subboxes of type I. To see this, note that if we have three or more such subboxes, then some pair $J \times K$ and $J \times K'$ are not vertical neighbors, and so any hypothetical pair of points of the graph,

one in each subbox, are less than distance δ apart horizontally but at least distance ε' apart vertically. But by (6.5), this is impossible. (See figure 6.17. The horizontal direction in the figure is only a schematic for \mathbf{R}^n , but the vertical direction accurately depicts the one-dimensional codomain of φ .)




Figure 6.17. The graph meets at most two boxes over each base

Now, working with subboxes $J \times K$ of $P \times Q$, compute that

$$\begin{aligned} \sum_{\text{type I}} \text{vol}(J \times K) &= \sum_{\text{type I}} \text{vol}(J) \cdot \varepsilon' && \text{since length}(K) = \varepsilon' \\ &\leq 2 \sum_J \text{vol}(J) \cdot \varepsilon' && \text{by the preceding paragraph} \\ &= 2\text{vol}(B) \cdot \varepsilon' < \varepsilon && \text{since } \varepsilon' < \varepsilon / (2\text{vol}(B)), \end{aligned}$$

and the proof is complete by the Volume Zero Criterion. \square

An exercise shows that any finite union of sets of volume zero also has volume zero, and another exercise shows that any subset of a set of volume zero also has volume zero. These facts and the preceding proposition are enough to demonstrate that many regions have boundaries of volume zero. The boundary of a set consists of all points simultaneously near the set and near its complement—roughly speaking, its edge. (Unfortunately, the mathematical terms *bounded* and *boundary* need have nothing to do with each other. A set with a boundary need not be bounded, and a bounded set need not have any boundary points nor contain any of its boundary points if it does have them.) For example, the set in figure 6.18 has a boundary consisting of four graphs of functions on one-dimensional boxes, i.e., on intervals. Two of the boundary pieces are graphs of functions $y = f(x)$, and the other two are graphs of functions $x = f(y)$. Two of the four functions are constant functions.




Figure 6.18. Boundary with area zero

The main result of this section is that discontinuity on a set of volume zero does not interfere with integration.

Theorem 6.5.4 (Near-continuity Implies Integrability). *Let $B \subset \mathbf{R}^n$ be a box. Let $f : B \rightarrow \mathbf{R}$ be bounded, and continuous except on a set $S \subset B$ of volume zero. Then $\int_B f$ exists.*

Proof. Let $\varepsilon > 0$ be given.

The proof involves two partitions. Since f is bounded there exists a positive number R such that $|f(x)| < R$ for all $x \in B$. Take a partition P of B whose subboxes J of type I (those intersecting the set S where f is discontinuous) have volumes adding to less than $\varepsilon/(4R)$. (See figure 6.19, in which the function f is the dome shape over the unit disk but is 0 outside the unit disk, making it discontinuous on the unit circle.)




Figure 6.19. Type I subboxes of small total area

Consider some yet unspecified refinement P' of P , dividing each subbox J of P into further subboxes J' . (See figure 6.20, in which the original boxes J of type I remain shaded, but each box J of either type has been further partitioned.) On any J' , $M_{J'}(f) - m_{J'}(f) \leq 2R$, and so a short calculation shows that regardless how the refinement P' is to be specified, its subboxes J' that sit inside type I subboxes J of P make only a small contribution to the difference between the lower sum and the upper sum of f over P' ,

$$\begin{aligned} & \sum_{J: \text{type I}} \sum_{J' \subset J} (M_{J'}(f) - m_{J'}(f)) \text{vol}(J') \\ & \leq 2R \sum_{J: \text{type I}} \sum_{J' \subset J} \text{vol}(J') = 2R \sum_{J: \text{type I}} \text{vol}(J) < 2R \frac{\varepsilon}{4R} = \frac{\varepsilon}{2}. \end{aligned} \quad (6.6)$$




Figure 6.20. Refinement of the partition

To specify the refinement P' of P that we need, consider the type II subboxes J of P , i.e., the union of the unshaded boxes in figure 6.19. The function f is continuous on each such subbox and hence integrable over it by Theorem 6.3.1. Let the number of these subboxes be denoted N . By (\implies) of the Integrability Criterion, each type II subbox J has a partition P'_J such that

$$U(f, P'_J) - L(f, P'_J) < \frac{\varepsilon}{2N}.$$

Let P' be a partition of the full box B that refines the original partition P and also incorporates all the partitions P'_J of the type II subboxes J . Thus the intersection of P' with any particular type II subbox J refines P'_J . Since refinement can not increase the distance between lower and upper sums, another short calculation shows that the subboxes J' of P' that sit inside type II subboxes J of P also make only a small contribution to the difference between the lower sum and the upper sum of f over P' ,

$$\begin{aligned} \sum_{J : \text{type II}} \sum_{J' \subset J} (M_{J'}(f) - m_{J'}(f)) \text{vol}(J') \\ \leq \sum_{J : \text{type II}} U(f, P'_J) - L(f, P'_J) < N \cdot \frac{\varepsilon}{2N} = \frac{\varepsilon}{2}. \end{aligned} \quad (6.7)$$

Finally, combining (6.6) and (6.7) shows that $U(f, P') - L(f, P') < \varepsilon$, and so by (\Leftarrow) of the Integrability Criterion, $\int_B f$ exists. \square

To recapitulate the argument: The fact that f is bounded means that its small set of discontinuities can't cause much difference between lower and upper sums, and the continuity of f on the rest of its domain poses no problem either. The only difficulty was making the ideas fit into our box-counting definition of the integral. The reader could well object that proving Theorem 6.5.4 shouldn't be this hard. Indeed, the theory of integration being presented in this chapter, *Riemann integration*, involves laborious proofs precisely because it uses such crude technology: finite sums over boxes. More powerful theories of integration exist, with stronger theorems and more graceful proofs. However, these theories also entail the startup cost of assimilating a larger, more abstract set of working ideas, making them inappropriate for a course at this level.

Now we can discuss integration over nonboxes.

Definition 6.5.5 (Known-Integrable Function). *A function*

$$f : K \longrightarrow \mathbf{R}$$

is known-integrable if K is a compact subset of \mathbf{R}^n having boundary of volume zero, and if f is bounded and is continuous off a subset of K having boundary zero.

Definition 6.5.6 (Integral Over a Nonbox). *Let*

$$f : K \longrightarrow \mathbf{R}$$

be a known-integrable function. Extend its domain to \mathbf{R}^n by defining a new function

$$\tilde{f} : \mathbf{R}^n \longrightarrow \mathbf{R}, \quad \tilde{f}(x) = \begin{cases} f(x) & \text{if } x \in K, \\ 0 & \text{if } x \notin K. \end{cases}$$

Then the integral of f over K is

$$\int_K f = \int_B \tilde{f} \quad \text{where } B \text{ is any box containing } K.$$

This integral on the right side of the equality in the definition exists because \tilde{f} is bounded and discontinuous on a set of volume zero, as required for Theorem 6.5.4. In particular, the definition of volume is now, sensibly enough,

$$\text{vol}(K) = \int_K 1.$$

Naturally, the result of Proposition 6.2.4, that the integral over the whole is the sum of the integrals over the pieces, is not particular to boxes and subboxes.

Proposition 6.5.7. *Let $K \subset \mathbf{R}^n$ be a compact set whose boundary has volume zero. Let $f : K \rightarrow \mathbf{R}$ be continuous. Further, let $K = K_1 \cup K_2$ where each K_j is compact and the intersection $K_1 \cap K_2$ has volume zero. Then f is integrable over K_1 and K_2 , and*

$$\int_{K_1} f + \int_{K_2} f = \int_K f.$$

Proof. Define

$$f_1 : K \rightarrow \mathbf{R}, \quad f_1(x) = \begin{cases} f(x) & \text{if } x \in K_1, \\ 0 & \text{otherwise.} \end{cases}$$

Then f_1 is nearly continuous on K , and so $\int_K f_1$ exists and equals $\int_{K_1} f_1$. Define a corresponding function $f_2 : K \rightarrow \mathbf{R}$, for which the corresponding conclusions hold. It follows that

$$\int_{K_1} f_1 + \int_{K_2} f_2 = \int_K f_1 + \int_K f_2 = \int_K (f_1 + f_2).$$

But $f_1 + f_2$ equals f except on the volume-zero set $K_1 \cap K_2$, which contributes nothing to the integral. The result follows. \square

Exercises

6.5.1. (a) Suppose that $I_1 = [a_1, b_1]$, $I_2 = [a_2, b_2]$, \dots are intervals in \mathbf{R} . Show that their intersection $I_1 \cap I_2 \cap \dots$ is another interval (possibly empty).

(b) Suppose that $S = S_1 \times \dots \times S_n$, $T = T_1 \times \dots \times T_n$, $U = U_1 \times \dots \times U_n$, \dots are cartesian products of sets. Show that their intersection is

$$S \cap T \cap U \cap \dots = (S_1 \cap T_1 \cap U_1 \cap \dots) \times \dots \times (S_n \cap T_n \cap U_n \cap \dots).$$

(c) Show that any intersection of boxes in \mathbf{R}^n is another box (possibly empty).

(d) If S is a set and T_1, T_2, T_3, \dots are all sets that contain S , show that $T_1 \cap T_2 \cap T_3 \cap \dots$ contains S .

6.5.2. Let S be a nonempty bounded subset of \mathbf{R}^n , let B be any box containing S , and let B' be the intersection of all boxes containing S . By the preceding problem, B' is also a box containing S . Use Proposition 6.2.4 to show that if either of $\int_B \chi_S$ and $\int_{B'} \chi_S$ exist then both exist and they are equal. It follows, as remarked in the text, that the definition of the volume of S is independent of the containing box B .

6.5.3. Let $B \subset \mathbf{R}^n$ be a box. Show that its volume under Definition 6.5.1 equals its volume under Definition 6.1.4. (Hint: Exercise 6.2.3.)

6.5.4. Let S be the set of rational numbers in $[0, 1]$. Show that S does not have a volume (i.e., a length) under Definition 6.5.1.

6.5.5. Prove the Volume Zero Criterion.

6.5.6. If $S \subset \mathbf{R}^n$ has volume zero and R is a subset of S , show R has volume zero. (Hint: $0 \leq \chi_R \leq \chi_S$.)

6.5.7. Prove that if S_1, \dots, S_N have volume zero, then so does $S_1 \cup \dots \cup S_N$. (Hints: Induction on N ; $\chi_{S_1 \cup S_2} \leq \chi_{S_1} + \chi_{S_2}$.)

6.5.8. Find an unbounded set with a nonempty boundary, and a bounded set with empty boundary.

6.5.9. Use Theorem 6.5.4, the discussion immediately after its proof, Proposition 6.5.3, and any other results necessary to explain why for each set K and function $f : K \rightarrow \mathbf{R}$ below, the integral $\int_K f$ exists.

- (a) $K = \{(x, y) : 2 \leq y \leq 3, 0 \leq x \leq 1 + \log y/y\}$, $f(x, y) = e^{xy}$.
- (b) $K = \{(x, y) : 1 \leq x \leq 4, 1 \leq y \leq \sqrt{x}\}$, $f(x, y) = e^{x/y^2}/y^5$.
- (c) $K =$ the region between the curves $y = 2x^2$ and $x = 4y^2$, $f(x, y) = 1$.
- (d) $K = \{(x, y) : 1 \leq x^2 + y^2 \leq 2\}$, $f(x, y) = x^2$.
- (e) $K =$ the pyramid with vertices $(0, 0, 0)$, $(3, 0, 0)$, $(0, 3, 0)$, $(0, 0, 3/2)$, $f(x, y, z) = x$.
- (f) $K = \{x \in \mathbf{R}^n : |x| \leq 1\}$ (the solid unit ball in \mathbf{R}^n), $f(x_1, \dots, x_n) = x_1 \cdots x_n$.

6.6 Fubini's Theorem

With existence theorems for the integral now in hand, this section and the next one present tools to compute integrals.

An **n -fold iterated integral** is n one-dimensional integrals nested inside each other, such as

$$\int_{x_1=a_1}^{b_1} \int_{x_2=a_2}^{b_2} \cdots \int_{x_n=a_n}^{b_n} f(x_1, x_2, \dots, x_n),$$

for some function $f : [a_1, b_1] \times \cdots \times [a_n, b_n] \rightarrow \mathbf{R}$. This is definitely not the same sort of object as an n -dimensional integral. We can evaluate the iterated integral by working from the inside out. For the innermost integral, f is to be viewed as a function of the variable x_n with its other inputs treated as constants, and so on outwards. For example,

$$\int_{x=0}^1 \int_{y=0}^2 xy^2 = \int_{x=0}^1 \left. \frac{1}{3}xy^3 \right|_{y=0}^2 = \int_{x=0}^1 \frac{8}{3}x = \left. \frac{4}{3}x^2 \right|_{x=0}^1 = \frac{4}{3}.$$

There are $n!$ different orders in which one can iterate n integrals, e.g., the example just worked is not the same object as $\int_{y=0}^2 \int_{x=0}^1 xy^2$. Regardless of order, each one-dimensional integral requires varying its particular input to f while holding the other inputs fixed. The upshot of all this variable-dependence is that there is no reasonable alternative to naming and writing the variables in an iterated integral.

In an inner integral, outer variables may figure not only as inputs to the integrand, but also in the limits of integration. For example, in the calculation

$$\int_{x=0}^{\pi} \int_{y=0}^x \cos(x+y) = \int_{x=0}^{\pi} \sin(x+y) \Big|_{y=0}^x = \int_{x=0}^{\pi} \sin(2x) - \sin(x) = -2,$$

each inner integral over y is being taken over a segment of x -dependent length as the outer variable x varies from 0 to π . (See figure 6.21.)




Figure 6.21. Variable Range of Inner Integration

Fubini's Theorem says that under suitable conditions, the n -dimensional integral is equal to the n -fold iterated integral. This provides an essential calculational tool for multivariable integration.

Theorem 6.6.1 (Fubini's Theorem). *Let $B = [a, b] \times [c, d] \subset \mathbf{R}^2$, and let $f : B \rightarrow \mathbf{R}$ be bounded, and continuous except on a subset $S \subset B$ of area zero, so $\int_B f$ exists. Suppose that for each $x \in [a, b]$, S contains only finitely many points (possibly none) with first coordinate x . Then the iterated integral $\int_{x=a}^b \int_{y=c}^d f(x, y)$ also exists, and*

$$\int_B f = \int_{x=a}^b \int_{y=c}^d f(x, y).$$

For notational convenience, the theorem is stated only in two dimensions. Replacing $[a, b]$ and $[c, d]$ by boxes gives a more general version with a virtually identical proof. Thinking geometrically in terms of area and volume makes the




Figure 6.22. Inner integral as cross-sectional area

theorem plausible in two dimensions since each inner integral is the area of a cross-section of the volume under the graph of f . (See figure 6.22.)

However, since the multiple integral and the iterated integral are defined analytically as limits of sums, the only available tactic for proving the theorem is analytic: we must compare approximating sums for the two integrals. A lower sum for the integral $\int_B f$ is shown geometrically in the left side of figure 6.23. A partition $P \times Q$ divides the box $B = [a, b] \times [c, d]$ into subboxes $I \times J$, and the volume of each solid region in the figure is the area of a subbbox times the minimum height of the graph over the subbbox. By contrast, letting $g(x) = \int_{y=c}^d f(x, y)$ be the area of the cross-section at x , the right side of figure 6.23 shows a lower sum for the integral $\int_{x=a}^b g(x)$. The partition P divides the interval $[a, b]$ into subintervals I , and the volume of each bread-slice in the figure is the length of a subinterval times the minimum area of the cross-sections orthogonal to I . Because integrating in the y -direction is a finer diagnostic than summing minimal box-areas in the y -direction, the bread-slices in the right side of the figure are a superset of the boxes in the left side. Consequently, the volume beneath the bread-slices is at least the volume of the boxes,

$$L(f, P \times Q) \leq L(g, P).$$

By similar reasoning for upper sums, in fact we expect that

$$L(f, P \times Q) \leq L(g, P) \leq U(g, P) \leq U(f, P \times Q). \quad (6.8)$$

Since $L(f, P \times Q)$ and $U(f, P \times Q)$ converge to $\int_B f$ under suitable refinement of $P \times Q$, so do $L(g, P)$ and $U(g, P)$. Thus the iterated integral exists and equals the double integral as desired. The details of turning the geometric intuition of this paragraph into a proof of Fubini's Theorem work out fine provided that we carefully tend to matters in just the right order. However,

the need for care is genuine. A subtle point not illustrated by figure 6.23 is that

- although the boxes lie entirely beneath the bread-slices (this is a relation between two sets),
- and although the boxes lie entirely beneath the graph (so is this),
- and although the volume of the bread-slices is at most the volume beneath the graph (but this is a relation between two numbers),
- the bread-slices need not lie entirely beneath the graph.

Since the bread-slices need not lie entirely beneath the graph, the fact that their volume $L(g, P)$ estimates the integral $\int_B f$ from below does not follow from pointwise considerations. The proof finesses this point by establishing the inequalities (6.8) without reference to the integral, only then bringing the integral into play as the limit of the extremal sums in (6.8).

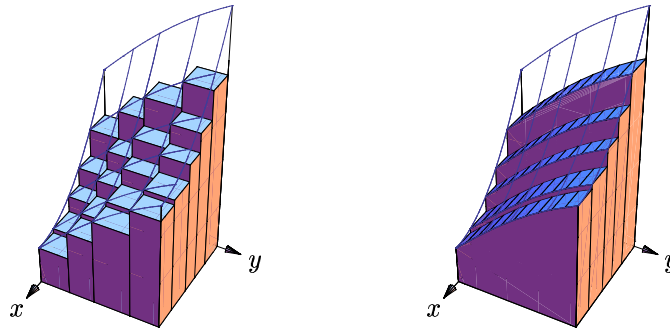


Figure 6.23. Geometry of two lower sums

Proof. For each $x \in [a, b]$, define the cross-sectional function

$$\varphi_x : [c, d] \longrightarrow \mathbf{R}, \quad \varphi_x(y) = f(x, y).$$

The hypotheses of Fubini's Theorem ensure that as x varies from a to b , each cross-sectional function φ_x is continuous except at finitely many points and hence it is integrable on $[c, d]$. Give the cross-sectional integral a name,

$$g : [a, b] \longrightarrow \mathbf{R}, \quad g(x) = \int_c^d \varphi_x.$$

The iterated integral $\int_{x=a}^b \int_{y=c}^d f(x, y)$ is precisely the integral $\int_a^b g$. We need to show that this exists and equals $\int_B f$.

Consider any partition $P \times Q$ of B into subboxes $J \times K$. Thus P partitions $[a, b]$ into subintervals J , and Q partitions $[c, d]$ into subintervals K . Take any subinterval J of $[a, b]$, and take any point x of J . The lower sum of the cross-sectional function φ_x over the y -partition Q is a lower bound for the cross-sectional integral $g(x)$,

$$\sum_K m_K(\varphi_x) \text{length}(K) = L(\varphi_x, Q) \leq \int_c^d \varphi_x = g(x).$$

Note that φ_x on each K samples f only on a cross-section of $J \times K$, and so f has more opportunity to be small on $J \times K$ than φ_x has on K . That is,

$$m_{J \times K}(f) \leq m_K(\varphi_x).$$

This previous two displays combine to give a lower bound for the cross-sectional integral $g(x)$ that depends on the interval J but is independent of the point x of J ,

$$\sum_K m_{J \times K}(f) \text{length}(K) \leq g(x) \quad \text{for all } x \in J.$$

That is, the left side of this last display is a lower bound of all values $g(x)$ where $x \in J$. So it is at most the greatest lower bound,

$$\sum_K m_{J \times K}(f) \text{length}(K) \leq m_J(g).$$

Multiply through by the length of J to get

$$\sum_K m_{J \times K}(f) \text{area}(J \times K) \leq m_J(g) \text{length}(J).$$

(This inequality says that each y -directional row of boxes in the left half of figure 6.23 has at most the volume of the corresponding bread-slice in the right half of the figure.) As noted at the end of the preceding paragraph, the iterated integral is the integral of g . The estimate just obtained puts us in a position to compare lower sums for the double integral and the iterated integral,

$$L(f, P \times Q) = \sum_{J, K} m_{J \times K}(f) \text{area}(J \times K) \leq \sum_J m_J(g) \text{length}(J) = L(g, P).$$

Concatenating a virtually identical argument with upper sums gives the anticipated chain of inequalities,

$$L(f, P \times Q) \leq L(g, P) \leq U(g, P) \leq U(f, P \times Q).$$

The outer terms converge to $\int_B f$ under suitable refinement of $P \times Q$, and hence so do the inner terms, showing that $\int_a^b g$ exists and equals $\int_B f$. \square

Since we will use Fubini's Theorem to evaluate actual examples, all the notational issues discussed in section 6.4 arise here again. A typical notation for examples is

$$\int_B f(x, y) = \int_{x=a}^b \int_{y=c}^d f(x, y),$$

where the left side is a 2-dimensional integral, the right side is an iterated integral, and $f(x, y)$ is an expression defining f . For example, by Fubini's Theorem and the calculation at the beginning of this section,

$$\int_{[0,1] \times [0,2]} xy^2 = \int_{x=0}^1 \int_{y=0}^2 xy^2 = \frac{4}{3}.$$

Of course, an analogous theorem asserts that $\int_B f(x, y) = \int_{y=c}^d \int_{x=a}^b f(x, y)$ provided that the set S of discontinuity meets horizontal segments at only finitely many points too. In other words, the double integral also equals the other iterated integral, and consequently the two iterated integrals agree. For example, $\int_{y=0}^2 \int_{x=0}^1 xy^2$ also works out easily to $4/3$.

In many applications, the integral over B is really an integral over a non-rectangular compact set K , as defined at the end of the previous section. If K is the area between the graphs of continuous functions $\varphi_1, \varphi_2 : [a, b] \rightarrow \mathbf{R}$, i.e., if

$$K = \{(x, y) : a \leq x \leq b, \varphi_1(x) \leq y \leq \varphi_2(x)\},$$

then one iterated integral takes the form $\int_{x=a}^b \int_{y=\varphi_1(x)}^{\varphi_2(x)} f(x, y)$. Similarly, if

$$K = \{(x, y) : c \leq y \leq d, \theta_1(y) \leq x \leq \theta_2(y)\},$$

then the other iterated integral is $\int_{y=c}^d \int_{x=\theta_1(y)}^{\theta_2(y)} f(x, y)$. (See figure 6.24.)

The interchangeability of the order of integration leads to a fiendish class of iterated integral problems where one switches order to get a workable integrand. For example, the iterated integral

$$\int_{y=0}^2 \int_{x=y/2}^1 e^{-x^2}$$

looks daunting because the integrand e^{-x^2} has no convenient antiderivative, but after exchanging the order of the integrations and then carrying out a change of variable, it becomes

$$\int_{x=0}^1 \int_{y=0}^{2x} e^{-x^2} = \int_{x=0}^1 2xe^{-x^2} = \int_{u=0}^1 e^{-u} = 1 - e^{-1}.$$




Figure 6.24. Setting up nonrectangular double integrals

Interchanging the order of integration can be tricky in such cases; often one has to break K up into several pieces first, e.g.,

$$\int_{x=1}^2 \int_{y=1/x}^2 = \int_{y=1/2}^1 \int_{x=1/y}^2 + \int_{y=1}^2 \int_{x=1}^2 .$$

A carefully labeled diagram facilitates this process. For example, figure 6.25 shows the sketch that arises from the integral on the left side, and then the resulting sketch that leads to the sum of two integrals on the right side.




Figure 6.25. Sketches for iterated integrals

Interchanging the outer two integrals in a triply iterated integral is no different from the double case, but interchanging the inner two is tricky because of the constant-but-unknown value taken by the outer variable. Sketching a

generic two-dimensional cross-section usually makes the substitutions clear. For example, consider the iterated integral

$$\int_{x=0}^1 \int_{y=x^3}^{x^2} \int_{z=y}^{x^2} . \quad (6.9)$$

(The function being integrated is irrelevant to this discussion of how to exchange the order of integration, so it is omitted from the notation.) Exchanging the outer two integrals is carried out via the first diagram in figure 6.26. The diagram leads to the iterated integral

$$\int_{y=0}^1 \int_{x=\sqrt[3]{y}}^{\sqrt{y}} \int_{z=y}^{x^2} .$$

On the other hand, to exchange the inner integrals of (6.9), think of x as fixed but generic between 0 and 1 and consider the second diagram in figure 6.26. This diagram shows that (6.9) is also the iterated integral

$$\int_{x=0}^1 \int_{z=x^3}^{x^2} \int_{y=x^3}^z . \quad (6.10)$$




Figure 6.26. Sketches for a triply-iterated integral

Switching the outermost and innermost integrals of (6.9) while leaving the middle one in place requires three successive switches of adjacent integrals. For instance, switching the inner integrals as we just did and then doing an outer exchange on (6.10) virtually identical to the outer exchange of a moment earlier (substitute z for y in the first diagram of figure 6.26) shows that (6.9) is also

$$\int_{z=0}^1 \int_{x=\sqrt{z}}^{\sqrt[3]{z}} \int_{y=x^3}^z .$$

Finally, the first diagram of figure 6.27 shows how to exchange the inner integrals once more. The result is

$$\int_{z=0}^1 \int_{y=z^{3/2}}^z \int_{x=\sqrt{z}}^{\sqrt[3]{y}} .$$

The second diagram of figure 6.27 shows the three-dimensional figure that our iterated integral has traversed in various fashions. It is satisfying to see how this picture is compatible with the cross-sectional sketches, and to determine which axis is which. However, the three-dimensional figure is unnecessary for exchanging the order of integration. The author of these notes finds using two-dimensional cross-sections easier and more reliable than trying to envision an entire volume at once. Also, the two-dimensional cross-section technique will work in an n -fold iterated integral for any $n \geq 3$, even when the whole situation is hopelessly beyond visualizing.




Figure 6.27. Another cross-section and the three-dimensional region

The **unit simplex in \mathbf{R}^3** is the set

$$S = \{(x, y, z) : x \geq 0, y \geq 0, z \geq 0, x + y + z \leq 1\}$$

(see figure 6.28). Its **centroid** is $(\bar{x}, \bar{y}, \bar{z})$, where

$$\bar{x} = \frac{\int_S x}{\text{vol}(S)}, \quad \bar{y} = \frac{\int_S y}{\text{vol}(S)}, \quad \bar{z} = \frac{\int_S z}{\text{vol}(S)}.$$

Fubini's Theorem lets us treat the integrals as iterated, giving

$$\begin{aligned}
\int_S x &= \int_{x=0}^1 \int_{y=0}^{1-x} \int_{z=0}^{1-x-y} x \\
&= \int_{x=0}^1 \int_{y=0}^{1-x} x(1-x-y) \\
&= \int_{x=0}^1 \frac{1}{2}x(1-x)^2 = \frac{1}{24},
\end{aligned}$$

where the routine one-variable calculations are not shown in detail. Similarly, $\text{vol}(S) = \int_S 1$ works out to $1/6$, so $\bar{x} = 1/4$. By symmetry, $\bar{y} = \bar{z} = 1/4$ also. See the exercises for an n -dimensional generalization of this result.




Figure 6.28. Unit simplex

To find the volume between the two paraboloids $z = 8 - x^2 - y^2$ and $z = x^2 + 3y^2$, first set $8 - x^2 - y^2 = x^2 + 3y^2$ to find that the graphs intersect over the ellipse $\{(x, y) : (x/2)^2 + (y/\sqrt{2})^2 = 1\}$. (See figure 6.29.) By Fubini's Theorem the volume is

$$V = \int_{y=-\sqrt{2}}^{\sqrt{2}} \int_{x=-\sqrt{4-2y^2}}^{\sqrt{4-2y^2}} \int_{z=x^2+3y^2}^{8-x^2-y^2} 1 = \pi 8\sqrt{2}$$

where again the one-dimensional calculations are omitted.

Another example is to find the volume of the region K common to the cylinders $x^2 + y^2 = 1$ and $x^2 + z^2 = 1$. For each x -value between -1 and 1 , y and z vary independently between $-\sqrt{1-x^2}$ and $\sqrt{1-x^2}$. That is, the intersection of the two cylinders is a union of squares, bounded by two tilted circles. (See figure 6.30.) By the methods of this section, the integral has the same value as the iterated integral, which is

$$\int_{x=-1}^1 \int_{y=-\sqrt{1-x^2}}^{\sqrt{1-x^2}} \int_{z=-\sqrt{1-x^2}}^{\sqrt{1-x^2}} 1 = 4 \int_{x=-1}^1 (1-x^2) = \frac{16}{3}.$$

Finally, we end the section with a more theoretical example.




Figure 6.29. Volume between two graphs




Figure 6.30. Volume common to two cylinders

Proposition 6.6.2 (Differentiation under the Integral Sign). *Consider a function*

$$f : [a, b] \times [c, d] \longrightarrow \mathbf{R}.$$

Suppose that f and $D_1 f$ are continuous. Also consider the cross-sectional integral function,

$$g : [a, b] \longrightarrow \mathbf{R}, \quad g(x) = \int_{y=c}^d f(x, y).$$

Then g is differentiable, and $g'(x) = \int_{y=c}^d D_1 f(x, y)$. That is,

$$\frac{d}{dx} \int_{y=c}^d f(x, y) = \int_{y=c}^d \frac{\partial}{\partial x} f(x, y).$$

Proof. Compute for any $x \in [a, b]$, using the Fundamental Theorem of Integral Calculus (Theorem 6.4.2) for the second equality and then Fubini's Theorem for the fourth,

$$\begin{aligned} g(x) &= \int_{y=c}^d f(x, y) \\ &= \int_{y=c}^d \left(\int_{t=a}^x D_1 f(t, y) + f(a, y) \right) \\ &= \int_{y=c}^d \int_{t=a}^x D_1 f(t, y) + C \quad \text{where } C = \int_{y=c}^d f(a, y) \\ &= \int_{t=a}^x \int_{y=c}^d D_1 f(t, y) + C. \end{aligned}$$

It follows from Theorem 6.4.1 that the derivative equals the integrand evaluated at x ,

$$g'(x) = \int_{y=c}^d D_1 f(x, y),$$

as desired. \square

See exercise 6.6.10 for another example in this spirit.

Exercises

6.6.1. Let S be the set of points $(x, y) \in \mathbf{R}^2$ between the x -axis and the sine curve as x varies between 0 and 2π . Since the sine curve has two arches between 0 and 2π , and since the area of an arch of sine is 2,

$$\int_S 1 = 4.$$

On the other hand,

$$\int_{x=0}^{2\pi} \int_{y=0}^{\sin x} 1 = \int_{x=0}^{2\pi} \sin x = 0.$$

Why doesn't this contradict Fubini's Theorem?

6.6.2. Exchange the order of integration in $\int_{x=a}^b \int_{y=a}^x f(x, y)$.

6.6.3. Exchange the inner order of integration in $\int_{x=0}^1 \int_{y=0}^{1-x} \int_{z=0}^{x+y} f$.

6.6.4. Exchange the inner order of integration in $\int_{x=0}^1 \int_{y=0}^1 \int_{z=0}^{x^2+y^2} f$. Sketch the region of integration.

6.6.5. Evaluate $\int_K f$ from in parts (a), (b), (c), (f) of exercise 6.5.9.

6.6.6. Find the volume of the region K in the first octant bounded by $x = 0$, $z = 0$, $z = y$, and $x = 4 - y^2$. Sketch K .

6.6.7. Evaluate $\int_K (1 + x + y + z)^{-3}$ where K is the unit simplex.

6.6.8. Find the volume of the region K bounded by the coordinate planes, $x + y = 1$, and $z = x^2 + y^2$. Sketch K .

6.6.9. Find the volume of the region K between $z = x^2 + 9y^2$ and $z = 18 - x^2 - 9y^2$. Sketch K .

6.6.10. Let $f : \mathbf{R}^2 \rightarrow \mathbf{R}$ have continuous mixed second order partial derivatives, i.e., let $D_{12}f$ and $D_{21}f$ exist and be continuous. Rederive the familiar fact that $D_{12}f = D_{21}f$ as follows. If $D_{12}f(p, q) - D_{21}f(p, q) > 0$ at some point (p, q) then $D_{12}f - D_{21}f > 0$ on some rectangle $B = [a, b] \times [c, d]$ containing (p, q) , so $\int_B (D_{12}f - D_{21}f) > 0$. Obtain a contradiction by evaluating this integral.

6.6.11. Let K and L be compact subsets of \mathbf{R}^n with boundaries of volume zero. Suppose that for each $x_1 \in \mathbf{R}$, the cross sectional sets

$$\begin{aligned} K_{x_1} &= \{(x_2, \dots, x_n) : (x_1, x_2, \dots, x_n) \in K\} \\ L_{x_1} &= \{(x_2, \dots, x_n) : (x_1, x_2, \dots, x_n) \in L\} \end{aligned}$$

have equal $(n - 1)$ -dimensional volumes. Show that K and L have the same volume. Illustrate for $n = 2$.

6.6.12. Let x_0 be a positive real, and let $f : [0, x_0] \rightarrow \mathbf{R}$ be continuous. Show that

$$\int_{x_1=0}^{x_0} \int_{x_2=0}^{x_1} \cdots \int_{x_n=0}^{x_{n-1}} f(x_n) = \frac{1}{(n-1)!} \int_{t=0}^{x_0} (x_0 - t)^{n-1} f(t).$$

(Use induction. The base case $n = 1$ is easy, then the induction hypothesis applies to the inner $(n - 1)$ -fold integral.)

6.6.13. Let $n \in \mathbf{Z}^+$ and $r \in \mathbf{R}_{\geq 0}$. The n -dimensional simplex of side r is

$$S_n(r) = \{(x_1, \dots, x_n) : 0 \leq x_1, \dots, 0 \leq x_n, x_1 + \cdots + x_n \leq r\}.$$

(a) Make sketches for $n = 2$ and $n = 3$ showing that $S_n(r)$ is a union of cross-sectional $(n - 1)$ -dimensional simplices of side $r - x_n$ at height x_n as x_n varies from 0 to r . Prove this symbolically for general $n > 1$. That is, prove that

$$S_n(r) = \{S_{n-1}(r - x_n) \times \{x_n\} : 0 \leq x_n \leq r\}.$$

(b) Prove that $\text{vol}(S_1(r)) = r$. Use part (a) and Fubini's Theorem to prove that

$$\text{vol}(S_n(r)) = \int_{x_n=0}^r \text{vol}(S_{n-1}(r - x_n)) \quad \text{for } n > 1,$$

and show by induction that $\text{vol}(S_n(r)) = r^n/n!$.

(c) Use Fubini's Theorem to show that $\int_{S_n(r)} x_n = \int_{x_n=0}^r x_n \frac{(r - x_n)^{n-1}}{(n-1)!}$.

Work this integral by parts to get $\int_{S_n(r)} x_n = r^{n+1}/(n+1)!$.

(d) The centroid of $S_n(r)$ is $(\bar{x}_1, \dots, \bar{x}_n)$, where $\bar{x}_j = \int_{S_n(r)} x_j / \text{vol}(S_n(r))$ for each j . What are these coordinates explicitly? (Make sure your answer agrees with the case in the text.)

6.7 Change of Variable

Any point $p \in \mathbf{R}^2$ with cartesian coordinates (x, y) is also specified by its **polar coordinates** (r, θ) , where r is the distance from the origin to p and θ is the angle from the positive x -axis to p . (See figure 6.31.)




Figure 6.31. Polar coordinates

The angle θ is defined only up to multiples of 2π , and it isn't defined at all when $p = (0, 0)$. Trigonometry expresses (x, y) in terms of (r, θ) ,

$$x = r \cos \theta, \quad y = r \sin \theta. \quad (6.11)$$

But expressing (r, θ) in terms of (x, y) is a little more subtle. Certainly

$$r = \sqrt{x^2 + y^2}.$$

Also, $\tan \theta = y/x$ provided that $x \neq 0$, but this doesn't mean that $\theta = \arctan(y/x)$. Indeed, \arctan isn't even a well-defined function until its range

is specified, e.g., as $(-\pi/2, \pi/2)$. With this particular restriction, the actual formula for θ , even given that not both x and y are 0, is not $\arctan(y/x)$, but

$$\theta = \begin{cases} \arctan(y/x) & \text{if } x > 0 \text{ and } y \geq 0 \text{ (this lies in } [0, \pi/2)), \\ \pi/2 & \text{if } x = 0 \text{ and } y > 0, \\ \arctan(y/x) + \pi & \text{if } x < 0 \text{ (this lies in } (\pi/2, 3\pi/2)), \\ 3\pi/2 & \text{if } x = 0 \text{ and } y < 0, \\ \arctan(y/x) + 2\pi & \text{if } x > 0 \text{ and } y < 0 \text{ (this lies in } (3\pi/2, 2\pi)). \end{cases}$$

The formula is unwieldy, to say the least. (The author probably would not read through the whole thing if he were instead a reader. In any case, see figure 6.32.) A better approach is that given (x, y) , the polar radius r is the unique nonnegative number such that

$$r^2 = x^2 + y^2,$$

and then, if $r \neq 0$, the polar angle θ is the unique number in $[0, 2\pi)$ such that (6.11) holds. But still, going from polar coordinates (r, θ) to cartesian coordinates (x, y) as in (6.11) is considerably more convenient than conversely. This is good since, as we will see, doing so is also more natural.




Figure 6.32. The angle θ between 0 and 2π

The change of variable mapping from polar to cartesian coordinates is

$$\Phi : \mathbf{R}_{\geq 0} \times [0, 2\pi] \longrightarrow \mathbf{R}^2, \quad \Phi(r, \theta) = (r \cos \theta, r \sin \theta).$$

The mapping is injective except that the half-lines $\mathbf{R}_{\geq 0} \times \{0\}$ and $\mathbf{R}_{\geq 0} \times \{2\pi\}$ both map to the nonnegative x -axis, and the vertical segment $\{0\} \times [0, 2\pi]$ is squashed to the point $(0, 0)$. Each horizontal half-line $\mathbf{R}_{\geq 0} \times \{\theta\}$ maps to the ray of angle θ with the positive x -axis, and each vertical segment $\{r\} \times [0, 2\pi]$ maps to the circle of radius r . (See figure 6.33.)




Figure 6.33. The polar coordinate mapping

It follows that regions in the (x, y) -plane defined by radial or angular constraints are images under Φ of (r, θ) -regions defined by rectangular constraints. For example, the cartesian disk

$$D_b = \{(x, y) : x^2 + y^2 \leq b^2\}$$

is the Φ -image of the polar rectangle

$$R_b = \{(r, \theta) : 0 \leq r \leq b, 0 \leq \theta \leq 2\pi\}.$$

(See figure 6.34.) Similarly the cartesian annulus and quarter disk,

$$A_{a,b} = \{(x, y) : a^2 \leq x^2 + y^2 \leq b^2\},$$

$$Q_b = \{(x, y) : x \geq 0, y \geq 0, x^2 + y^2 \leq b^2\},$$

are the images of rectangles. (See figures 6.35 and 6.36.)




Figure 6.34. Rectangle to disk under the polar coordinate mapping

Iterated integrals over rectangles are especially convenient to evaluate, because the limits of integration for the two one-variable integrals are constants rather than variables that interact. For example,

$$\int_{r=a}^b \int_{\theta=0}^{2\pi} = \int_{\theta=0}^{2\pi} \int_{r=a}^b .$$




Figure 6.35. Rectangle to annulus under the polar coordinate mapping




Figure 6.36. Rectangle to quarter disk under the polar coordinate mapping

These tidy (r, θ) limits describe the (x, y) annulus $A_{a,b}$ indirectly via Φ , while the more direct approach of an (x, y) -iterated integral over $A_{a,b}$ requires four messy pieces,

$$\int_{x=-b}^{-a} \int_{y=-\sqrt{b^2-x^2}}^{\sqrt{b^2-x^2}} + \int_{x=-a}^a \left[\int_{y=-\sqrt{a^2-x^2}}^{-\sqrt{b^2-x^2}} + \int_{y=\sqrt{a^2-x^2}}^{\sqrt{b^2-x^2}} \right] + \int_{x=a}^b \int_{y=-\sqrt{b^2-x^2}}^{\sqrt{b^2-x^2}}.$$

Thus, since Fubini's Theorem equates integrals over two-dimensional regions to twofold iterated integrals, it would be a real convenience to reduce integrating over the (x, y) -annulus to integrating over the (r, θ) rectangle that maps to it under Φ . The Change of Variable Theorem will do so. This is the sense in which it is natural to map from polar to cartesian coordinates rather than in the other direction.

The Change of Variable Theorem says in some generality how to transform an integral from one coordinate system to another. Recall that given a set $A \subset \mathbf{R}^n$ and a differentiable mapping $\Phi : A \rightarrow \mathbf{R}^n$, the n -by- n matrix of partial derivatives of Φ is denoted Φ' ,

$$\Phi' = [D_j \Phi_i]_{i,j=1 \dots n}.$$

A differentiable mapping whose partial derivatives are all continuous is called a \mathcal{C}^1 -mapping. Also, for any set $K \subset \mathbf{R}^n$, an **interior point** of K is a point

of K that is not a boundary point, and the **interior** of K is the set of all such points,

$$K^\circ = \{\text{interior points of } S\}.$$

We will discuss boundary points and interior points more carefully in the next section. In the specific sorts of examples that arise in calculus, they are easy enough to recognize.

Theorem 6.7.1 (Change of Variable Theorem for Multiple Integrals). *Let $K \subset \mathbf{R}^n$ be a compact and connected set having boundary of volume zero. Let A be an open superset of K , and let*

$$\Phi : A \longrightarrow \mathbf{R}^n$$

be a C^1 -mapping such that

$$\Phi \text{ is injective on } K^\circ \quad \text{and} \quad \det \Phi' \neq 0 \text{ on } K^\circ.$$

Let

$$f : \Phi(K) \longrightarrow \mathbf{R}$$

be a continuous function. Then

$$\int_{\Phi(K)} f = \int_K (f \circ \Phi) \cdot |\det \Phi'|.$$

This section will end with a heuristic argument to support Theorem 6.7.1, and then section 6.9 will prove the theorem after some preliminaries in section 6.8. In particular, section 6.8 will explain why the left side integral in the theorem exists. (The right side integral exists because the integrand is continuous on K , which is compact and has boundary of volume zero, but the fact that $\Phi(K)$ is nice enough for the left side integral to exist requires some discussion.) From now to the end of this section, the focus is on how the theorem is used. Generally, the idea is to carry out substitutions of the sort that were called inverse substitutions in the one-variable discussion of section 6.4. That is, to apply the theorem to an integral $\int_D f$, find a suitable set K and mapping Φ such that $D = \Phi(K)$ and the integral $\int_K (f \circ \Phi) \cdot |\det \Phi'|$ is easier to evaluate instead. The new integral most likely will be easier because K has a nicer shape than D (this wasn't an issue in the one-variable case), but also possibly because the new integrand is more convenient.

For example, to integrate the function $f(x, y) = x^2 + y^2$ over the annulus $A_{a,b}$, recall the polar coordinate mapping $\Phi(r, \theta) = (r \cos \theta, r \sin \theta)$, and recall that under this mapping, the annulus is the image of a box,

$$A_{a,b} = \Phi([a, b] \times [0, 2\pi]).$$

The composition of the integrand with Φ is

$$(f \circ \Phi)(r, \theta) = r^2,$$

and the polar coordinate has derivative matrix

$$\Phi' = \begin{bmatrix} \cos \theta & -r \sin \theta \\ \sin \theta & r \cos \theta \end{bmatrix},$$

with absolute determinant

$$|\det \Phi'| = r.$$

So by the Change of Variable Theorem, the desired integral is instead an integral over a box in polar coordinate space,

$$\int_{A_{a,b}} f = \int_{[a,b] \times [0, 2\pi]} r^2 \cdot r$$

By Fubini's Theorem, the latter integral can be evaluated as an iterated integral,

$$\int_{[a,b] \times [0, 2\pi]} r^3 = \int_{\theta=0}^{2\pi} \int_{r=a}^b r^3 = \frac{\pi}{2}(b^4 - a^4).$$

Similarly, the quarter disk $Q_b = \Phi([0, b] \times [0, \pi/2])$ has centroid (\bar{x}, \bar{y}) where

$$\bar{x} = \frac{\int_Q x}{\text{area}(Q)} = \frac{\int_{\theta=0}^{\pi/2} \int_{r=0}^b r \cos \theta \cdot r}{\pi b^2/4} = \frac{b^3/3}{\pi b^2/4} = \frac{4b}{3\pi},$$

and \bar{y} takes the same value by symmetry. Indeed $4/(3\pi)$ is somewhat less than $1/2$, in conformance with our physical intuition of the centroid of a region as its balancing point.

The sharp-eyed reader has noticed that a subtle aspect of Theorem 6.7.1 was in play for this example. Although the polar change of coordinate mapping $\Phi(r, \theta)$ is defined on all of \mathbf{R}^2 , it fails to be injective on all of the box $K = [a, b] \times [0, 2\pi]$: the 2π -periodic behavior of Φ as a function of θ maps the top and bottom edges of the box to the same segment of the x -axis. Furthermore, if the annulus has inner radius $a = 0$, i.e., if the annulus is a disk, then Φ not only collapses the left edge of the box to the origin in the (x, y) -plane, but also $\det \Phi' = 0$ on the left edge of the box. Thus we need the theorem's hypotheses that Φ need be injective only on the interior of K , and that the condition $\det \Phi' \neq 0$ need hold only on the interior of K .

Just as polar coordinates are convenient for radial symmetry in \mathbf{R}^2 , **cylindrical coordinates** in \mathbf{R}^3 conveniently describe regions with symmetry about the z -axis. A point $p \in \mathbf{R}^3$ with cartesian coordinates (x, y, z) has cylindrical coordinates (r, θ, z) where (r, θ) are the polar coordinates for the point (x, y) . (See figure 6.37.)

The cylindrical change of variable mapping is thus

$$\Phi : \mathbf{R}_{\geq 0} \times [0, 2\pi] \times \mathbf{R} \longrightarrow \mathbf{R}^3$$




Figure 6.37. Cylindrical coordinates

given by

$$\Phi(r, \theta, z) = (r \cos \theta, r \sin \theta, z).$$

This is just the polar coordinate mapping on z cross-sections, so like the polar map, it is mostly injective. Its derivative matrix is

$$\Phi' = \begin{bmatrix} \cos \theta & -r \sin \theta & 0 \\ \sin \theta & r \cos \theta & 0 \\ 0 & 0 & 1 \end{bmatrix},$$

and again

$$|\det \Phi'| = r.$$

So, for example, to integrate $f(x, y, z) = y^2 z$ over the cylinder $C : x^2 + y^2 \leq 1$, $0 \leq z \leq 2$, note that $C = \Phi([0, 1] \times [0, 2\pi] \times [0, 2])$, and therefore by the Change of Variable Theorem and then Fubini's Theorem,

$$\int_C f = \int_{\theta=0}^{2\pi} \int_{r=0}^1 \int_{z=0}^2 r^2 \sin^2 \theta \cdot z \cdot r = \int_{\theta=0}^{2\pi} \sin^2 \theta \cdot \frac{r^4}{4} \Big|_{r=0}^1 \cdot \frac{z^2}{2} \Big|_{z=0}^2 = \frac{\pi}{2}.$$

From now on, Fubini's Theorem no longer necessarily warrants comment.

For another example, we evaluate the integral $\int_S \sqrt{x^2 + y^2}$ where S is the region bounded by $z^2 = x^2 + y^2$, $z = 0$, and $z = 1$. (This looks like an ice cream cone with the ice cream licked down flat.) The Change of Variable Theorem transforms the integral into (r, θ, z) -coordinates,

$$\int_S \sqrt{x^2 + y^2} = \int_{r=0}^1 r^2 \int_{\theta=0}^{2\pi} \int_{z=r}^1 1 = \frac{\pi}{6}.$$

Spherical coordinates in \mathbf{R}^3 are designed to exploit symmetry about the origin. A point $p = (x, y, z) \in \mathbf{R}^3$ has spherical coordinates (ρ, θ, φ) where

the spherical radius ρ is the distance from the origin to p , the longitude θ is the angle from the positive x -axis to the (x, y) -projection of p , and the colatitude φ is the angle from the positive z -axis to p . By some geometry, the spherical coordinate mapping is

$$\Phi : R_{\geq 0} \times [0, 2\pi] \times [0, \pi] \longrightarrow \mathbf{R}^3$$

given by

$$\Phi(\rho, \theta, \varphi) = (\rho \cos \theta \sin \varphi, \rho \sin \theta \sin \varphi, \rho \cos \varphi).$$

This has derivative matrix

$$\Phi' = \begin{bmatrix} \cos \theta \sin \varphi & -\rho \sin \theta \sin \varphi & \rho \cos \theta \cos \varphi \\ \sin \theta \sin \varphi & \rho \cos \theta \sin \varphi & \rho \sin \theta \cos \varphi \\ \cos \varphi & 0 & -\rho \sin \varphi \end{bmatrix},$$

with determinant $\det \Phi' = -\rho^2 \sin \varphi$, so that since $0 \leq \varphi \leq \pi$,

$$|\det \Phi'| = \rho^2 \sin \varphi.$$

That is, the spherical coordinate mapping reverses orientation. It can be re-defined to preserve orientation by changing φ to the latitude angle, varying from $-\pi/2$ to $\pi/2$, rather than the colatitude.

Figure 6.38 shows the image under the spherical coordinate mapping of some (θ, φ) -rectangles, each having a fixed value of ρ , and similarly for figure 6.39 for some fixed values of θ , and figure 6.40 for some fixed values of φ . Thus the spherical coordinate mapping takes boxes to regions with these sorts of walls, such as the half ice cream cone with a bite taken out of its bottom in figure 6.41.

For an example of the Change of Variable Theorem using spherical coordinates, the solid ball of radius r in \mathbf{R}^3 is

$$B_3(r) = \Phi([0, r] \times [0, 2\pi] \times [0, \pi]),$$

and therefore its volume is

$$\text{vol}(B_3(r)) = \int_{B_3(r)} 1 = \int_{\theta=0}^{2\pi} \int_{\rho=0}^r \int_{\varphi=0}^{\pi} \rho^2 \sin \varphi = 2\pi \cdot \frac{1}{3}r^3 \cdot 2 = \frac{4}{3}\pi r^3.$$

It follows that the cylindrical shell $B_3(b) - B_3(a)$ has volume $4\pi(b^3 - a^3)/3$. See exercises 6.7.12 through 6.7.14 for the lovely formula giving the volume of the n -ball for arbitrary n .

The Change of Variable Theorem and spherical coordinates work together to integrate over the solid ellipsoid of (positive) axes a, b, c ,

$$E_{a,b,c} = \{(x, y, z) : (x/a)^2 + (y/b)^2 + (z/c)^2 \leq 1\}.$$

For example, to compute the integral




Figure 6.38. Spherical coordinates for some fixed spherical radii




Figure 6.39. Spherical coordinates for some fixed longitudes

$$\int_{E_{a,b,c}} (Ax^2 + By^2 + Cz^2),$$

first define a change of variable mapping that stretches the unit sphere into the ellipsoid,

$$\Phi : B_3(1) \longrightarrow E_{a,b,c}, \quad \Phi(u, v, w) = (au, bv, cw).$$

Thus

$$\Phi' = \begin{bmatrix} a & 0 & 0 \\ 0 & b & 0 \\ 0 & 0 & c \end{bmatrix}, \quad |\det \Phi'| = abc.$$




Figure 6.40. Spherical coordinates for some fixed colatitudes




Figure 6.41. The spherical coordinate mapping on a box

Let $f(x, y, z) = Cz^2$. Then since $E_{a,b,c} = \Phi(B_3(1))$ and $(f \circ \Phi)(u, v, w) = Cc^2w^2$, part of the integral is

$$\int_{\Phi(B_3(1))} f = \int_{B_3(1)} (f \circ \Phi) \cdot |\det \Phi'| = abc^3 C \int_{B_3(1)} w^2.$$

Apply the Change of Variable Theorem again, this time using the spherical coordinate mapping into (u, v, w) -space,

$$abc^3 C \int_{B_3(1)} w^2 = abc^3 C \int_{\rho=0}^1 \int_{\theta=0}^{2\pi} \int_{\varphi=0}^{\pi} \rho^2 \cos^2 \varphi \cdot \rho^2 \sin \varphi = \frac{4\pi}{15} abc^3 C.$$

By the symmetry of the symbols in the original integral, its overall value is therefore

$$\int_{E_{a,b,c}} (Ax^2 + By^2 + Cz^2) = \frac{4\pi}{15} abc(a^2A + b^2B + c^2C).$$

Another example is to find the centroid of upper hemispherical shell

$$S = (B_3(b) - B_3(a)) \cap \{z \geq 0\}.$$

By symmetry, $\bar{x} = \bar{y} = 0$. As for \bar{z} , compute using spherical coordinates that

$$\int_S z = \int_{\rho=a}^b \int_{\theta=0}^{2\pi} \int_{\varphi=0}^{\pi/2} \rho \cos \varphi \cdot \rho^2 \sin \varphi = \frac{\pi}{4}(b^4 - a^4).$$

This needs to be divided by the volume $2\pi(b^3 - a^3)/3$ of S to give

$$\bar{z} = \frac{3(b^4 - a^4)}{8(b^3 - a^3)}.$$

In particular, the centroid of the solid hemisphere is $3/8$ of the way up. It is perhaps surprising that π does not figure in this formula, as it did in the two-dimensional case.

Here is a heuristic argument to support the Change of Variable Theorem. Suppose that K is a box. Recall the assertion: under certain conditions,

$$\int_{\Phi(K)} f = \int_K (f \circ \Phi) \cdot |\det \Phi'|.$$

Take a partition P dividing K into subboxes J , and in each subbox choose a point x_J . If the partition is fine enough, then each J maps under Φ to a small patch A of volume $\text{vol}(A) \approx |\det \Phi'(x_J)| \text{vol}(J)$ (cf. section 3.8), and each x_J maps to a point $y_A \in A$. (See figure 6.42.) Since the integral is a limit of weighted sums,

$$\begin{aligned} \int_{\Phi(K)} f &\approx \sum_A f(y_A) \text{vol}(A) \\ &\approx \sum_J f(\Phi(x_J)) |\det \Phi'(x_J)| \text{vol}(J) \\ &\approx \int_K (f \circ \Phi) \cdot |\det \Phi'|, \end{aligned}$$

and these should become equalities in the limit as P becomes finer. What makes this reasoning incomplete is that the patches A are not boxes, as are required for our theory of integration.

Recall from sections 3.8 and 3.9 that the absolute value of $\det \Phi'(x)$ describes how the mapping Φ scales volume at x , while the sign of $\det \Phi'(x)$




Figure 6.42. Change of variable

says whether the mapping locally preserves or reverses orientation. The factor $|\det \Phi'|$ in the n -dimensional Change of Variable Theorem (rather than the signed $\det \Phi'$) reflects the fact that n -dimensional integration does not take orientation into account. This is less satisfying than the one-variable theory, which does consider orientation and therefore comes with a signed change of variable theorem, $\int_{\phi(a)}^{\phi(b)} f = \int_a^b (f \circ \phi) \cdot \phi'$. An orientation-sensitive n -dimensional integration theory will be developed in chapter 8.

Exercises

6.7.1. Evaluate $\int_S x^2 + y^2$ where S is the region bounded by $x^2 + y^2 = 2z$ and $z = 2$. Sketch S .

6.7.2. Find the volume of the region S between $x^2 + y^2 = 4z$ and $x^2 + y^2 + z^2 = 5$. Sketch S .

6.7.3. Find the volume of the region between the graphs of $z = x^2 + y^2$ and $z = (x^2 + y^2 + 1)/2$.

6.7.4. Derive the spherical coordinate mapping.

6.7.5. Let Φ be the spherical coordinate mapping. Describe $\Phi(K)$ where

$$K = \{(\rho, \theta, \varphi) : 0 \leq \theta \leq 2\pi, 0 \leq \varphi \leq \pi/2, 0 \leq \rho \leq \cos \varphi\}.$$

Same question for

$$K = \{(\rho, \theta, \varphi) : 0 \leq \theta \leq 2\pi, 0 \leq \varphi \leq \pi, 0 \leq \rho \leq \sin \varphi\}.$$

6.7.6. Evaluate $\int_S xyz$ where S is the first octant of $B_3(1)$.

6.7.7. Find the mass of a solid figure filling the spherical shell

$$S = B_3(b) - B_3(a)$$

with density $\delta(x, y, z) = x^2 + y^2 + z^2$.

6.7.8. A solid sphere of radius r has density $\delta(x, y, z) = e^{-(x^2+y^2+z^2)^{3/2}}$. Find its mass, $\int_{B_3(r)} \delta$.

6.7.9. Find the centroid of the region $S = B_3(a) \cap \{x^2 + y^2 \leq z^2\} \cap \{z \geq 0\}$. Sketch S .

6.7.10. Prove the **change of scale principle**: If the set $K \subset \mathbf{R}^n$ has volume v then for any $r \geq 0$, the set $rK = \{rx : x \in K\}$ has volume $r^n v$. (Change variables by $\Phi(x) = rx$.)

6.7.11. (a) Prove **Pappas's Theorem**: Let K be a compact set in the (x, z) -plane lying to the right of the z -axis and with boundary of area zero. Let S be the solid obtained by rotating K about the z -axis in \mathbf{R}^3 . Then

$$\text{vol}(S) = 2\pi\bar{x} \cdot \text{area}(K),$$

where as always, $\bar{x} = \int_K x / \text{area}(K)$. (Use cylindrical coordinates.)

(b) What is the volume of the torus $T_{a,b}$ of cross-sectional radius a and major radius b ? (See figure 6.43.)




Figure 6.43. Torus

6.7.12. (Volume of the n -ball, first version.) Let $n \in \mathbf{Z}^+$ and $r \in \mathbf{R}_{\geq 0}$. The n -dimensional ball of radius r is

$$B_n(r) = \{x : x \in \mathbf{R}^n \mid |x| \leq r\} = \{(x_1, \dots, x_n) : x_1^2 + \dots + x_n^2 \leq r^2\}.$$

Let

$$v_n = \text{vol}(B_n(1)).$$

(a) Explain how exercise 6.7.10 reduces computing the volume of $B_n(r)$ to computing v_n .

(b) Explain why $v_1 = 2$ and $v_2 = \pi$.

(c) Let D denote the unit disk $B_2(1)$. Prove that for $n > 2$,

$$B_n(1) = \{(x_1, x_2) \times B_{n-2}(\sqrt{1 - x_1^2 - x_2^2}) : (x_1, x_2) \in D\}.$$

That is, the unit n -ball is a union of cross-sectional $(n-2)$ -dimensional balls of radius $\sqrt{1-x_1^2-x_2^2}$ as (x_1, x_2) varies through the unit disk. Make a sketch for $n=3$, the only value of n for which we can see this.

(d) Explain why for $n > 2$,

$$\begin{aligned} v_n &= v_{n-2} \int_{(x_1, x_2) \in D} (1-x_1^2-x_2^2)^{\frac{n}{2}-1} \\ &= v_{n-2} \int_{\theta=0}^{2\pi} \int_{r=0}^1 (1-r^2)^{\frac{n}{2}-1} \cdot r \\ &= v_{n-2} \pi / (n/2). \end{aligned}$$

(Use the definition of volume at the end of section 6.5, Fubini's Theorem, the definition of volume again, the change of scale principle from the previous exercise, and the Change of Variable Theorem.)

(e) Show by induction only the *for n even* case of the formula

$$v_n = \begin{cases} \frac{\pi^{n/2}}{(n/2)!} & \text{for } n \text{ even,} \\ \frac{\pi^{(n-1)/2} 2^n ((n-1)/2)!}{n!} & \text{for } n \text{ odd.} \end{cases}$$

(The *for n odd* case can be shown by induction as well, but the next two exercises provide a better, more conceptual approach to the volumes of odd-dimensional balls.)

6.7.13. This exercise computes the “improper” integral $I = \int_{x=0}^{\infty} e^{-x^2}$, defined as the limit $\lim_{R \rightarrow \infty} \int_{x=0}^R e^{-x^2}$. Let $I(R) = \int_{x=0}^R e^{-x^2}$ for any $R \geq 0$.

(a) Use Fubini's Theorem to show that $I(R)^2 = \int_{S(R)} e^{-x^2-y^2}$, where $S(R)$ is the square

$$S(R) = \{(x, y) : 0 \leq x \leq R, 0 \leq y \leq R\}.$$

(b) Let $Q(R)$ be the quarter disk

$$Q(R) = \{(x, y) : 0 \leq x, 0 \leq y, x^2 + y^2 \leq R^2\},$$

and similarly for $Q(\sqrt{2}R)$. Explain why

$$\int_{Q(R)} e^{-x^2-y^2} \leq \int_{S(R)} e^{-x^2-y^2} \leq \int_{Q(\sqrt{2}R)} e^{-x^2-y^2}.$$

(c) Change variables, and evaluate $\int_{Q(R)} e^{-x^2-y^2}$ and $\int_{Q(\sqrt{2}R)} e^{-x^2-y^2}$. What are the limits of these two quantities as $R \rightarrow \infty$?

(d) What is I ?

6.7.14. (Volume of the n -ball, improved version) Define the gamma function as an integral,

$$\Gamma(s) = \int_{x=0}^{\infty} x^{s-1} e^{-x} dx, \quad s > 0.$$

(This “improper” integral is well-behaved, even though it is not being carried out over a bounded region and even though the integrand is unbounded near $x = 0$ when $0 < s < 1$.)

(a) Show: $\Gamma(1) = 1$, $\Gamma(1/2) = \sqrt{\pi}$, $\Gamma(s+1) = s\Gamma(s)$. (Substitute and see the previous exercise for the second identity, integrate by parts for the third.)

(b) Use part (a) to show that $n! = \Gamma(n+1)$ for $n = 0, 1, 2, \dots$. Accordingly, define $x! = \Gamma(x+1)$ for all real numbers $x > -1$, not only nonnegative integers.

(c) Use exercise 6.7.12(b), exercise 6.7.12(d), and the extended definition of the factorial in part (b) of this exercise to obtain a uniform formula for the volume of the unit n -ball,

$$v_n = \frac{\pi^{n/2}}{(n/2)!}, \quad n = 1, 2, 3, \dots$$

(We already have this formula for n even. For n odd, the argument is essentially identical to exercise 6.7.12(d) but starting at the base case $n = 1$.) Thus the n -ball of radius r has volume

$$\text{vol}(B_n(r)) = \frac{\pi^{n/2}}{(n/2)!} r^n, \quad n = 1, 2, 3, \dots$$

6.8 Topological Preliminaries for the Change of Variable Theorem

In preparation for proving the Change of Variable Theorem (Theorem 6.7.1), we review its statement. The statement includes the terms *boundary* and *interior*, which we have considered only informally so far, but we soon will discuss them more carefully. The statement also includes the term *open*, and the reader is reminded that a set is called open if its complement is closed; we soon will review the definition of a closed set. The statement includes the term \mathcal{C}^1 -*mapping*, meaning a mapping such that all partial derivatives of all of its component functions exist and are continuous. And the statement includes the notation K° for the interior of a set K . The theorem says:

Let $K \subset \mathbf{R}^n$ be a compact and connected set having boundary of volume zero. Let A be an open superset of K , and let

$$\Phi: A \longrightarrow \mathbf{R}^n$$

be a \mathcal{C}^1 -mapping such that

$$\Phi \text{ is injective on } K^\circ \quad \text{and} \quad \det \Phi' \neq 0 \text{ on } K^\circ.$$

Let

$$f : \Phi(K) \rightarrow \mathbf{R}$$

be a continuous function. Then

$$\int_{\Phi(K)} f = \int_K (f \circ \Phi) \cdot |\det \Phi'|.$$

Thus the obvious data for the theorem are K , Φ , and f . (The description of Φ subsumes A , and in any case the role of A is auxiliary.) But also, although the dimension n is conceptually generic-but-fixed, in fact the proof of the theorem will entail induction on n , so that we should view n as a variable part of the setup as well. Here are some comments about the data.

- The continuous image of a compact set is compact (Theorem 2.4.14), so that $\Phi(K)$ is again compact. Similarly, by an invocation in section 2.4, the continuous image of a connected set is connected, so that $\Phi(K)$ is again connected. The reader who wants to minimize invocation may instead assume that K is path-connected, so that $\Phi(K)$ is again path-connected (see exercise 2.4.9 for the definition of path-connectedness and the fact that path-connectedness is a topological property); the distinction between connectedness and path-connectedness is immaterial for any example that will arise in calculus. We soon will see that also the image $\Phi(K)$ again has boundary of volume zero, so that in fact $\Phi(K)$ inherits all of the assumed properties of K .
- Thus both integrals in the Change of Variable Theorem exist, because in each case the integrand is continuous on the domain of integration and the domain of integration is compact and has boundary of volume zero.
- The hypotheses of the theorem can be weakened or strengthened in various ways with no effect on the outcome. Indeed, the proof of the theorem proceeds partly by strengthening the hypotheses. The hypotheses in Theorem 6.7.1 were chosen to make the theorem fit the applications that arise in calculus. Especially, parameterizations by polar, cylindrical, or spherical coordinates often degenerate on the boundary of the parameter-box, hence the conditions that Φ is injective and $\det \Phi' \neq 0$ being required only on the interior K° . On the other hand, while the hypotheses about Φ are weaker than necessary in order to make the theorem easier to use, the hypothesis that f is continuous is stronger than necessary in order to make the theorem easier to prove. The theorem continues to hold if f is assumed only to be integrable, but then the proof is more work. In calculus examples, f is virtually always continuous.

This section places a few more topological ideas into play to set up the proof of the Change of Variable Theorem in the next section. The symbols K , A , Φ , and f denoting the set, the open superset, the change of variable, and the function in the theorem will retain their meanings throughout the discussion. Symbols such as S will denote other sets, symbols such as Ψ will denote other transformations, and symbols such as g will denote other functions.

Recall some topological ideas that we have already discussed.

- For any point $a \in \mathbf{R}^n$ and any positive real number $r > 0$, the open ball centered at a of radius r is the set

$$B(a, r) = \{x \in \mathbf{R}^n : |x - a| < r\}.$$

- A point $a \in \mathbf{R}^n$ is called a limit point of a set $S \in \mathbf{R}^n$ if every open ball centered at a contains some point $x \in S$ such that $x \neq a$. The subset A of \mathbf{R}^n is called closed if it contains all of its limit points.

Definition 6.8.1. Let S be a subset of \mathbf{R}^n . Its **closure** \overline{S} is the smallest closed superset of S .

Here *smallest* is taken in the sense of set-containment. The intersection of closed sets is closed (exercise 6.8.1(a)), and so \overline{S} is the intersection of all closed supersets of S , including \mathbf{R}^n . This shows that \overline{S} exists and is unique. The special-case definition

$$\overline{B(a, r)} = \{x \in \mathbf{R}^n : |x - a| \leq r\}$$

from section 5.1 is consistent with Definition 6.8.1.

Closed sets can also be described in terms of *boundary points* rather than limit points.

Definition 6.8.2. Let S be a subset of \mathbf{R}^n . A point $p \in \mathbf{R}^n$ is called a **boundary point** of S if for every $r > 0$ the open ball $B(p, r)$ contains a point from S and a point from the complement S^c . The **boundary** of S , denoted ∂S , is the set of boundary points of S .

A boundary point of a set need not be a limit point of the set, and a limit point of a set need not be a boundary point of the set (exercise 6.8.1(b)). Nonetheless, similarly to the definition of closed set in the second bullet before Definition 6.8.1, a set is closed if and only if it contains all of its boundary points (exercise 6.8.1(c)). The boundary of any set is closed (exercise 6.8.1(d)). Since the definition of boundary point is symmetric in the set and its complement, the boundary of the set is also the boundary of the complement,

$$\partial S = \partial(S^c).$$

In general, the closure of a set is the union of the set and its boundary (exercise 6.8.2(a)),

$$\overline{S} = S \cup \partial S.$$

If S is bounded then so is its closure \overline{S} (exercise 6.8.2(b)), and therefore the closure of a bounded set is compact. The special-case definition

$$\partial \overline{B(a, r)} = \{x \in \mathbf{R}^n : |x - a| = r\}$$

from section 6.1 is consistent with Definition 6.8.2.

Definition 6.8.3. An open box in \mathbf{R}^n is a set of the form

$$J = (a_1, b_1) \times (a_2, b_2) \times \cdots \times (a_n, b_n).$$

The word box, unmodified, continues to mean a closed box.

Proposition 6.8.4 (Finiteness Property of Compact Sets). Consider a compact set $K \subset \mathbf{R}^n$. Suppose that some collection of open boxes J_i covers K . Then a finite collection of the open boxes J_i covers K .

Proof (Sketch). Suppose that no finite collection of the open boxes J_i covers K . Let B_1 be a box that contains K . Partition B_1 into 2^n subboxes \tilde{B} by bisecting it in each direction. If for each subbox \tilde{B} , some finite collection of the open boxes J_i covers $K \cap \tilde{B}$, then the 2^n -fold collection of these finite collections in fact covers all of K . Thus no finite collection of the open boxes J_i covers $K \cap \tilde{B}$ for at least one subbox \tilde{B} of B_1 . Name some such subbox B_2 , repeat the argument with B_2 in place of B_1 , and continue in this fashion, obtaining nested boxes

$$B_1 \supset B_2 \supset B_3 \supset \cdots$$

whose sides are half as long at each succeeding generation, and such that no $K \cap B_j$ is covered by a finite collection of the open boxes J_i . The intersection $B_1 \cap B_2 \cap \cdots$ contains at most one point because the boxes B_j eventually shrink smaller than the distance between any two given distinct points. On the other hand, since each B_j is compact and the B_j are nested, the sequence $\{c_j\}$ of the centerpoints of the B_j has a subsequence that converges in each B_j and hence converges in the intersection. Thus the intersection is nonempty, consisting of a single point c . Some open box J_i covers c , and so it covers B_j for all high enough indices j , and thus in fact it covers $K \cap B_j$ for all high enough indices j . This is a contradiction, making the initial supposition that no finite collection of the open boxes J_i covers K untenable. \square

The following lemma is similar to the Difference Magnification Lemma (Lemma 5.1.3).

Lemma 6.8.5 (Box-Volume Magnification Lemma). Let B be a box in \mathbf{R}^n whose longest side is at most twice its shortest side. Let g be a differentiable mapping from an open superset of B in \mathbf{R}^n back to \mathbf{R}^n . Suppose that there is a number c such that $|D_j g_i(x)| \leq c$ for all $i, j \in \{1, \dots, n\}$ and all $x \in B$. Then $g(B)$ sits in a box B' such that $\text{vol}(B') \leq (2nc)^n \text{vol}(B)$.

Proof. Let x be the centerpoint of B and let \tilde{x} be any point of B . Make the line segment connecting x to \tilde{x} the image of a function of one variable,

$$\gamma : [0, 1] \longrightarrow \mathbf{R}^n, \quad \gamma(t) = x + t(\tilde{x} - x).$$

Fix any $i \in \{1, \dots, n\}$. Identically to the proof of the Difference Magnification Lemma, we have for some $t \in (0, 1)$,

$$g_i(\tilde{x}) - g_i(x) = \langle g'_i(\gamma(t)), \tilde{x} - x \rangle.$$

For each j , the j th entry of the vector $g'_i(\gamma(t))$ is $D_j g_i(\gamma(t))$, and we are given that $|D_j g_i(\gamma(t))| \leq c$. Also, the j th entry of the vector $\tilde{x} - x$ satisfies $|\tilde{x}_j - x_j| \leq \ell/2$ where ℓ is the longest side of B . Thus

$$|g_i(\tilde{x}) - g_i(x)| \leq n c \ell / 2,$$

and so

$$g_i(B) \subset [g_i(x) - n c \ell / 2, g_i(x) + n c \ell / 2].$$

Apply this argument for each $i \in \{1, \dots, n\}$ to show that $g(B)$ lies in the box B' centered at $g(x)$ having sides $n c \ell$ and therefore having volume

$$\text{vol}(B') = (n c \ell)^n.$$

On the other hand, since the shortest side of B is at least $\ell/2$,

$$\text{vol}(B) \geq (\ell/2)^n.$$

The result follows. \square

Using the previous two results, we can show that the property of having volume zero is preserved under mappings that are well enough behaved. However, we need to assume more than just continuity. Having volume zero is not a topological property.

Proposition 6.8.6 (Volume Zero Preservation Under C^1 -Mappings).

Let $S \subset \mathbf{R}^n$ be a compact set having volume zero. Let A be an open superset of S , and let

$$\Phi: A \longrightarrow \mathbf{R}^n$$

be a C^1 -mapping. Then $\Phi(S)$ again has volume zero.

Proof. For each $s \in S$ there exists an $r_s > 0$ such that the copy of the box $[-r_s, r_s]^n$ centered at s lies in A (exercise 6.8.5(a)). Let J_s denote the corresponding open box, i.e., a copy of $(-r_s, r_s)^n$ centered at s . By the Finiteness Property of compact sets, a collection of finitely many of the open boxes J_s covers S , so certainly the corresponding collection U of the closed boxes does so as well. As a finite union of compact sets, U is compact (exercise 6.8.1(f)). Therefore the partial derivatives $D_j \Phi_i$ for $i, j = 1, \dots, n$ are uniformly continuous on U , and so some constant c bounds all $D_j \Phi_i$ on U .

Let $\varepsilon > 0$ be given. Cover S by finitely many boxes B_i having total volume less than $\varepsilon / (2nc)^n$. After replacing each box by its intersections with the boxes of U , we may assume that the boxes all lie in U . (Here it is relevant that the intersection of two boxes is a box.) And after further subdividing the boxes if necessary, we may assume that the longest side of each box is at most twice the shortest side (exercise 6.8.6(b)). By the Box-Volume Magnification Lemma, the Φ -images of the boxes lie in a union of boxes B'_i having volume

$$\sum_i \text{vol}(B'_i) \leq (2nc)^n \sum_i \text{vol}(B_i) < \varepsilon.$$

□

The last topological preliminary that we need is the formal definition of *interior*.

Definition 6.8.7 (Interior Point, Interior of a Set). *Let $S \subset \mathbf{R}^n$ be a set. Any nonboundary point of S is an interior point of S . Thus x is an interior point of S if some open ball $B(x, r)$ lies entirely in S . The interior of S is*

$$S^\circ = \{\text{interior points of } S\}.$$

The interior of any set S is open (exercise 6.8.6(a)). Any set decomposes as the disjoint union of its interior and its boundary (exercise 6.8.6(b)),

$$S = S^\circ \cup \partial S, \quad S^\circ \cap \partial S = \emptyset.$$

As anticipated at the beginning of the section, we now can complete the argument that the properties of the set K in the Change of Variable Theorem are preserved by the mapping Φ in the theorem.

Proposition 6.8.8. *Let $K \subset \mathbf{R}^n$ be a compact and connected set having boundary of volume zero. Let A be an open superset of K , and let $\Phi : A \rightarrow \mathbf{R}^n$ be a C^1 -mapping such that $\det \Phi' \neq 0$ everywhere on K° . Then $\Phi(K)$ is again a compact and connected set having boundary of volume zero.*

Proof. We have discussed the fact that $\Phi(K)$ is again compact and connected. Restrict Φ to K . The Inverse Function Theorem says that Φ maps interior points of K to interior points of $\Phi(K)$, and thus $\partial(\Phi(K)) \subset \Phi(\partial K)$. By the Volume-Zero Preservation proposition $\text{vol}(\Phi(\partial K)) = 0$, and so $\text{vol}(\partial(\Phi(K))) = 0$ as well. □

Exercises

6.8.1. (a) Show that every intersection—not just twofold intersections and not even just finite-fold intersections—of closed sets is closed. (Recall from Proposition 2.4.5 that a set S is closed if and only if every sequence in S that converges in \mathbf{R}^n in fact converges in S .)

(b) Show by example that a boundary point of a set need not be a limit point of the set. Show by example that a limit point of a set need not be a boundary point of the set.

(c) Show that a set is closed if and only if it contains each of its boundary points. (Again recall the characterization of closed sets mentioned in part (a).)

(d) Show that the boundary of any set is closed.

(e) Show that every union of two closed sets is closed. It follows that any union of finitely many closed sets is closed. Recall that by definition a set is

open if its complement is closed. Explain why consequently every intersection of finitely many open sets is open.

(f) Explain why any union of finitely many compact sets is compact.

6.8.2. Let S be any subset of \mathbf{R}^n .

(a) Show that its closure is its union with its boundary, $\bar{S} = S \cup \partial S$.

(b) Show that if S is bounded then so is \bar{S} .

6.8.3. (a) Which points of the proof of Proposition 6.8.4 are sketchy? Fill in the details.

(b) Let S be an unbounded subset of \mathbf{R}^n , meaning that S is not contained in any ball. Find a collection of open boxes J_i that covers S but such that no finite subcollection of the open boxes J_i covers S .

(c) Let S be an bounded but non-closed subset of \mathbf{R}^n , meaning that S is bounded but missing a limit point. Find a collection of open boxes J_i that covers S but such that no finite subcollection of the open boxes J_i covers S .

6.8.4. Let $\varepsilon > 0$. Consider the box $B = [0, 1] \times [0, \varepsilon] \subset \mathbf{R}^2$, and consider the mapping $g : \mathbf{R}^2 \rightarrow \mathbf{R}^2$ given by $g(x, y) = (x, x)$. What is the smallest box B' containing $g(B)$? What is the ratio $\text{vol}(B')/\text{vol}(B)$? Discuss the relation between this example and Lemma 6.8.5.

6.8.5. The following questions are about the proof of Proposition 6.8.6.

(a) Explain why for each $s \in S$ there exists an $r_s > 0$ such that the copy of the box $[-r_s, r_s]^n$ centered at s lies in A .

(b) Explain why any box (with all sides assumed to be positive) can be subdivided into boxes whose longest side is at most twice the shortest side.

6.8.6. Let $S \subset \mathbf{R}^n$ be any set.

(a) Show that the interior S° is open.

(b) Show that S decomposes as the disjoint union of its interior and its boundary.

6.9 Proof of the Change of Variable Theorem

Again recall the statement of the Change of Variable Theorem:

Let $K \subset \mathbf{R}^n$ be a compact and connected set having boundary of volume zero. Let A be an open superset of K , and let $\Phi : A \rightarrow \mathbf{R}^n$ be a C^1 -mapping such that Φ is injective on K° and $\det \Phi' \neq 0$ on K° . Let $f : \Phi(K) \rightarrow \mathbf{R}$ be a continuous function. Then

$$\int_{\Phi(K)} f = \int_K (f \circ \Phi) \cdot |\det \Phi'|.$$

We begin chipping away at the theorem by strengthening its hypotheses.

Proposition 6.9.1 (Optional Hypothesis-Strengthening). *To prove the Change of Variable Theorem, it suffices to prove the theorem subject to any combination of the following additional hypotheses:*

- K is a box.
- Φ is injective on all of A ,
- $\det \Phi' \neq 0$ on all of A .

Before proceeding to the proof of the proposition, it deserves comment that we will not always want K to be a box. But once the proposition is proved, we may take K to be a box or not as convenient.

Proof. Let $\varepsilon > 0$ be given.

Let B be a box containing K , and let P be a partition of B into subboxes J . Define three types of subbox,

- type I : J such that $J \subset K^\circ$,
 type II : J such that $J \cap \partial K \neq \emptyset$ (and thus $J \cap \partial(B \setminus K) \neq \emptyset$),
 type III : J such that $J \subset (B \setminus K)^\circ$.

These three types of box are exclusive and exhaustive (exercise 6.9.2(a)). Also define a function

$$g : B \longrightarrow \mathbf{R}, \quad g(x) = \begin{cases} (f \circ \Phi)(x) \cdot |\det \Phi'(x)| & \text{if } x \in K, \\ 0 & \text{if } x \notin K. \end{cases}$$

The continuous function f is necessarily bounded on $\Phi(K)$, say by R . The partial derivatives $D_j \Phi_i$ of the component functions of Φ are continuous on K , and so the continuous function $|\det \Phi'|$ is bounded on the compact set K , say by \tilde{R} . Thus $R\tilde{R}$ bounds g on B .

As in the proof of the Volume Zero Preservation proposition (Proposition 6.8.6), we can cover the subset K of A by a collection U of finitely many boxes that is again a subset of A , and so the continuous partial derivatives $D_j \Phi_i$ of the component functions of Φ are bounded on the compact set U , say by c . We may assume that the partition P is fine enough that all subboxes J of type I and type II lie in U (exercise 6.9.2(b)). And we may assume that the longest side of each subbox J is at most twice the shortest side. Recall that $\varepsilon > 0$ has been given. Because the boundary of K has volume zero, we may further assume that the partition P is fine enough that

$$\sum_{J:\text{type II}} \text{vol}(J) < \min \left\{ \frac{\varepsilon}{R(2nc)^n}, \frac{\varepsilon}{R\tilde{R}} \right\}$$

(exercise 6.9.2(c)).

Let

$$\Phi(K)_{\text{I}} = \bigcup_{J:\text{type I}} \Phi(J), \quad \Phi(K)_{\text{II}} = \Phi(K) \setminus \Phi(K)_{\text{I}}.$$

Then the integral on the left side of the equality in the Change of Variable Theorem decomposes into two parts,

$$\int_{\Phi(K)} f = \int_{\Phi(K)_I} f + \int_{\Phi(K)_{II}} f,$$

and because Φ is injective on K , the previous display rewrites as

$$\int_{\Phi(K)} f = \sum_{J : \text{type I}} \int_{\Phi(J)} f + \int_{\Phi(K)_{II}} f. \tag{6.12}$$

Also,

$$\Phi(K)_{II} \subset \bigcup_{J : \text{type II}} \Phi(J),$$

so that

$$\left| \int_{\Phi(K)_{II}} f \right| \leq \int_{\Phi(K)_{II}} |f| \leq \sum_{J : \text{type II}} \int_{\Phi(J)} |f|.$$

By the Box-Volume Magnification Lemma (Lemma 6.8.5), for each box J of type II, $\text{vol}(\Phi(J)) \leq (2nc)^n \text{vol}(J)$. Thus, by the bounds on f and on the sum of the type II box-volumes, it follows that

$$\left| \int_{\Phi(K)_{II}} f \right| < \varepsilon.$$

That is, the second term on the right side of (6.12) contributes as negligibly as desired to the integral on the left side, which is the integral on the left side of the Change of Variable Theorem.

Meanwhile, the integral on the right side of the equality in the Change of Variable Theorem also decomposes into two parts,

$$\int_K (f \circ \Phi) \cdot |\det \Phi'| = \sum_{J : \text{type I}} \int_J g + \sum_{J : \text{type II}} \int_J g. \tag{6.13}$$

By the bounds on g and on the sum of the type II box-volumes,

$$\left| \sum_{J : \text{type II}} \int_J g \right| \leq \sum_{J : \text{type II}} \int_J |g| < \varepsilon.$$

That is, the second term on the right side of (6.13) contributes as negligibly as desired to the integral on the left side, which is the integral on the right side of the Change of Variable Theorem.

The type I subboxes J of the partition of the box B containing the original K (which is not assumed to be a box) satisfy all of the additional hypotheses in the statement of the proposition: each J is a box, and we may

shrink the domain of Φ to the open superset K° of each J , where Φ is injective and where $\det \Phi' \neq 0$. Thus, knowing the Change of Variable Theorem subject to any of the additional hypotheses says that the first terms on the right sides of (6.12) and (6.13) are equal, making the integrals on the left sides lie within ε of each other. Since ε is arbitrary, the integrals are in fact equal. In sum, it suffices to prove the Change of Variable Theorem assuming any of the additional hypotheses as desired. \square

Proposition 6.9.2 (Alternative Optional Hypothesis-Strengthening).
To prove the Change of Variable Theorem, it suffices to prove the theorem subject to the following additional hypotheses:

- $\Phi(K)$ is a box (but now we may not assume that K is a box).
- Φ is injective on all of A .
- $\det \Phi' \neq 0$ on all of A .

Similarly to the remark after Proposition 6.9.1, we will not always want the additional hypotheses.

Proof. With the previous proposition in play, the idea now is to run through its proof in reverse, starting from the strengthened hypotheses that it grants us. Thus we freely assume that K is a box, that the change of variable mapping Φ is injective on all of A , and that $\det \Phi' \neq 0$ on all of A . By the Inverse Function Theorem, the superset $\Phi(A)$ of $\Phi(K)$ is open and $\Phi : A \rightarrow \Phi(A)$ has a C^1 inverse

$$\Phi^{-1} : \Phi(A) \rightarrow A.$$

Let $\varepsilon > 0$ be given.

Let B be a box containing $\Phi(K)$, and let P be a partition of B into subboxes J . Define three types of subbox,

- type I : J such that $J \subset \Phi(K)^\circ$,
- type II : J such that $J \cap \partial\Phi(K) \neq \emptyset$ (and thus $J \cap \partial(B \setminus \Phi(K)) \neq \emptyset$),
- type III : J such that $J \subset (B \setminus \Phi(K))^\circ$.

These three types of box are exclusive and exhaustive. Also, define as before

$$g : B \rightarrow \mathbf{R}, \quad g(x) = \begin{cases} (f \circ \Phi)(x) \cdot |\det \Phi'(x)| & \text{if } x \in K, \\ 0 & \text{if } x \notin K. \end{cases}$$

Again, f is bounded on $\Phi(K)$, say by R , and $|\det \Phi'|$ is bounded on K , say by \tilde{R} , so that $R\tilde{R}$ bounds g on B .

Cover the subset $\Phi(K)$ of $\Phi(A)$ by a collection U of finitely many boxes that is again a subset of $\Phi(A)$. Then the continuous partial derivatives $D_j \Phi_i^{-1}$ of the component functions of Φ^{-1} are bounded on the compact set U , say by c . We may assume that the partition P is fine enough that all subboxes J of type I and type II lie in U . And we may assume that the longest side of each

subbox J is at most twice the shortest side. Recall that $\varepsilon > 0$ has been given. Because the boundary of $\Phi(K)$ has volume zero, we may further assume that the partition P is fine enough that

$$\sum_{J:\text{type II}} \text{vol}(J) < \min \left\{ \frac{\varepsilon}{R}, \frac{\varepsilon}{R\tilde{R}(2nc)^n} \right\}.$$

Let

$$K_{\text{I}} = \bigcup_{J:\text{type I}} \Phi^{-1}(J), \quad K_{\text{II}} = K \setminus K_{\text{I}}.$$

Then the integral on the left side of the equality in the Change of Variable Theorem decomposes into two parts,

$$\int_{\Phi(K)} f = \sum_{J:\text{type I}} \int_J f + \sum_{J:\text{type II}} \int_J f. \tag{6.14}$$

By the bounds on f and on the sum of the type II box-volumes,

$$\left| \sum_{J:\text{type II}} \int_J f \right| \leq \sum_{J:\text{type II}} \int_J |f| < \varepsilon.$$

That is, the second term on the right side of (6.14) contributes as negligibly as desired to the integral on the left side, which is the integral on the left side of the Change of Variable Theorem.

Meanwhile, the integral on the right side of the equality in the Change of Variable Theorem also decomposes into two parts,

$$\int_K (f \circ \Phi) \cdot |\det \Phi'| = \int_{K_{\text{I}}} g + \int_{K_{\text{II}}} g,$$

and because Φ^{-1} is injective, the previous display rewrites as

$$\int_K (f \circ \Phi) \cdot |\det \Phi'| = \sum_{J:\text{type I}} \int_{\Phi^{-1}(J)} g + \int_{K_{\text{II}}} g. \tag{6.15}$$

Also,

$$K_{\text{II}} \subset \bigcup_{J:\text{type II}} \Phi^{-1}(J),$$

so that

$$\left| \int_{K_{\text{II}}} g \right| \leq \int_{K_{\text{II}}} |g| \leq \sum_{J:\text{type II}} \int_{\Phi^{-1}(J)} |g|.$$

For each box J of type II, $\text{vol}(\Phi^{-1}(J)) \leq (2nc)^n \text{vol}(J)$. Thus, by the bounds on g and on the sum of the type II box-volumes, it follows that

$$\left| \int_{K_{\text{II}}} g \right| < \varepsilon.$$

That is, the second term on the right side of (6.15) contributes as negligibly as desired to the integral on the left side, which is the integral on the right side of the Change of Variable Theorem.

The type I subboxes J of the partition of the box B containing the original $\Phi(K)$ (which is not assumed to be a box) satisfy the new additional hypothesis in the statement of the proposition. The other two additional hypotheses in the statement of the proposition are already assumed. Thus, knowing the Change of Variable Theorem subject to the additional hypotheses says that the first terms on the right sides of (6.14) and (6.15) are equal, making the integrals on the left sides lie within ε of each other. Since ε is arbitrary, the integrals are in fact equal. In sum, it suffices to prove the Change of Variable Theorem assuming the additional hypotheses as desired. \square

Proposition 6.9.3 (Further Optional Hypothesis-Strengthening). *To prove the Change of Variable Theorem, it suffices to prove the theorem subject to the additional hypothesis that f is identically 1.*

As with the other hypothesis-strengthenings, we will not always want f to be identically 1, but we may take it to be so when convenient.

Proof. We assume the strengthened hypotheses given us by Proposition 6.9.2. Let P be a partition of the box $\Phi(K)$ into subboxes J . For each subbox J , view the quantity $M_J(f) = \sup \{f(x) : x \in J\}$ both as a number and as a constant function. Assume that the Change of Variable Theorem holds for the constant function 1 and therefore for any constant function, and compute

$$\begin{aligned} \int_K (f \circ \Phi) \cdot |\det \Phi'| &= \sum_J \int_{\Phi^{-1}(J)} (f \circ \Phi) \cdot |\det \Phi'| \\ &\leq \sum_J \int_{\Phi^{-1}(J)} (M_J(f) \circ \Phi) \cdot |\det \Phi'| \\ &= \sum_J \int_J M_J(f) \quad \text{by the assumption} \\ &= \sum_J M_J(f) \operatorname{vol}(J) \\ &= U(f, P). \end{aligned}$$

As a lower bound of the upper sums, $\int_K (f \circ \Phi) \cdot |\det \Phi'|$ is at most the integral,

$$\int_K (f \circ \Phi) \cdot |\det \Phi'| \leq \int_{\Phi(K)} f.$$

A similar argument gives the opposite inequality, making the integrals equal as desired. \square

The next result will allow the proof of the Change of Variable Theorem to decompose the change of variable mapping.

Proposition 6.9.4 (Persistence Under Composition). *In the Change of Variable Theorem, suppose that the change of variable mapping is a composition*

$$\Phi = \Gamma \circ \Psi$$

where the mappings

$$\Psi : A \longrightarrow \mathbf{R}^n$$

and

$$\Gamma : \tilde{A} \longrightarrow \mathbf{R}^n \quad (\text{where } \tilde{A} \text{ is an open superset of } \Psi(K))$$

satisfy the hypotheses of the Change of Variable Theorem. If

$$\int_{\Psi(K)} g = \int_K (g \circ \Psi) \cdot |\det \Psi'| \quad \text{for continuous functions } g : \Psi(K) \longrightarrow \mathbf{R}$$

and

$$\int_{\Gamma(\Psi(K))} 1 = \int_{\Psi(K)} |\det \Gamma'|$$

then also

$$\int_{\Phi(K)} 1 = \int_K |\det \Phi'|.$$

Proof. This is a straightforward calculation using the definition of Φ , the second given equality, the first given equality, the multiplicativity of the determinant, the Chain Rule, and again the definition of Φ ,

$$\begin{aligned} \int_{\Phi(K)} 1 &= \int_{\Gamma(\Psi(K))} 1 = \int_{\Psi(K)} |\det \Gamma'| \\ &= \int_K |\det(\Gamma' \circ \Psi)| \cdot |\det \Psi'| \\ &= \int_K |\det((\Gamma' \circ \Psi) \cdot \Psi')| \\ &= \int_K |\det(\Gamma \circ \Psi)'| = \int_K |\det \Phi'|. \end{aligned}$$

□

Proposition 6.9.5 (Linear Change of Variable). *The Change of Variable Theorem holds for invertible linear mappings.*

Proof. Let

$$T : \mathbf{R}^n \longrightarrow \mathbf{R}^n$$

be an invertible linear mapping having matrix M . Thus $T'(x) = M$ for all x . Also, T is a composition of recombiners, scales, and transpositions, and so

by the persistence of the Change of Variable Theorem under composition it suffices to prove the theorem assuming that T is a recombine or a scale or a transposition. In each case, Propositions 6.9.1 and 6.9.3 allow us to assume that K is a box B and $f = 1$. Thus the desired result is simply

$$\text{vol}(T(B)) = |\det M| \cdot \text{vol}(B),$$

and we established this formula back in section 3.8. \square

The Change of Variable Theorem is proved partially by induction on the dimension n .

Proposition 6.9.6 (Base Case for the Induction). *The Change of Variable Theorem holds if $n = 1$.*

Proof. Because $n = 1$, K is an interval $[a, b] \subset \mathbf{R}$ where $a \leq b$. Here is where we use the hypothesis that K is connected. Since we have not studied connected sets closely, the reader is being asked to take for granted that any compact and connected subset of \mathbf{R} is a closed and bounded interval. (Or see exercise 6.9.1 for a proof that any compact and *path-connected* subset of \mathbf{R} is a closed and bounded interval.) The continuous function

$$\Phi' : [a, b] \longrightarrow \mathbf{R}$$

can take the value 0 only at a and b . Thus by the Intermediate Value Theorem, Φ' never changes sign on $[a, b]$. If $\Phi' \geq 0$ on $[a, b]$ then Φ is increasing, and so (using Theorem 6.4.3 for the second equality),

$$\int_{\Phi([a, b])} f = \int_{\Phi(a)}^{\Phi(b)} f = \int_a^b (f \circ \Phi) \cdot \Phi' = \int_{[a, b]} (f \circ \Phi) \cdot |\Phi'|.$$

If $\Phi' \leq 0$ on $[a, b]$ then Φ is decreasing, and so

$$\int_{\Phi([a, b])} f = \int_{\Phi(b)}^{\Phi(a)} f = - \int_{\Phi(a)}^{\Phi(b)} f = - \int_a^b (f \circ \Phi) \cdot \Phi' = \int_{[a, b]} (f \circ \Phi) \cdot |\Phi'|.$$

Thus in either case the desired result holds. \square

Proposition 6.9.7 (Bootstrap Induction Step). *For any $n > 1$, if the Change of Variable Theorem holds in dimension $n - 1$ then it holds in dimension n subject to the additional hypothesis that the transformation Φ fixes at least one coordinate.*

Proof. Propositions 6.9.1 and 6.9.3 allow us to assume that K is a box B , that Φ' is injective on B , that $\det \Phi' \neq 0$ on B , and that $f = 1$. Also, we may assume that the coordinate fixed by Φ is the last coordinate. There is a box $B_{n-1} \subset \mathbf{R}^{n-1}$ and an interval $I = [a, b] \subset \mathbf{R}$ such that

$$B = \bigcup_{t \in I} B_{n-1} \times \{t\}.$$

By assumption, Φ is a \mathcal{C}^1 -mapping on an open superset A of B . For each $t \in I$ let A_t denote the cross-section of A with last coordinate t ,

$$A_t = \{x \in \mathbf{R}^{n-1} : (x, t) \in A\}.$$

Then A_t is an open superset of B_{n-1} in \mathbf{R}^{n-1} . For each $t \in I$ define a mapping

$$\Psi_t : A_t \longrightarrow \mathbf{R}^{n-1}, \quad \Psi_t(x) = (\Phi_1(x, t), \dots, \Phi_{n-1}(x, t)).$$

Each Ψ_t is a \mathcal{C}^1 -mapping on an open superset of B_{n-1} , and

$$\Phi(B) = \bigcup_{t \in I} \Psi_t(B_{n-1}) \times \{t\}.$$

Since Φ is injective on B and $\det \Phi' \neq 0$ on B , it follows that each Ψ_t is injective on B_{n-1} , and the formula

$$|\det \Psi'_t(x)| = |\det \Phi'(x, t)|, \quad (x, t) \in B \quad (6.16)$$

(exercise 6.9.3) shows that $\det \Psi'_t \neq 0$ on B_{n-1} . Thus for each t , the set B_{n-1} and the transformation Ψ_t satisfy the Change of Variable Theorem hypotheses in dimension $n - 1$. Compute, using Fubini's Theorem, quoting the Change of Variable Theorem in dimension $n - 1$, and citing formula (6.16) and again using Fubini's Theorem, that

$$\int_{\Phi(B)} 1 = \int_{t \in I} \int_{\Psi_t(B_{n-1})} 1 = \int_{t \in I} \int_{B_{n-1}} |\det \Psi'_t| = \int_B |\det \Phi'|.$$

□

At long last we can prove the Change of Variable Theorem for $n > 1$.

Proof. We may assume the result for dimension $n - 1$, and we may assume that K is a box B , that A is an open superset of B , and that $\Phi : A \rightarrow \mathbf{R}^n$ is a \mathcal{C}^1 -mapping such that Φ is injective on A and $\det \Phi' \neq 0$ on A . We need to show that

$$\int_{\Phi(B)} 1 = \int_B |\det \Phi'|. \quad (6.17)$$

To prove the theorem, we will partition B into subboxes J , each J having an open superset A_J on which Φ is a composition

$$\Phi = T \circ \Gamma \circ \Psi,$$

where Ψ and Γ are \mathcal{C}^1 -mappings that fix at least one coordinate and T is a linear transformation. Note that Ψ , Γ , and T inherit injectivity and nonzero

determinant-derivatives from Φ , so that in particular T is invertible. Since the theorem holds for each of Ψ , Γ , and T , it holds for their composition. In more detail,

$$\begin{aligned}
 \int_{T(\Gamma(\Psi(J)))} 1 &= \int_{\Gamma(\Psi(J))} |\det T'| && \text{by Proposition 6.9.5} \\
 &= \int_{\Psi(J)} |\det(T' \circ \Gamma)| |\det \Gamma'| && \text{by Proposition 6.9.7} \\
 &= \int_{\Psi(J)} |\det(T \circ \Gamma)'| && \text{by the Chain Rule} \\
 &= \int_J |\det((T \circ \Gamma)' \circ \Psi)| |\det \Psi'| && \text{by Proposition 6.9.7} \\
 &= \int_J |\det(T \circ \Gamma \circ \Psi)'| && \text{by the Chain Rule.}
 \end{aligned}$$

That is, for each J ,

$$\int_{\Phi(J)} 1 = \int_J |\det \Phi'|,$$

and so summing over all subboxes J finally gives (6.17).

To obtain the subboxes J , proceed as follows. For each point $x \in B$, let

$$T = D\Phi_x^{-1}$$

and define

$$\tilde{\Phi} = T^{-1} \circ \Phi,$$

so that $D\tilde{\Phi}_x = \text{id}_n$ is the n -dimensional identity map. Let id denote the one-dimensional identity map and further define

$$\Psi : A \longrightarrow \mathbf{R}^n, \quad \Psi = (\tilde{\Phi}_1, \dots, \tilde{\Phi}_{n-1}, id),$$

so that $D\Psi_x = \text{id}_n$ as well. By the Inverse Function Theorem, Ψ is locally invertible. Let J_x be a subbox of B containing x having an open superset A_x such that Ψ^{-1} exists on $\Psi(A_x)$. Now define

$$\Gamma : \Psi(A_x) \longrightarrow \mathbf{R}^n, \quad \Gamma = (id, \dots, id, \tilde{\Phi}_n \circ \Psi^{-1}).$$

Thus $\Phi = T \circ \Gamma \circ \Psi$ on A_x , and Ψ , Γ , and T have the desired properties.

Cover B by the collection of open interiors of the boxes J_x . By the finiteness property of B , some finite collection of the interiors covers B , and so certainly the corresponding finite collection of the boxes J_x themselves covers B . Partition B into subboxes J so that each J lies in one of the finitely many J_x , and the process is complete. \square

Exercises

6.9.1. Let K be a nonempty compact subset of \mathbf{R} . Explain why the quantities $a = \min\{x : x \in K\}$ and $b = \max\{x : x \in K\}$ exist. Now further assume that K is path-connected, so that in particular there is a continuous function

$$\gamma : [0, 1] \longrightarrow \mathbf{R}$$

such that $\gamma(0) = a$ and $\gamma(1) = b$. Explain why consequently $K = [a, b]$.

6.9.2. (a) Explain to yourself as necessary why the three types of rectangle in the proof of Proposition 6.9.1 are exclusive. Now suppose that the three types are not exhaustive, i.e., some rectangle J lies partly in K° and partly in $(B \setminus K)^\circ$ without meeting the set $\partial K = \partial(B \setminus K)$. Supply details as necessary for the following argument. Let $x \in J$ lie in K° and let $\tilde{x} \in J$ lie in $(B \setminus K)^\circ$. Define a function from the unit interval to \mathbf{R} by mapping the interval to the line segment from x to \tilde{x} , and then mapping each point of the segment to 1 if it lies in K and to -1 if it lies in $B \setminus K$. The resulting function is continuous on the interval, and it changes sign on the interval, but it does not take the value 0. This is impossible, so the rectangle J can not exist.

(b) In the proof of Proposition 6.9.1, show that we may assume that the partition P is fine enough that all subboxes J of type I and type II lie in U .

(c) In the proof of Proposition 6.9.1, show that given $\varepsilon > 0$, we may assume that the partition P is fine enough that

$$\sum_{J:\text{type II}} \text{vol}(J) < \min \left\{ \frac{\varepsilon}{R(2nc)^n}, \frac{\varepsilon}{RR} \right\}.$$

6.9.3. In the proof of Proposition 6.9.7, establish formula (6.16).

6.9.4. Here is a sketched variant of the endgame of the Change of Variable proof: *A slightly easier variant of Proposition 6.9.7 assumes that the transformation Φ changes at most one coordinate, and then the process of factoring Φ locally as a composition can be iterated until each factor is either linear or changes at most one coordinate.* Fill in the details.

6.10 Summary

Integration is a synthesis of many small elements into a whole. The integral of a continuous function on a reasonably-shaped region exists, as is shown by careful management of details and the technical point that continuity is uniform on compact sets. Fubini's Theorem, which is so apparent intuitively, is also easy to prove essentially from the definitions. However, the Change of Variable Theorem, which also seems plausible, requires a more elaborate proof.

Approximation by Smooth Functions

Let k be a nonnegative integer. Recall that a \mathcal{C}^k -function on \mathbf{R}^n is a function all of whose partial derivatives up to order k exist and are continuous. That is, to say that a function

$$f : \mathbf{R}^n \longrightarrow \mathbf{R}$$

is \mathcal{C}^k is to say that f , and $D_j f$ for $j = 1, \dots, n$, and $D_{jj'} f$ for $j, j' = 1, \dots, n$, and so on up to all $D_{j_1 \dots j_k} f$ all exist and are continuous on \mathbf{R}^n . Various ideas that we have discussed so far have required different values of k :

- If f is \mathcal{C}^1 then f is differentiable in the multivariable sense of derivative (Theorem 4.3.3).
- If f is \mathcal{C}^2 then its mixed second order derivatives $D_{12}f$ and $D_{21}f$ are equal (Theorem 4.4.1).
- The multivariable max/min test (Proposition 4.5.8) assumes a \mathcal{C}^2 -function.
- If $f : \mathbf{R}^n \longrightarrow \mathbf{R}^n$ is componentwise \mathcal{C}^1 and its derivative Df_a is invertible at a point a then f is locally invertible about a , and the local inverse is again \mathcal{C}^1 (Theorem 5.1.2).
- If f (again scalar-valued now) is \mathcal{C}^0 then it is integrable over any compact set having boundary of volume zero (section 6.5).
- In the Change of Variable formula $\int_{\Phi(K)} f = \int_K (f \circ \Phi) \cdot |\det \Phi'|$ for multiple integrals (Theorem 6.7.1) the change of variable mapping Φ is assumed to be \mathcal{C}^1 and for now the integrand f is assumed to be \mathcal{C}^0 . We will return to this example at the very end of the chapter.

Meanwhile, a *smooth* function is a function on \mathbf{R}^n all of whose partial derivatives of all orders exist. Smooth functions are also called \mathcal{C}^∞ -functions, an appropriate notation because the derivatives of each order are continuous since the derivatives of one-higher order exist. This chapter briefly touches on the fact that for functions that vanish off a compact set, \mathcal{C}^0 -functions and \mathcal{C}^1 -functions and \mathcal{C}^2 -functions are well approximated by \mathcal{C}^∞ -functions.

The approximation technology is called *convolution*. One can see convolution in action visually by comparing graphs of convolutions against the graph

of the original function. And the conceptual framework for establishing the properties of convolution analytically is not difficult. Having discussed approximation by convolution, we will freely assume in the remaining chapters of these notes that our functions are C^∞ , i.e., that they are smooth.

7.1 Spaces of Functions

To begin, we quantify the phrase *functions that vanish off a compact set* from the chapter introduction.

Definition 7.1.1 (Support). *Consider a function*

$$f : \mathbf{R}^n \longrightarrow \mathbf{R}$$

The **support** of f is the closure of the set of its inputs that produce nonzero outputs,

$$\text{supp}(f) = \overline{\{x \in \mathbf{R}^n : f(x) \neq 0\}}.$$

The function f is **compactly supported** if its support is compact. The class of compactly supported C^k -functions is denoted $C_c^k(\mathbf{R}^n)$. Especially, $C_c^0(\mathbf{R}^n)$ denotes the class of compactly supported continuous functions.

Each class $C_c^k(\mathbf{R}^n)$ of functions forms a vector space over \mathbf{R} (exercise 7.1.1). Figure 7.1 shows a compactly supported C^0 -function on \mathbf{R} and its support. The graph has some corners, so the function is not C^1 .




Figure 7.1. Compactly supported continuous function on \mathbf{R} and its support

The spaces of compactly supported functions shrink as their member-functions are required to have more derivatives,

$$C_c^0(\mathbf{R}^n) \supset C_c^1(\mathbf{R}^n) \supset C_c^2(\mathbf{R}^n) \supset \cdots,$$

and we will see that all of the containments are proper.

Definition 7.1.2 (Test Function). A test function is a compactly supported smooth function. The class of test functions is denoted $\mathcal{C}_c^\infty(\mathbf{R}^n)$.

The class of test functions sits at the end of the chain of containments of function-spaces from a moment ago,

$$\mathcal{C}_c^\infty(\mathbf{R}^n) = \bigcap_{k \geq 0} \mathcal{C}_c^k(\mathbf{R}^n),$$

and as an intersection of vector spaces over \mathbf{R} , the test functions $\mathcal{C}_c^\infty(\mathbf{R}^n)$ again form a vector space over \mathbf{R} . In the chain of containments

$$\mathcal{C}_c^0(\mathbf{R}^n) \supset \mathcal{C}_c^1(\mathbf{R}^n) \supset \mathcal{C}_c^2(\mathbf{R}^n) \supset \cdots \supset \mathcal{C}_c^\infty(\mathbf{R}^n),$$

all of the containments are proper. Indeed, for the first containment, Weierstrass showed how to construct a function f of one variable, having support $[0, 1]$, that is continuous everywhere but differentiable nowhere on its support. The function of n variables

$$f_0(x_1, x_2, \dots, x_n) = f(x_1)$$

thus lies in $\mathcal{C}_c^0(\mathbf{R}^n)$ but not in $\mathcal{C}_c^1(\mathbf{R}^n)$. Next, the function

$$f_1(x_1, x_2, \dots, x_n) = \int_{t_1=0}^{x_1} f_0(t_1, x_2, \dots, x_n)$$

lies in $\mathcal{C}_c^1(\mathbf{R}^n)$ but not $\mathcal{C}_c^2(\mathbf{R}^n)$ because its first partial derivative is f_0 , which does not have a first partial derivative. Defining f_2 as a similar integral of f_1 gives a function that lies in $\mathcal{C}_c^2(\mathbf{R}^n)$ but not $\mathcal{C}_c^3(\mathbf{R}^n)$, and so on. Finally, none of the functions f_k just described lies in $\mathcal{C}_c^\infty(\mathbf{R}^n)$.

For any $k > 0$ and any $f \in \mathcal{C}_c^k(\mathbf{R}^n)$, the supports of the partial derivatives are contained in the support of the original function,

$$\text{supp}(D_j f) \subset \text{supp}(f), \quad j = 1, \dots, n.$$

Thus the partial derivative operators D_j take $\mathcal{C}_c^k(\mathbf{R}^n)$ to $\mathcal{C}_c^{k-1}(\mathbf{R}^n)$ as sets. The operators are linear because

$$D_j(f + \tilde{f}) = D_j f + D_j \tilde{f}, \quad f, \tilde{f} \in \mathcal{C}_c^k(\mathbf{R}^n)$$

and

$$D_j(cf) = cD_j f, \quad f \in \mathcal{C}_c^k(\mathbf{R}^n), \quad c \in \mathbf{R}.$$

In addition, more can be said about the D_j operators. Each space $\mathcal{C}_c^k(\mathbf{R}^n)$ of functions carries an absolute value function having properties similar to the absolute value on Euclidean space \mathbf{R}^n . With these absolute values in place, the partial differentiation operators are continuous.

Definition 7.1.3 ($C_c^k(\mathbf{R}^n)$ **Absolute Value**). *The absolute value function on $C_c^0(\mathbf{R}^n)$ is*

$$|\cdot| : C_c^0(\mathbf{R}^n) \longrightarrow \mathbf{R}, \quad |f| = \sup\{|f(x)| : x \in \mathbf{R}^n\}.$$

Let k be a nonnegative integer. The absolute value function on $C_c^k(\mathbf{R}^n)$ is

$$|\cdot|_k : C_c^k(\mathbf{R}^n) \longrightarrow \mathbf{R}$$

given by

$$|f|_k = \max \left\{ \begin{array}{l} |f|, \\ |D_j f| \text{ for } j = 1, \dots, n, \\ |D_{jj'} f| \text{ for } j, j' = 1, \dots, n, \\ \vdots \\ |D_{j_1 \dots j_k} f| \text{ for } j_1, \dots, j_k = 1, \dots, n \end{array} \right\}.$$

That is, $|f|_k$ is the largest absolute value of f or of any derivative of f up to order k . In particular, $|\cdot|_0 = |\cdot|$.

The largest absolute values mentioned in the definition exist by the Extreme Value Theorem since the relevant partial derivatives are compactly supported and continuous.

Proposition 7.1.4 ($C_c^k(\mathbf{R}^n)$ **Absolute Value Properties**).

- (A1) *Absolute value is positive: $|f|_k \geq 0$ for all $f \in C_c^k(\mathbf{R}^n)$, and $|f|_k = 0$ if and only if f is the zero function.*
- (A2) *Scaling Property: $|cf|_k = |c| |f|_k$ for all $c \in \mathbf{R}$ and $f \in C_c^k(\mathbf{R}^n)$.*
- (A3) *Triangle Inequality: $|f + g|_k \leq |f|_k + |g|_k$ for all $f, g \in C_c^k(\mathbf{R}^n)$.*

Proof. The first two properties are straightforward to check. For the third property, note that for any $f, g \in C_c^0(\mathbf{R}^n)$ and any $x \in \mathbf{R}^n$,

$$|(f + g)(x)| \leq |f(x)| + |g(x)| \leq |f| + |g|.$$

Thus $|f| + |g|$ is an upper bound of all values $|(f + g)(x)|$, so that

$$|f + g| \leq |f| + |g|.$$

That is, $|f + g|_0 \leq |f|_0 + |g|_0$. If $f, g \in C_c^1(\mathbf{R}^n)$ then the same argument shows that also $|D_j(f + g)| \leq |D_j f| + |D_j g|$ for $j = 1, \dots, n$, so that $|f + g|_1 \leq |f|_1 + |g|_1$. The argument for higher values of k is the same. \square

Now we can verify the anticipated continuity of the linear operators D_j from $C_c^k(\mathbf{R}^n)$ to $C_c^{k-1}(\mathbf{R}^n)$.

Proposition 7.1.5 (Continuity of Differentiation). *For any $k \geq 1$, the partial differentiation mappings*

$$D_j : \mathcal{C}_c^k(\mathbf{R}^n) \longrightarrow \mathcal{C}_c^{k-1}(\mathbf{R}^n), \quad j = 1, \dots, n$$

are continuous.

Proof. Consider any function $f \in \mathcal{C}_c^k(\mathbf{R}^n)$ and any sequence $\{f_m\}$ in $\mathcal{C}_c^k(\mathbf{R}^n)$. Suppose that

$$\lim_m |f_m - f|_k = 0.$$

Then

$$\begin{aligned} \lim_m |f_m - f| &= 0, \\ \lim_m |D_j f_m - D_j f| &= 0 \text{ for } j = 1, \dots, n, \\ \lim_m |D_{jj'} f_m - D_{jj'} f| &= 0 \text{ for } j, j' = 1, \dots, n, \\ &\vdots \\ \lim_m |D_{j_1 j_2 \dots j_k} f_m - D_{j_1 j_2 \dots j_k} f| &= 0 \text{ for } j_1, j_2, \dots, j_k = 1, \dots, n. \end{aligned}$$

Fix any $j \in \{1, \dots, n\}$. As a subset of the information in the previous display,

$$\begin{aligned} \lim_m |D_j f_m - D_j f| &= 0, \\ \lim_m |D_{jj'} f_m - D_{jj'} f| &= 0 \text{ for } j' = 1, \dots, n, \\ &\vdots \\ \lim_m |D_{jj_2 \dots j_k} f_m - D_{jj_2 \dots j_k} f| &= 0 \text{ for } j_2, \dots, j_k = 1, \dots, n. \end{aligned}$$

That is,

$$\lim_m |D_j f_m - D_j f|_{k-1} = 0.$$

The implication that we have just shown,

$$\lim_m |f_m - f|_k = 0 \implies \lim_m |D_j f_m - D_j f|_{k-1} = 0,$$

is exactly the assertion that $D_j : \mathcal{C}_c^k(\mathbf{R}^n) \longrightarrow \mathcal{C}_c^{k-1}(\mathbf{R}^n)$ is continuous, and the proof is complete. \square

Again let $k \geq 1$. The fact that $|f|_{k-1} \leq |f|_k$ for any $f \in \mathcal{C}_c^k(\mathbf{R}^n)$ (exercise 7.1.2) shows that for any $f \in \mathcal{C}_c^k(\mathbf{R}^n)$ and any sequence $\{f_m\}$ in $\mathcal{C}_c^k(\mathbf{R}^n)$, if $\lim_m |f_m - f|_k = 0$ then $\lim_m |f_m - f|_{k-1} = 0$. That is, the inclusion mapping

$$i : \mathcal{C}_c^k(\mathbf{R}^n) \longrightarrow \mathcal{C}_c^{k-1}(\mathbf{R}^n), \quad i(f) = f$$

is continuous.

The space $C_c^\infty(\mathbf{R}^n)$ of test functions is closed under partial differentiation, meaning that the partial derivatives of a test function are again test functions (exercise 7.1.3).

In this chapter we will show that just as any real number $x \in \mathbf{R}$ is approximated as closely as desired by rational numbers $q \in \mathbf{Q}$, any compactly supported continuous function $f \in C_c^k(\mathbf{R}^n)$ is approximated as closely as desired by test functions $g \in C_c^\infty(\mathbf{R}^n)$. More precisely, we will show that:

For any $f \in C_c^k(\mathbf{R}^n)$, there exists a sequence $\{f_m\}$ in $C_c^\infty(\mathbf{R}^n)$ such that $\lim_m |f_m - f|_k = 0$.

The fact that $\lim_m |f_m - f|_k = 0$ means that given any $\varepsilon > 0$, there exists a starting index m_0 such that f_m for all $m \geq m_0$ uniformly approximates f to within ε up to k th order. That is, for all $m \geq m_0$, simultaneously for all $x \in \mathbf{R}^n$,

$$\begin{aligned} |f_m(x) - f(x)| &< \varepsilon, \\ |D_j f_m(x) - D_j f(x)| &< \varepsilon \text{ for } j = 1, \dots, n, \\ |D_{jj'} f_m(x) - D_{jj'} f(x)| &< \varepsilon \text{ for } j, j' = 1, \dots, n, \\ &\vdots \\ |D_{j_1 \dots j_k} f_m(x) - D_{j_1 \dots j_k} f(x)| &< \varepsilon \text{ for } j_1, \dots, j_k = 1, \dots, n. \end{aligned}$$

The use of *uniform* here to connote that a condition holds simultaneously over a set of values is similar to its use in *uniform continuity*.

Exercises

7.1.1. Show that each class $C_c^k(\mathbf{R}^n)$ of functions forms a vector space over \mathbf{R} .

7.1.2. Verify that $|f|_{k-1} \leq |f|_k$ for any $f \in C_c^k(\mathbf{R}^n)$.

7.1.3. Explain why each partial derivative of a test function is again a test function.

7.2 Pulse Functions

A *pulse function* is a useful type of test function. To construct pulse functions, first consider the function

$$s : \mathbf{R} \rightarrow \mathbf{R}, \quad s(x) = \begin{cases} 0 & \text{if } x \leq 0, \\ e^{-1/x} & \text{if } x > 0. \end{cases}$$

(See figure 7.2.) Each $x < 0$ lies in an open interval on which s is the constant function 0, and each $x > 0$ lies in an open interval on which s is a composition

of smooth functions, so in either case all derivatives $s^{(k)}(x)$ exist. More specifically, for any nonnegative integer k , there exists a polynomial $p_k(x)$ such that the k th derivative of s takes the form

$$s^{(k)}(x) = \begin{cases} 0 & \text{if } x < 0, \\ p_k(x)x^{-2k}e^{-1/x} & \text{if } x > 0, \\ ? & \text{if } x = 0. \end{cases}$$

Only $s^{(k)}(0)$ is in question. However, $s^{(0)}(0) = 0$, and if we assume that $s^{(k)}(0) = 0$ for some $k \geq 0$ then it follows (because exponential behavior dominates polynomial behavior) that

$$\lim_{h \rightarrow 0^+} \frac{s^{(k)}(h) - s^{(k)}(0)}{h} = \lim_{h \rightarrow 0^+} p_k(h)h^{-2k-1}e^{-1/h} = 0.$$

That is, $s^{(k+1)}(0)$ exists and equals 0 as well. By induction, $s^{(k)}(0) = 0$ for all $k \geq 0$. Thus s is smooth: each derivative exists, and each derivative is continuous because the next derivative exists as well. But s is not a test function because its support is not compact: $\text{supp}(s) = [0, \infty)$.




Figure 7.2. Smooth function

Now the pulse function is defined in terms of the smooth function,

$$p : \mathbf{R} \longrightarrow \mathbf{R}, \quad p(x) = \frac{s(x+1)s(-x+1)}{\int_{x=-1}^1 s(x+1)s(-x+1)}.$$

The graph of p (figure 7.3) explains the name *pulse function*. As a product of compositions of smooth functions, p is smooth. The support of p is $[-1, 1]$, so p is a test function. Also, p is normalized so that

$$\int_{[-1,1]} p = 1.$$

The maximum pulse value $p(0)$ is therefore close to 1 because the pulse graph is roughly a triangle of base 2, but $p(0)$ is not exactly 1. The pulse function $p_2(x, y) = p(x)p(y)$ from \mathbf{R}^2 to \mathbf{R} , having support $[-1, 1]^2$, is shown in figure 7.4. A similar pulse function p_3 on \mathbf{R}^3 can be imagined as a concentration of density in a box about the origin.




Figure 7.3. Pulse function




Figure 7.4. Two-dimensional pulse function

Exercises

7.2.1. Since the function s in the section is smooth, it has n th degree Taylor polynomials $T_n(x)$ at $a = 0$ for all nonnegative integers n . (Here n does not denote the dimension of Euclidean space.) For what x does $s(x) = T_n(x)$?

7.2.2. Let p be the pulse function defined in the section. Explain why $\text{supp}(p) = [-1, 1]$.

7.2.3. Let $p : \mathbf{R} \rightarrow \mathbf{R}$ be the one-dimensional pulse function from the section.

- (a) Graph the function $q(x) = p(2a - b + x(b - a))$, where $a < b$.
- (b) Graph the function $r(x) = \int_{t=-1}^x p(t)$.
- (c) Use the function r from part (b) to give a formula for a test function that is 0 for $x < a$, climbs from 0 to 1 for $a \leq x \leq b$, is 1 for $b < x < c$, drops from 1 to 0 for $c \leq x \leq d$, and is 0 for $d < x$.

7.3 Convolution

This section shows how to construct test functions from $\mathcal{C}_c^0(\mathbf{R}^n)$ -functions. In preparation, we introduce a handy piece of notation.

Definition 7.3.1 (Sum, Difference of Two Sets). Let S and T be subsets of \mathbf{R}^n . Their **sum** is the set consisting of all sums of a point of S plus a point of T ,

$$S + T = \{s + t : s \in S, t \in T\}.$$

Their **difference** is similarly

$$S - T = \{s - t : s \in S, t \in T\}.$$

Visually, $S + T$ can be imagined as many copies of T , one based at each point of S , or vice versa. For example, if K is a three-dimensional box and B is a small ball about $\mathbf{0}_3$ then $K + B$ is slightly larger than K , again shaped like a box except that the edges and corners are rounded. Similarly, $\emptyset - T$ is the reflection of T through the origin. The sum or difference of two compact sets is compact (exercise 7.3.1(a)). The sum of the open balls $B(a, r)$ and $B(b, s)$ is $B(a+b, r+s)$ (exercise 7.3.1(b)). The reader is alerted that the set difference here is different from another, more common notion of set difference, that being the elements of one set that are not elements of another,

$$S \setminus T = \{s \in S : s \notin T\}.$$

Returning to $\mathcal{C}_c^0(\mathbf{R}^n)$ -functions, any such function can be integrated over all of \mathbf{R}^n .

Definition 7.3.2 (Integral of a \mathcal{C}_c^0 -Function). Let $f \in \mathcal{C}_c^0(\mathbf{R}^n)$. The **integral of f** is the integral of f over any box that contains its support,

$$\int f = \int_B f \quad \text{where } \text{supp}(f) \subset B.$$

In Definition 7.3.2 the integral on the right side exists by Theorem 6.3.1. Also, the integral on the right side is independent of the suitable box B , always being the integral over the intersection of all such boxes, the smallest suitable box. Thus the integral on the left side exists and is unambiguous. We do not bother writing $\int_{\mathbf{R}^n} f$ rather than $\int f$, because it is understood that by default we are integrating f over \mathbf{R}^n .

Definition 7.3.3 (Mollifying Kernel). Let $f \in \mathcal{C}_c^0(\mathbf{R}^n)$ be a compactly supported continuous function, and let $\varphi \in \mathcal{C}_c^\infty(\mathbf{R}^n)$ be a test function. The **mollifying kernel** associated to f and φ is the function

$$\kappa : \mathbf{R}^n \times \mathbf{R}^n \longrightarrow \mathbf{R}, \quad \kappa(x, y) = f(y)\varphi(x - y).$$

For any fixed $x \in \mathbf{R}^n$, the corresponding cross section of the mollifying kernel is denoted κ_x ,

$$\kappa_x : \mathbf{R}^n \longrightarrow \mathbf{R}, \quad \kappa_x(y) = \kappa(x, y).$$

For each $x \in \mathbf{R}^n$, the mollifying kernel $\kappa_x(y)$ can be nonzero only if $y \in \text{supp}(f)$ and $x - y \in \text{supp}(\varphi)$. It follows that

$$\text{supp}(\kappa_x) \subset \text{supp}(f) \cap (\{x\} - \text{supp}(\varphi)).$$

Therefore κ_x is compactly supported. (Figure 7.5 shows an example of the multiplicands $f(y)$ and $\varphi(x - y)$ of $\kappa_x(y)$, and figure 7.6 shows their compactly supported product.) Also, since f and φ are continuous, κ_x is continuous. That is, for each x , the mollifying kernel κ_x viewed as a function of y again lies in $\mathcal{C}_c^0(\mathbf{R}^n)$, making it integrable by Theorem 6.3.1.




Figure 7.5. Multiplicands of the mollifying kernel

The mollifying kernel is so named for good reason. First, it is a kernel in the sense that we integrate it to get a new function.

Definition 7.3.4 (Convolution). Let $f \in \mathcal{C}_c^0(\mathbf{R}^n)$ and let $\varphi \in \mathcal{C}_c^\infty(\mathbf{R}^n)$. The **convolution** of f and φ is the function defined by integrating the mollifying kernel,




Figure 7.6. The mollifying kernel is compactly supported

$$f * \varphi : \mathbf{R}^n \longrightarrow \mathbf{R}, \quad (f * \varphi)(x) = \int_y \kappa_x(y) = \int_y f(y)\varphi(x - y).$$

Second, although the mollifying kernel is only as well-behaved as f , integrating it indeed mollifies f in the sense that the integral is as well behaved as φ , i.e., the integral is a test function. Even if f is nowhere differentiable, $f * \varphi$ has all partial derivatives of all orders while remaining compactly supported. Furthermore, the derivatives have the natural formula obtained by passing them through the integral.

Proposition 7.3.5 (Derivatives of the Convolution). *Let $f \in \mathcal{C}_c^0(\mathbf{R}^n)$ and let $\varphi \in \mathcal{C}_c^\infty(\mathbf{R}^n)$. Then also $f * \varphi \in \mathcal{C}_c^\infty(\mathbf{R}^n)$. Specifically, the partial derivatives of the convolution are the convolutions with the partial derivatives,*

$$D_j(f * \varphi) = f * D_j\varphi, \quad j = 1, \dots, n,$$

and similarly for the higher order partial derivatives.

The following result helps to prove Proposition 7.3.5. In its statement, the symbol φ , which usually denotes a test function, instead denotes a $\mathcal{C}_c^1(\mathbf{R}^n)$ -function. The reason for the weaker hypothesis will appear soon in the proof of Corollary 7.3.7.

Lemma 7.3.6 (Uniformity Lemma for \mathcal{C}^1 -Functions). *Let $\varphi \in \mathcal{C}_c^1(\mathbf{R}^n)$. Given any $\varepsilon > 0$, there exists a corresponding $\delta > 0$ such that for all $a \in \mathbf{R}^n$ and all nonzero $h \in \mathbf{R}$, and for any $j \in \{1, \dots, n\}$,*

$$|h| < \delta \implies \left| \frac{\varphi(a + he_j) - \varphi(a)}{h} - D_j\varphi(a) \right| < \varepsilon.$$

Proof. Supposing first that for each $j \in \{1, \dots, n\}$ a corresponding δ_j exists, define $\delta = \min\{\delta_1, \dots, \delta_n\}$. Then for all nonzero $h \in \mathbf{R}$ and for any $j \in \{1, \dots, n\}$,

$$|h| < \delta \implies |h| < \delta_j.$$

Thus δ works simultaneously for all j . So it suffices to find δ_j for one fixed-but-arbitrary j , and from now on in this proof j is fixed

The Mean Value Theorem at the j th coordinate gives for all $a \in \mathbf{R}^n$ and all nonzero $h \in \mathbf{R}$,

$$\left| \frac{\varphi(a + he_j) - \varphi(a)}{h} - D_j\varphi(a) \right| = |D_j\varphi(a + te_j) - D_j\varphi(a)| \text{ where } |t| < |h|.$$

Since $D_j\varphi$ is continuous on \mathbf{R}^n and is compactly supported, it is uniformly continuous on \mathbf{R}^n , and so given any $\varepsilon > 0$ there exists a corresponding $\delta_j > 0$ such that for all $a \in \mathbf{R}^n$ and $t \in \mathbf{R}$,

$$|D_j\varphi(a + te_j) - D_j\varphi(a)| < \varepsilon \text{ if } |t| < \delta_j.$$

The desired result follows from the two displays. \square

Now we can establish the derivative formula for the convolution.

Proof (of Proposition 7.3.5). To see that $f * \varphi$ is compactly supported, recall the observation that for a given x , the mollifying kernel $\kappa_x(y) = f(y)\varphi(x - y)$ can be nonzero only at y -values such that

$$y \in \text{supp}(f) \cap (\{x\} - \text{supp}(\varphi)).$$

Such y can exist only if x takes the form

$$x = y + z, \quad y \in \text{supp}(f), \quad z \in \text{supp}(\varphi).$$

That is, the integrand is always zero if $x \notin \text{supp}(f) + \text{supp}(\varphi)$ (see figure 7.7). Hence,

$$\text{supp}(f * \varphi) \subset \text{supp}(f) + \text{supp}(\varphi).$$




Figure 7.7. The mollifying kernel is zero for x outside $\text{supp}(f) + \text{supp}(\varphi)$

To show that $D_j(f * \varphi)$ exists and equals $f * D_j\varphi$ for $j = 1, \dots, n$ is precisely to show that each x -derivative passes through the y -integral,

$$\frac{\partial}{\partial x_j} \int_y f(y)\varphi(x-y) = \int_y f(y) \frac{\partial \varphi}{\partial x_j}(x-y), \quad j = 1, \dots, n.$$

Since the integral is being taken over some box B , the equality follows from Proposition 6.6.2. But we prove it using other methods, for reasons that will emerge later in the chapter. The function f is bounded, say by R , so we can estimate that for any $x \in \mathbf{R}^n$ and any nonzero $h \in \mathbf{R}$ and any j ,

$$\begin{aligned} & \left| \frac{(f * \varphi)(x + he_j) - (f * \varphi)(x)}{h} - (f * D_j \varphi)(x) \right| \\ &= \left| \frac{\int_y f(y)\varphi(x + he_j - y) - \int_y f(y)\varphi(x - y)}{h} - \int_y f(y)D_j \varphi(x - y) \right| \\ &= \left| \int_y f(y) \left(\frac{\varphi(x - y + he_j) - \varphi(x - y)}{h} - D_j \varphi(x - y) \right) \right| \\ &\leq R \int_y \left| \frac{\varphi(x - y + he_j) - \varphi(x - y)}{h} - D_j \varphi(x - y) \right|. \end{aligned}$$

Assuming that $|h| < 1$, the support of the integrand as a function of y lies in the bounded set

$$\{x + te_j : -1 < t < 1\} - \text{supp}(\varphi),$$

and therefore the integral can be taken over some box B . By the Uniformity Lemma, given any $\varepsilon > 0$, for all small enough h the integrand is less than $\varepsilon/(R \text{vol}(B))$ uniformly in y . Consequently the integral is less than ε/R . In sum, given any $\varepsilon > 0$, for all small enough h we have

$$\left| \frac{(f * \varphi)(x + he_j) - (f * \varphi)(x)}{h} - (f * D_j \varphi)(x) \right| < \varepsilon.$$

Since x is arbitrary, this gives the desired result for first-order partial derivatives,

$$D_j(f * \varphi) = f * D_j \varphi, \quad j = 1, \dots, n.$$

As for higher-order partial derivatives, note that $D_j \varphi \in \mathcal{C}_c^\infty(\mathbf{R}^n)$ for each j . So the same result for second-order partial derivatives follows,

$$D_{jj'}(f * \varphi) = D_{j'}(f * D_j \varphi) = f * D_{jj'} \varphi, \quad j, j' = 1, \dots, n,$$

and so on. □

The proof of Proposition 7.3.5 required only that each κ_x be integrable, that f be bounded, and that φ lie in $\mathcal{C}_c^1(\mathbf{R}^n)$. We will make use of this observation in section 7.5.

If the function f lies in the subspace $\mathcal{C}_c^1(\mathbf{R}^n)$ of $\mathcal{C}_c^0(\mathbf{R}^n)$ then the partial derivatives of the convolution pass through the integral to f as well as to φ . That is, for differentiable functions, the derivative of the convolution is the convolution of the derivative.

Corollary 7.3.7. *Let $k \geq 1$, let $f \in \mathcal{C}_c^k(\mathbf{R}^n)$, and let $\varphi \in \mathcal{C}_c^\infty(\mathbf{R}^n)$. Then*

$$D_{j_1 \dots j_k}(f * \varphi) = D_{j_1 \dots j_k} f * \varphi, \quad j_1, \dots, j_k = 1, \dots, n.$$

Proof. Since

$$(f * \varphi)(x) = \int_y f(y)\varphi(x - y),$$

it follows by the Change of Variable Theorem (replace y by $x - y$) that also

$$(f * \varphi)(x) = \int_y f(x - y)\varphi(y).$$

Now the proof of the proposition works with the roles of f and φ exchanged to show that $D_j(f * \varphi) = D_j f * \varphi$ for $j = 1, \dots, n$. (Here is where it is relevant that the Uniformity Lemma requires only a $\mathcal{C}_c^1(\mathbf{R}^n)$ -function rather than a test function.) Similarly, if $f \in \mathcal{C}_c^2(\mathbf{R}^n)$ then because $D_j f \in \mathcal{C}_c^1(\mathbf{R}^n)$ for $j = 1, \dots, n$ it follows that

$$D_{jj'}(f * \varphi) = D_{jj'} f * \varphi, \quad j, j' = 1, \dots, n.$$

The argument for higher derivatives is the same. \square

Consider a function $f \in \mathcal{C}_c^0(\mathbf{R}^n)$. Now that we know that any convolution $f * \varphi$ (where $\varphi \in \mathcal{C}_c^\infty(\mathbf{R}^n)$) lies in $\mathcal{C}_c^\infty(\mathbf{R}^n)$, the next question is to what extent the test function $f * \varphi$ resembles the original compactly supported continuous function f . As already noted, for any x the integral

$$(f * \varphi)(x) = \int_y f(y)\varphi(x - y)$$

refers to values of f only on $\{x\} - \text{supp}(\varphi)$. Especially, if $\text{supp}(\varphi)$ is a small set about the origin then the convolution value $(f * \varphi)(x)$, depends only on the behavior of the original function f near x . The next section will construct useful test functions φ having small support, the idea being that convolutions $f * \varphi$ with such test functions will approximate the functions f being convolved. For example, in figure 7.5 $f(x)$ is small and positive, while the integral of the mollifying kernel shown in figure 7.6 is plausibly small and positive as well.

Exercises

7.3.1. (a) Show that the sum of two compact sets is compact.

(b) Let $B(a, r)$ and $B(b, s)$ be open balls. Show that their sum is $B(a + b, r + s)$.

(c) Recall that there are four standard axioms for addition, either in the context of a field or a vector space. Which of the four axioms are satisfied by set addition, and which are not?

(d) Let $0 < a < b$. Let A be the circle of radius b in the (x, y) -plane, centered at the origin. Let B be the closed disk of radius a in the (x, z) -plane, centered at $(b, 0, 0)$. Describe the sum $A + B$.

7.3.2. Let $f \in C_c^0(\mathbf{R}^n)$, and let $\varphi \in C_c^\infty(\mathbf{R}^n)$. Assume that $\varphi \geq 0$, i.e., all output values of φ are nonnegative, and assume that $\int \varphi = 1$. Suppose that R bounds f , meaning that $|f(x)| < R$ for all x . Show that also R bounds $f * \varphi$.

7.4 Smooth Approximate Identity and Convolution

Our next technical tool is a sequence of test functions whose graphs are ever taller and more narrow, each enclosing volume 1.

Definition 7.4.1 (Smooth Approximate Identity). *A smooth approximate identity is a sequence of test functions*

$$\{\varphi_m\} = \{\varphi_1, \varphi_2, \varphi_3, \dots\}$$

such that:

- (1) Each φ_m is nonnegative, i.e., each φ_m maps \mathbf{R}^n to $\mathbf{R}_{\geq 0}$.
- (2) Each φ_m has integral 1, i.e., $\int \varphi_m = 1$ for each m .
- (3) The supports of the φ_m shrink to $\{\mathbf{0}\}$, i.e.,

$$\text{supp}(\varphi_1) \supset \text{supp}(\varphi_2) \supset \dots, \quad \bigcap_{m=1}^{\infty} \text{supp}(\varphi_m) = \{\mathbf{0}\}.$$

We can construct a smooth approximate identity using the pulse function p from section 7.2. Define for $m = 1, 2, 3, \dots$

$$\varphi_m : \mathbf{R}^n \longrightarrow \mathbf{R}, \quad \varphi_m(x) = m^n p(mx_1) p(mx_2) \cdots p(mx_n).$$

Then $\text{supp}(\varphi_m) = [-1/m, 1/m]^n$ for each m . Here the coefficient m^n is chosen so that $\int \varphi_m = 1$ (exercise 7.4.1). Figure 7.8 shows the graphs of $\varphi_2, \varphi_4, \varphi_8$, and φ_{15} when $n = 1$. The first three graphs have the same vertical scale, but not the fourth. Figure 7.9 shows the graphs of φ_1 through φ_4 when $n = 2$, all having the same vertical scale.

The *identity* being approximated by the sequence of test functions $\{\varphi_m\}$ is the n -dimensional *Dirac delta function*, conceptually a unit point mass at the origin. That is, the n -dimensional Dirac function δ should have the properties that

$$\text{supp}(\delta) = \{\mathbf{0}\}, \quad \int \delta = 1.$$

No such function exists in the orthodox sense of the word *function*. But regardless of sense, for any function $f : \mathbf{R}^n \longrightarrow \mathbf{R}$ and any $x \in \mathbf{R}^n$, the mollifying kernel associated to f and δ ,

$$\kappa_x(y) = f(y)\delta(x - y),$$

is conceptually a point of mass $f(x)$ at each x . That is, its properties should be




Figure 7.8. The functions φ_2 , φ_4 , φ_8 , and φ_{15} from an approximate identity




Figure 7.9. The functions φ_1 through φ_4 from a two-dimensional approximate identity

$$\text{supp}(\kappa_x) = \{x\}, \quad (f * \delta)(x) = \int_y \kappa_x(y) = f(x).$$

Under a generalized notion of function, the Dirac delta makes perfect sense as an object called a *distribution*, defined by the integral in the previous display

but only for a limited class of functions:

$$\text{For all } x, \quad (f * \delta)(x) = f(x) \text{ for test functions } f.$$

Yes, now it is f that is restricted to be a test function. The reason for this is that δ is not a test function, not being a function at all, and to get a good theory of distributions such as the Dirac delta, we need to restrict the functions that they convolve with. In sum, the Dirac delta function is an identity in the sense that

$$f * \delta = f \text{ for test functions } f.$$

Distribution theory is beyond the scope of these notes, but we may conceive of the defining property of the Dirac delta function as the expected limiting behavior of any smooth approximate identity. That is, returning to the environment of $f \in C_c^0(\mathbf{R}^n)$ and taking any smooth approximate identity $\{\varphi_m\}$, we expect that

$$\lim_m (f * \varphi_m) = f \text{ for } C_c^0(\mathbf{R}^n)\text{-functions } f.$$

As explained in section 7.1, this limit will be uniform, meaning that the values $(f * \varphi_m)(x)$ will converge to $f(x)$ at one rate simultaneously for all x in \mathbf{R}^n . See exercise 7.4.3 for an example of nonuniform convergence.

For an example of convolution with elements on a smooth approximate identity, consider the sawtooth function

$$f : \mathbf{R} \longrightarrow \mathbf{R}, \quad f(x) = \begin{cases} |x| & \text{if } |x| \leq 1/4, \\ 1/2 - |x| & \text{if } 1/4 < |x| \leq 1/2, \\ 0 & \text{if } 1/2 < |x|. \end{cases}$$

Recall the smooth approximate identity $\{\varphi_m\}$ from after Definition 7.4.1. Figure 7.10 shows f and its convolutions with $\varphi_2, \varphi_4, \varphi_8$, and φ_{15} . The convolutions approach the original function while smoothing its corners, and the convolutions are bounded by the bound on the original function as shown in exercise 7.3.2. Also, the convolutions have larger supports than the original function, but the supports shrink toward the original support as m grows.

The following lemma says that if compact sets shrink to a point, then eventually they lie inside any given ball about the point. Specifically, the sets that we have in mind are the supports of a smooth approximate identity.

Lemma 7.4.2 (Shrinking Sets Lemma). *Let*

$$\{S_m\} = \{S_1, S_2, S_3, \dots\}$$

be a sequence of compact subsets of \mathbf{R}^n such that

$$S_1 \supset S_2 \supset S_3 \supset \dots, \quad \bigcap_{m=1}^{\infty} S_m = \{\mathbf{0}\}.$$




Figure 7.10. The sawtooth function convolved with various φ_m

Then for any $\delta > 0$ there exists some positive integer m_0 such that

$$\text{for all } m \geq m_0, \quad S_m \subset B(\mathbf{0}, \delta).$$

Proof. Let $\delta > 0$ be given. If no S_m lies in $B(\mathbf{0}, \delta)$ then there exist points

$$\begin{aligned} x_1 &\in S_1 \setminus B(\mathbf{0}, \delta), \\ x_2 &\in S_2 \setminus B(\mathbf{0}, \delta), \\ x_3 &\in S_3 \setminus B(\mathbf{0}, \delta), \end{aligned}$$

and so on. The sequence $\{x_m\}$ lies in S_1 , so it has a convergent subsequence. The containments $S_1 \supset S_2 \supset \cdots$ show that replacing the sequence by the subsequence preserves the displayed conditions, so we may assume that the original sequence converges. Let x denote its limit. For any $m \geq 1$, the terms of the sequence from index m onward lie in S_m , so $x \in S_m$. Thus $x \in \bigcap_m S_m = \{\mathbf{0}\}$, i.e., $x = \mathbf{0}$. But also, $|x_m| \geq \delta$ for each m , so $|x| \geq \delta$. This is a contradiction, so we are done. \square

The hypothesis of compactness is necessary in the Shrinking Sets Lemma (exercise 7.4.2).

Theorem 7.4.3 ($C_c^0(\mathbf{R}^n)$ -Approximation by Convolutions). *Let $f \in C_c^0(\mathbf{R}^n)$ and let $\{\varphi_m\} : \mathbf{R}^n \rightarrow \mathbf{R}$ be a smooth approximate identity. Given $\varepsilon > 0$, there exists a positive integer m_0 such that for all integers m ,*

$$m \geq m_0 \implies |f * \varphi_m - f| < \varepsilon.$$

*That is, the convolutions $f * \varphi_m$ converge uniformly to the original function f .*

Proof. Let $\varepsilon > 0$ be given. Since the support of f is compact, f is uniformly continuous on its support, and hence f is uniformly continuous on all of \mathbf{R}^n . So there exists some $\delta > 0$ such that for all $x, y \in \mathbf{R}^n$,

$$|y - x| < \delta \implies |f(y) - f(x)| < \varepsilon.$$

Since the supports of the approximate identity functions shrink to $\{\mathbf{0}\}$, the Shrinking Sets Lemma says that there exists some positive integer m_0 such that for all integers $m \geq m_0$, $\text{supp}(\varphi_m) \subset B(\mathbf{0}, \delta)$. Note that m_0 depends only on δ , which in turn depends only on ε , all of this with no reference to any particular $x \in \mathbf{R}^n$. Now, for all $x, y \in \mathbf{R}^n$, and all $m \geq m_0$,

$$\begin{aligned} y \in x - \text{supp}(\varphi_m) &\implies y \in x - B(\mathbf{0}, \delta) = x + B(\mathbf{0}, \delta) \\ &\implies |y - x| < \delta \\ &\implies |f(y) - f(x)| < \varepsilon. \end{aligned}$$

Because the approximate identity functions φ_m have integral 1, we have for all $x \in \mathbf{R}^n$ and all positive integers m ,

$$f(x) = \int_y f(x) \varphi_m(x - y).$$

Use the fact that the approximate identity functions φ_m are nonnegative to estimate that for all $x \in \mathbf{R}^n$ and all positive integers m ,

$$\begin{aligned} |(f * \varphi_m)(x) - f(x)| &= \left| \int_y (f(y) - f(x)) \varphi_m(x - y) \right| \\ &\leq \int_y |f(y) - f(x)| \varphi_m(x - y). \end{aligned}$$

We may integrate only over y -values in $x - \text{supp}(\varphi_m)$, so that if $m \geq m_0$ then the integrand is less than $\varepsilon \varphi_m(x - y)$. That is, since the approximate identity functions have integral 1 we have for all $x \in \mathbf{R}^n$ and all positive integers m ,

$$m \geq m_0 \implies |(f * \varphi_m)(x) - f(x)| < \varepsilon \int_y \varphi_m(x - y) = \varepsilon.$$

This is the desired result. Note how the argument has used all three defining properties of the approximate identity. \square

Corollary 7.4.4 ($\mathcal{C}_c^k(\mathbf{R}^n)$ -Approximation by Convolutions). *Let k be a positive integer. Let $f \in \mathcal{C}_c^k(\mathbf{R}^n)$ and let $\{\varphi_m\} : \mathbf{R}^n \rightarrow \mathbf{R}$ be a smooth approximate identity. Given $\varepsilon > 0$, there exists a positive integer m_0 such that for all integers m ,*

$$m \geq m_0 \implies |f * \varphi_m - f|_k < \varepsilon.$$

That is, the convolutions and their derivatives converge uniformly to the original function and its derivatives up to order k .

Proof. Recall from Corollary 7.3.7 that if $f \in \mathcal{C}_c^1(\mathbf{R}^n)$ then for any test function φ the derivative of the convolution is the convolution of the derivative,

$$D_j(f * \varphi) = D_j f * \varphi, \quad j = 1, \dots, n.$$

Since the derivatives $D_j f$ lie in $\mathcal{C}_c^0(\mathbf{R}^n)$, the theorem says that their convolutions $D_j f * \varphi_m$ converge uniformly to the derivatives $D_j f$ as desired. The argument for higher derivatives is the same. \square

Exercises

7.4.1. Recall that $\int p = 1$ where $p : \mathbf{R}^n \rightarrow \mathbf{R}$ is the pulse function from section 7.2. Let m be any positive integer and recall the definition in the section,

$$\varphi_m(x) = m^n p(mx_1) p(mx_2) \cdots p(mx_n).$$

Explain why consequently $\int \varphi_m = 1$.

7.4.2. Find a sequence $\{S_m\}$ of subsets of \mathbf{R} satisfying all of the hypotheses of the Shrinking Sets Lemma except for compactness, and such that no S_m is a subset of the interval $B(0, 1) = (-1, 1)$.

7.4.3. This exercise illustrates a nonuniform limit. For each positive integer m , define

$$f_m : [0, 1] \rightarrow \mathbf{R}, \quad f_m(x) = x^m.$$

Also define

$$f : [0, 1] \rightarrow \mathbf{R}, \quad f(x) = \begin{cases} 0 & \text{if } 0 \leq x < 1, \\ 1 & \text{if } x = 1. \end{cases}$$

- (a) Using one set of axes, graph f_1, f_2, f_3, f_{10} , and f .
 (b) Show that for any $x \in [0, 1]$, $\lim_m f_m(x) = f(x)$. That is, given $\varepsilon > 0$, there exists some positive integer m_0 such that for all positive integers m ,

$$m \geq m_0 \implies |f_m(x) - f(x)| < \varepsilon.$$

Thus the function f is the limit of the sequence of functions $\{f_m\}$. That is:

$$\text{For each } x, \quad f(x) = \lim_m \{f_m(x)\}.$$

(c) Now let $\varepsilon = 1/2$. Show that for any positive integer m , no matter how large, there exists some corresponding $x \in [0, 1]$ such that $|f_m(x) - f(x)| \geq \varepsilon$. That is:

$$\text{For each } m, \quad |f_m(x) - f(x)| \text{ fails to be small for some } x.$$

Thus the convergence of $\{f_m\}$ to f is not uniform, i.e., the functions do not converge to the limit-function at one rate simultaneously for all $x \in [0, 1]$.

7.5 Known-Integrable Functions

Recall that the slogan-title of Theorem 6.5.4 is *near-continuity implies integrability*. The largest space of functions that we have considered so far in this chapter is $C_c^0(\mathbf{R}^n)$, so we have not yet discussed the entire class of functions that we know to be integrable. This section gives some results about convolution and approximation for such functions.

Recall also that a function is called bounded if its outputs form a bounded set.

Definition 7.5.1 (Known-Integrable Function). *A function*

$$f : \mathbf{R}^n \longrightarrow \mathbf{R}.$$

is known-integrable if it is bounded, compactly supported, and continuous except on a set of volume zero. The class of known-integrable functions is denoted $\mathcal{I}_c(\mathbf{R}^n)$.

Unsurprisingly, the class $\mathcal{I}_c(\mathbf{R}^n)$ forms a vector space over \mathbf{R} .

Let $f \in \mathcal{I}_c(\mathbf{R}^n)$. The integral of f is the integral of f over any box that contains its support,

$$\int f = \int_B f \quad \text{where } \text{supp}(f) \subset B.$$

Similarly to the remarks after Definition 7.3.2, the integral on the right side exists, but this time by Theorem 6.5.4. The integral on the right side is independent of the box B , and so the integral on the left side exists, is unambiguous, and is understood to be the integral of f over all of \mathbf{R}^n .

The convolution remains sensible when f is known-integrable. That is, if $f \in \mathcal{I}_c(\mathbf{R}^n)$ and $\varphi \in C_c^\infty(\mathbf{R}^n)$ then for each $x \in \mathbf{R}^n$ the mollifying kernel

$$\kappa_x : \mathbf{R}^n \longrightarrow \mathbf{R}, \quad \kappa_x(y) = f(y)\varphi(x - y)$$

again lies in $\mathcal{I}_c(\mathbf{R}^n)$. And so we may continue to define the convolution of f and φ as

$$f * \varphi : \mathbf{R}^n \longrightarrow \mathbf{R}, \quad (f * \varphi)(x) = \int_y \kappa_x(y).$$

The formulas for convolution derivatives remain valid as well. That is, if $f \in \mathcal{I}_c(\mathbf{R}^n)$ and $\varphi \in C_c^\infty(\mathbf{R}^n)$ then also $f * \varphi \in C_c^\infty(\mathbf{R}^n)$, and

$$\begin{aligned} D_j(f * \varphi) &= f * \varphi_j, & j = 1, \dots, n, \\ D_{jj'}(f * \varphi) &= f * D_{jj'}\varphi_j, & j, j' = 1, \dots, n, \end{aligned}$$

and so on. Here is where it is relevant that our proof of Proposition 7.3.5 required only that each κ_x be integrable, that f be bounded, and that φ lie in $C_c^1(\mathbf{R}^n)$.

Given a known-integrable function $f \in \mathcal{I}_c(\mathbf{R}^n)$ and a smooth approximate identity $\{\varphi_m\}$, we would like the convolutions $\{f * \varphi_m\}$ to approximate f uniformly as m grows. But the following proposition shows that this is impossible when f has discontinuities.

Proposition 7.5.2 (The Uniform Limit of Continuous Functions is Continuous). *Let*

$$\{f_m\} : \mathbf{R}^n \longrightarrow \mathbf{R}$$

be a sequence of continuous functions that converges uniformly to a limit function

$$f : \mathbf{R}^n \longrightarrow \mathbf{R}.$$

Then f is continuous as well.

Proof. For any two points $x, \tilde{x} \in \mathbf{R}^n$ and for any positive integer m we have

$$\begin{aligned} |f(\tilde{x}) - f(x)| &= |f(\tilde{x}) - f_m(\tilde{x}) + f_m(\tilde{x}) - f_m(x) + f_m(x) - f(x)| \\ &\leq |f(\tilde{x}) - f_m(\tilde{x})| + |f_m(\tilde{x}) - f_m(x)| + |f_m(x) - f(x)|. \end{aligned}$$

Let $\varepsilon > 0$ be given. For all m large enough, the first and third terms are less than $\varepsilon/3$ regardless of the values of x and \tilde{x} . Fix such a value of m , and fix x . Then since f_m is continuous, the middle term is less than $\varepsilon/3$ if \tilde{x} is close enough to x . It follows that

$$|f(\tilde{x}) - f(x)| < \varepsilon \quad \text{for all } \tilde{x} \text{ close enough to } x.$$

That is, f is continuous. □

Thus the convergence property of convolutions must become more technical for known-integrable functions rather than compactly supported continuous functions. In preparation for proving the convergence property, the following lemma says that if K is a compact subset of an open set then so is the sum of K and some closed ball.

Lemma 7.5.3 (Thickening Lemma). *Let K and A be subsets of \mathbf{R}^n such that*

$$K \subset A, \quad K \text{ is compact, } \quad A \text{ is open.}$$

Then

$$\text{for some } r > 0, \quad K + \overline{B(\mathbf{0}, r)} \subset A.$$

Proof. Since K is compact, it lies in some ball $B(\mathbf{0}, R)$. Solving the problem with the open set $A \cap B(\mathbf{0}, R)$ in place of A also solves the original problem.

Having replaced A by $A \cap B(\mathbf{0}, R)$, define a function on K that takes positive real values,

$$d : K \longrightarrow \mathbf{R}_{>0}, \quad d(a) = \sup\{r : B(a, r) \subset A\}.$$

The fact that we have shrunk A (if necessary) to lie inside the ball has ensured that d is finite because specifically $d(a) \leq R$ for all a . Fix some $a \in K$ and let $r = d(a)$. Let $\{r_m\}$ be a strictly increasing sequence of positive real numbers such that $\lim_m \{r_m\} = r$. Then $B(a, r_m) \subset A$ for each m , and so

$$B(a, r) = \bigcup_{m=1}^{\infty} B(a, r_m) \subset A.$$

This argument shows that in fact

$$d(a) = \max\{r : B(a, r) \subset A\}.$$

The function d is continuous. To see this, fix some point $a \in K$ and let $r = d(a)$. Consider also a second point $\tilde{a} \in K$ such that $|\tilde{a} - a| < r$, and let $\tilde{r} = d(\tilde{a})$. Then

$$B(\tilde{a}, r - |\tilde{a} - a|) \subset B(a, r) \subset A,$$

showing that $\tilde{r} \geq r - |\tilde{a} - a|$. Either $\tilde{r} \leq r + |\tilde{a} - a|$, or $\tilde{r} > r + |\tilde{a} - a| \geq r$ so that also $|\tilde{a} - a| < \tilde{r}$ and the same argument shows that $r \geq \tilde{r} - |\tilde{a} - a|$, i.e., $\tilde{r} \leq r + |\tilde{a} - a|$ after all. That is, we have shown that for any $a \in K$,

$$\left\{ \begin{array}{l} \tilde{a} \in K \\ |\tilde{a} - a| < r(a) \end{array} \right\} \implies |d(\tilde{a}) - d(a)| \leq |\tilde{a} - a|.$$

Thus d is continuous at a (given $\varepsilon > 0$, let $\delta = \min\{r(a), \varepsilon/2\}$), and since $a \in K$ is arbitrary, d is continuous on K as claimed.

Since K is compact and d is continuous, d takes a minimum value $\tilde{r} > 0$. Thus $K + \overline{B(\mathbf{0}, \tilde{r})} \subset A$. Finally, let $r = \tilde{r}/2$. Then $K + \overline{B(\mathbf{0}, r)} \subset A$ as desired. \square

Now we can establish the convergence property of convolutions for known-integrable functions.

Theorem 7.5.4 ($\mathcal{I}_c(\mathbf{R}^n)$ -Approximation by Convolutions). *Let $f \in \mathcal{I}_c(\mathbf{R}^n)$ and let $\{\varphi_m\} : \mathbf{R}^n \rightarrow \mathbf{R}$ be a smooth approximate identity. Let K be a compact subset of \mathbf{R}^n such that f is continuous on an open superset of K . Given $\varepsilon > 0$, there exists a positive integer m_0 such that for all integers m ,*

$$m \geq m_0 \implies |(f * \varphi_m)(x) - f(x)| < \varepsilon \text{ for all } x \in K.$$

That is, the convolutions converge uniformly to the original function on compact sets where the function is continuous.

Proof. Let $\varepsilon > 0$ be given. By the Thickening Lemma, there exists some $r > 0$ such that f is continuous on $K + \overline{B(\mathbf{0}, r)}$. Hence f is uniformly continuous, on $K + \overline{B(\mathbf{0}, r)}$. That is, there exists $\delta > 0$ (with $\delta < r$) such that for all $x \in K$ and all $y \in \mathbf{R}^n$,

$$|y - x| < \delta \implies |f(y) - f(x)| < \varepsilon.$$

There exists some positive integer m_0 such that for all integers $m \geq m_0$, $\text{supp}(\varphi_m) \subset B(\mathbf{0}, \delta)$. For all $x \in K$, all $y \in \mathbf{R}^n$, and all $m \geq m_0$,

$$\begin{aligned} y \in x - \text{supp}(\varphi_m) &\implies y \in x - B(\mathbf{0}, \delta) = x + B(\mathbf{0}, \delta) \\ &\implies |y - x| < \delta \\ &\implies |f(y) - f(x)| < \varepsilon. \end{aligned}$$

From here, the proof is virtually identical to the proof of Theorem 7.4.3. \square

For example, consider the truncated squaring function

$$f : \mathbf{R} \longrightarrow \mathbf{R}, \quad f(x) = \begin{cases} x^2 & \text{if } |x| \leq 1/2, \\ 0 & \text{if } 1/2 < |x|. \end{cases}$$

Note that f lies in $\mathcal{I}_c(\mathbf{R}^n)$ rather than in $\mathcal{C}_c^0(\mathbf{R}^n)$ because of its discontinuities at $x = \pm 1/2$. Figure 7.11 shows f and its convolutions with $\varphi_2, \varphi_4, \varphi_8,$ and φ_{15} . The convolutions converge uniformly to the truncated parabola on compact sets away from the two points of discontinuity. But the convergence is not well behaved at or near those two points. Indeed, the function value $f(\pm 1/2) = 1/4$ rather than $f(\pm 1/2) = 0$ is arbitrary and has no effect on the convolution in any case. And again the convolutions are bounded by the bound on the original function and their supports shrink toward the original support as m grows.




Figure 7.11. The truncated squaring function convolved with various φ_m

In consequence of $\mathcal{I}_c(\mathbf{R}^n)$ -approximation by convolutions, any integral of a known-integrable function is approximated as closely as desired by the integral of a test function. Thus the hypothesis of a continuous integrand f in the Change of Variable Theorem for multiple integrals (Theorem 6.7.1), mentioned in the last bullet of the chapter introduction, can now be weakened to a known-integrable integrand.

Integration of Differential Forms

The integration of differential forms over surfaces is characteristic of a fully developed mathematical theory: it starts from carefully preconfigured definitions and proceeds to one central theorem, whose proof is purely mechanical because of how the definitions are rigged. Furthermore, much of the work is algebraic, even though the theorem appears analytical. Since the motivation for the definitions is not immediately obvious, the early stages of working through such a body of material can feel unenlightening, but the payoff lies in the lucidity of the later arguments and the power of the end result. The main theorem here is called often called Stokes's Theorem, but in fact it is a generalization not only of the classical Stokes's Theorem (which is not due to Stokes; he just liked to put it on his exams), but also of other nineteenth century results called the Divergence Theorem (or Gauss' Theorem) and Green's Theorem, and even of the Fundamental Theorem of Integral Calculus. In fact, a better name for the theorem to be presented here is the Generalized FTIC.

The definitions of a surface and of the integral of a function over a surface are given in section 8.1. Formulas for particular integrals called flow and flux integrals are derived in section 8.2. The theory to follow is designed partly to handle such integrals easily. The definitions of a differential form and of the integral of a differential form over a surface are given in section 8.3, and the definitions are illustrated by examples in sections 8.4 and 8.5. Sections 8.6 through 8.9 explain the algebraic rules of how to add differential forms and multiply them by scalars, how to multiply differential forms, how to differentiate them, and how to pass them through changes of variable. A Change of Variable Theorem for differential forms follows automatically in section 8.10. Returning to surfaces, sections 8.11 and 8.12 define a special class of surfaces called cubes, and a geometric boundary operator from cubes to cubes of lower dimension. The General FTIC is proved in section 8.13. Finally, section 8.14 explains how the classical vector integration theorems are special cases of the General FTIC, and section 8.15 takes a closer look at some of the quantities that arise in this context.

8.1 Integration of Functions Over Surfaces

Having studied integration over solid regions in \mathbf{R}^n , i.e., over subsets of \mathbf{R}^n with positive n -dimensional volume, we face the new problem of how to integrate over surfaces of lower dimension in \mathbf{R}^n . For example the circle in \mathbf{R}^2 is one-dimensional, and the torus surface in \mathbf{R}^3 is two-dimensional. Each of these sets has volume zero as a subset of its ambient space, in which it is curving around. In general, whatever the yet-undefined notion of a k -dimensional subset of \mathbf{R}^n means, such objects will have volume zero when $k < n$, and so any attempt to integrate over them in the sense of chapter 6 will give an integral of zero and a dull state of affairs. Instead, the idea is to *parameterize* surfaces in \mathbf{R}^n and then define integration over a parameterized surface in terms of integration over a non-curved parameter space.

Definition 8.1.1 (Parameterized Surface). *Let A be an open subset of \mathbf{R}^n . A k -surface in A is a smooth mapping*

$$\Phi : D \rightarrow A,$$

where D is a compact connected subset of \mathbf{R}^k whose boundary has volume zero. The set D is called the **parameter domain** of Φ .

See figure 8.1. Here are some points to note about Definition 8.1.1:

- Recall that a subset A of \mathbf{R}^n is called *open* if its complement is closed. The definitions in this chapter need the environment of an open subset rather than all of \mathbf{R}^n in order to allow for functions that are not defined everywhere. For instance, the reciprocal modulus function is defined only on surfaces that avoid the origin. In most of the examples, A will be all of \mathbf{R}^n , but exercise 8.8.7 will touch on how the subject becomes more nuanced when it is not.
- Recall also that *compact* means closed and bounded. *Connected* means that D consists of only one piece, as discussed informally in section 2.4. And as discussed informally in section 6.5 and formally in section 6.8, the boundary of a set consists of all points simultaneously near the set and near its complement—roughly speaking, its edge. Typically D will be some region that is easy to integrate over, such as a box, whose compactness, connectedness, and small boundary are self-evident.
- The word *smooth* in the definition means that the mapping Φ extends to some open superset of D in \mathbf{R}^k , on which it has continuous partial derivatives of all orders. Each such partial derivative is therefore again smooth. All mappings in this chapter are assumed to be smooth.
- When we compute, coordinates in parameter space will usually be written as (u_1, \dots, u_k) , and coordinates in \mathbf{R}^n as (x_1, \dots, x_n) .
- It may be disconcerting that a surface is by definition a mapping rather than a set, but this is for good reason. Just as the integration of the previous chapter was facilitated by distinguishing between functions and

their outputs, the integration of this chapter is facilitated by viewing the surfaces over which we integrate as mappings rather than their images.




Figure 8.1. A surface

When $k = 0$, Definition 8.1.1 is a little tricky. By convention, \mathbf{R}^0 is the set of all points with no coordinates, each of the no coordinates being a real number. (Our definition of \mathbf{R}^n at the beginning of chapter 2 danced around this issue by requiring that n be positive.) There is exactly one such point, the point $()$. That is, \mathbf{R}^0 consists of a single point, naturally called $\mathbf{0}$ even though it is *not* (0) . A 0-surface in \mathbf{R}^n is thus a mapping

$$\Phi_p : \mathbf{R}^0 \longrightarrow \mathbf{R}^n, \quad \Phi_p(\mathbf{0}) = p,$$

where p is some point in \mathbf{R}^n . In other words, Φ_p simply parameterizes the point p . At the other dimensional extreme, if $k = n$ then any compact connected subset D of \mathbf{R}^n naturally defines a corresponding n -surface in \mathbf{R}^n by trivially parameterizing itself,

$$\Delta : D \longrightarrow \mathbf{R}^n, \quad \Delta(u) = u \text{ for all } u \in D.$$

Thus Definition 8.1.1 of a surface as a mapping is silly in the particular cases of $k = 0$ and $k = n$, when it amounts to parameterizing points by using the empty point as a parameter domain, or parameterizing solids by taking them to be their own parameter domains and having the identity mapping map them to themselves. But for intermediate values of k , i.e., $0 < k < n$, we

are going to integrate over k -dimensional subsets of \mathbf{R}^n by traversing them, and parameterizing is the natural way to do so. Especially, a 1-surface is a parameterized curve, and a 2-surface is a parameterized surface in the usual sense of surface as in figure 8.1.

Let A be an open subset of \mathbf{R}^n , let $\Phi : D \rightarrow A$ be a k -surface in A , and let $f : A \rightarrow \mathbf{R}$ be a smooth function. As mentioned above, if $k < n$ then the integral of f over $\Phi(D)$ in the sense of chapter 6 is zero because $\Phi(D)$ is lower-dimensional than its ambient space \mathbf{R}^n . However, the integral of f over Φ can be defined more insightfully.

For each point u of the parameter domain D , the n -by- k derivative matrix $\Phi'(u)$ has as its columns vectors that are naturally viewed as tangent vectors to Φ at $\Phi(u)$, the j th column being tangent to the curve in Φ that arises from motion in the j th direction of the parameter domain. In symbols, the matrix is

$$\Phi'(u) = [v_1 \cdots v_k]_{n \times k},$$

where each column vector v_j is

$$v_j = D_j \Phi(u) = \begin{bmatrix} D_j \Phi_1(u) \\ \vdots \\ D_j \Phi_n(u) \end{bmatrix}_{n \times 1}.$$

The parallelepiped spanned by these vectors (see figure 8.2) has a naturally defined k -dimensional volume.




Figure 8.2. Tangent parallelepiped

Definition 8.1.2 (Volume of a Parallelepiped). Let v_1, \dots, v_k be vectors in \mathbf{R}^n . Let V be the n -by- k matrix with these vectors as its columns. Then the

k -**volume** of the parallelepiped spanned by the $\{v_j\}$ is

$$\text{vol}_k(\mathcal{P}(v_1, \dots, v_k)) = \sqrt{\det(V^t V)}. \quad (8.1)$$

In coordinates, this formula is

$$\text{vol}_k(\mathcal{P}(v_1, \dots, v_k)) = \sqrt{\det([v_i \cdot v_j]_{i,j=1,\dots,k})}, \quad (8.2)$$

where $v_i \cdot v_j$ is the inner product of v_i and v_j .

The matrix V in this definition is n -by- k and its transpose V^t is k -by- n , so neither of them need be square. But the product $V^t V$ is square, k -by- k , and this is the matrix whose determinant is being taken. Equation (8.2) follows immediately from (8.1) because

$$V^t V = \begin{bmatrix} v_1^t \\ \vdots \\ v_k^t \end{bmatrix} [v_1 \cdots v_k] = \begin{bmatrix} v_1 \cdot v_1 & \cdots & v_1 \cdot v_k \\ \vdots & \ddots & \vdots \\ v_k \cdot v_1 & \cdots & v_k \cdot v_k \end{bmatrix} = [v_i \cdot v_j]_{i,j=1,\dots,k}.$$

For example, if $k = 1$ and $\gamma : [a, b] \rightarrow \mathbf{R}^n$ is a 1-surface (i.e., a curve) in \mathbf{R}^n , then its derivative matrix at a point u of $[a, b]$ has one column,

$$\gamma'(u) = \begin{bmatrix} \gamma'_1(u) \\ \vdots \\ \gamma'_n(u) \end{bmatrix}.$$

Consequently, formula (8.2) is

$$\text{length}(\gamma'(u)) = \sqrt{\gamma'(u) \cdot \gamma'(u)}.$$

That is, Definition 8.1.2 for $k = 1$ specializes to the definition of $|\gamma'|$ as $\sqrt{\gamma' \cdot \gamma'}$ from section 2.2. At the other extreme, if $k = n$ then formula (8.1) is

$$\text{vol}_n(\mathcal{P}(v_1, \dots, v_n)) = |\det(v_1, \dots, v_n)|.$$

That is, Definition 8.1.2 for $k = n$ recovers the interpretation of $|\det|$ as volume from section 3.8. When $k = 2$, formula (8.2) is

$$\begin{aligned} \text{area}(\mathcal{P}(v_1, v_2)) &= \sqrt{|v_1|^2 |v_2|^2 - (v_1 \cdot v_2)^2} \\ &= \sqrt{|v_1|^2 |v_2|^2 (1 - \cos^2 \theta_{12})} \\ &= |v_1| |v_2| |\sin \theta_{12}|, \end{aligned}$$

giving the familiar formula for the area of a parallelogram. When $k = 2$ and also $n = 3$, we can study the formula further by working in coordinates. Consider two vectors $u = (x_u, y_u, z_u)$ and $v = (x_v, y_v, z_v)$. An elementary

calculation shows that the quantity under the square root in the previous display works out to

$$|u|^2|v|^2 - (u \cdot v)^2 = |u \times v|^2.$$

So when $k = 2$ and $n = 3$, Definition 8.1.2 subsumes the familiar formula

$$\text{area}(\mathcal{P}(v_1, v_2)) = |v_1 \times v_2|.$$

Here is an argument that (8.2) is the appropriate formula for the k -dimensional volume of the parallelepiped spanned by the vectors v_1, \dots, v_k in \mathbf{R}^n . (The fact that the vectors are tangent vectors to a k -surface is irrelevant to this discussion.) Results from linear algebra guarantee that there exist vectors v_{k+1}, \dots, v_n in \mathbf{R}^n such that

- each of v_{k+1} through v_n is a unit vector orthogonal to all the other v_j ,
- $\det(v_1, \dots, v_n) \geq 0$.

Recall the notation in Definition 8.1.2 that V is the n -by- k matrix with columns v_1, \dots, v_k . Augment V to an n -by- n matrix W by adding the remaining v_j as columns too,

$$W = [v_1 \cdots v_n] = [V \ v_{k+1} \cdots v_n].$$

The scalar $\det(W)$ is the n -dimensional volume of the parallelepiped spanned by v_1, \dots, v_n . But by the properties of v_{k+1} through v_n , this scalar should also be the k -dimensional volume of the the parallelepiped spanned by v_1, \dots, v_k . That is, the natural definition is (using the second property of v_1, \dots, v_n for the second equality to follow)

$$\begin{aligned} \text{vol}_k(\mathcal{P}(v_1, \dots, v_k)) &= \det(W) = \sqrt{(\det W)^2} = \sqrt{\det(W^t) \det(W)} \\ &= \sqrt{\det(W^t W)}. \end{aligned}$$

The first property shows that

$$W^t W = \begin{bmatrix} V^t V & \mathbf{0}_{k \times (n-k)} \\ \mathbf{0}_{(n-k) \times k} & I_{n-k} \end{bmatrix},$$

so that $\det(W^t W) = \det(V^t V)$ and the natural definition becomes the desired formula,

$$\text{vol}_k(\mathcal{P}(v_1, \dots, v_k)) = \sqrt{\det(V^t V)}.$$

The argument here generalizes the ideas used in section 3.10 to suggest a formula for the area of a 2-dimensional parallelogram in \mathbf{R}^3 as a 3-by-3 determinant. Thus the coordinate calculation sketched in the previous paragraph to recover the relation between parallelogram area and cross product length in \mathbf{R}^3 was unnecessary.

With k -dimensional volume in hand, we can naturally define the integral of a function over a k -surface.

Definition 8.1.3 (Integral of a Function over a Surface). Let A be an open subset of \mathbf{R}^n . Let $\Phi: D \rightarrow A$ be a k -surface in A . Let $f: \Phi(D) \rightarrow \mathbf{R}$ be a function such that $f \circ \Phi$ is smooth. Then the **integral of f over Φ** is

$$\int_{\Phi} f = \int_D (f \circ \Phi) \operatorname{vol}_k(\mathcal{P}(D_1\Phi, \dots, D_k\Phi)).$$

In particular, the k -dimensional volume of Φ is

$$\operatorname{vol}_k(\Phi) = \int_{\Phi} 1 = \int_D \operatorname{vol}_k(\mathcal{P}(D_1\Phi, \dots, D_k\Phi)).$$

By Definition 8.1.2 the k -volume factor in the surface integral is

$$\operatorname{vol}_k(\mathcal{P}(D_1\Phi, \dots, D_k\Phi)) = \sqrt{\det(\Phi'^t \Phi')} = \sqrt{\det([D_i\Phi \cdot D_j\Phi]_{i,j=1,\dots,k})}.$$

The idea of Definition 8.1.3 is that as a parameter u traverses the parameter domain D , the composition $f \circ \Phi$ samples the function f over the surface, while the k -volume factor makes the integral the limit of sums of many f -weighted small tangent parallelepiped k -volumes over the surface rather than the limit of sums of many $(f \circ \Phi)$ -weighted small box volumes over the parameter domain. (See figure 8.3.) The k -volume factor itself is not small, as seen in figure 8.2, but it is the ratio of the small parallelepiped k -volume to the small box volume shown in figure 8.3.




Figure 8.3. Integrating over a surface

For example, let r be a positive real number and consider a 2-surface in \mathbf{R}^3 ,

$$\Phi : [0, 2\pi] \times [0, \pi] \longrightarrow \mathbf{R}^3, \quad \Phi(\theta, \varphi) = (r \cos \theta \sin \varphi, r \sin \theta \sin \varphi, r \cos \varphi).$$

This is the 2-sphere of radius r . Since the sphere is a surface of revolution, its area is readily computed by methods from a first calculus course, but we do so with the ideas of this section to demonstrate their use. The derivative vectors are

$$v_1 = \begin{bmatrix} -r \sin \theta \sin \varphi \\ r \cos \theta \sin \varphi \\ 0 \end{bmatrix}, \quad v_2 = \begin{bmatrix} r \cos \theta \cos \varphi \\ r \sin \theta \cos \varphi \\ -r \sin \varphi \end{bmatrix},$$


and so the integrand of the surface area integral is

$$\sqrt{|v_1|^2 |v_2|^2 - (v_1 \cdot v_2)^2} = \sqrt{r^4 \sin^2 \varphi} = r^2 \sin \varphi$$

(note that $\sin \varphi \geq 0$ since $\varphi \in [0, \pi]$). Therefore the area is

$$\text{area}(\Phi) = r^2 \int_{\theta=0}^{2\pi} \int_{\varphi=0}^{\pi} \sin \varphi = 4\pi r^2.$$

The integral in Definition 8.1.3 seems to depend on the surface Φ as a parameterization rather than merely as a set, but in fact the integral is unaffected by reasonable changes of parameterization, because of the Change of Variable Theorem. To see this, let A be an open subset of \mathbf{R}^n , and let $\Phi : D \rightarrow A$ and $\Psi : \tilde{D} \rightarrow A$ be k -surfaces in A . Suppose that there exists a smoothly invertible mapping $T : D \rightarrow \tilde{D}$ such that $\Psi \circ T = \Phi$. In other words, T is smooth, T is invertible, its inverse is also smooth, and the following diagram commutes (meaning that either path around the triangle yields the same result):



When such a mapping T exists, Ψ is called a **reparameterization** of Φ . Let $f : A \rightarrow \mathbf{R}$ be any smooth function. Then the integral of f over the reparameterization Ψ of Φ is

$$\int_{\tilde{D}} (f \circ \Psi) \sqrt{\det(\Psi'^t \Psi')}.$$

By the Change of Variable Theorem, since $\tilde{D} = T(D)$, this integral is

$$\int_D (f \circ \Psi \circ T) \sqrt{\det((\Psi' \circ T)^t (\Psi' \circ T))} |\det(T')|.$$

But $|\det(T')| = \sqrt{\det(T')^2} = \sqrt{\det(T'^t) \det(T')}$, so this becomes

$$\int_D (f \circ \Psi \circ T) \sqrt{\det(T'^t (\Psi' \circ T)^t (\Psi' \circ T) T')},$$

and by the general matrix rule $B^t A^t A B = (AB)^t A B$, this is in turn

$$\int_D (f \circ \Psi \circ T) \sqrt{\det((\Psi' \circ T) T')^t (\Psi' \circ T) T')}.$$

Finally, since $\Psi \circ T = \Phi$, the Chain Rule shows that we have

$$\int_D (f \circ \Phi) \sqrt{\det(\Phi'^t \Phi')},$$

giving the integral of f over the original surface Φ as desired.

Exercises

8.1.1. Consider two vectors $u = (x_u, y_u, z_u)$ and $v = (x_v, y_v, z_v)$. Calculate that $|u|^2|v|^2 - (u \cdot v)^2 = |u \times v|^2$.

8.1.2. Let $f(x, y, z) = x^2 + yz$.

(a) Integrate f over the box $B = [0, 1]^3$.

(b) Integrate f over the parameterized curve

$$\gamma : [0, 2\pi] \rightarrow \mathbf{R}^3, \quad \gamma(t) = (\cos t, \sin t, t).$$

(c) Integrate f over the parameterized surface

$$S : [0, 1]^2 \rightarrow \mathbf{R}^3, \quad S(u, v) = (u + v, u - v, v).$$

(d) Integrate f over the parameterized solid

$$V : [0, 1]^3 \rightarrow \mathbf{R}^3, \quad V(u, v, w) = (u + v, v - w, u + w).$$

8.1.3. (a) Let $D \subset \mathbf{R}^k$ be a parameter domain, and let $f : D \rightarrow \mathbf{R}$ be a smooth function. Recall from exercise 2.4.3 that the graph of f is a subset of \mathbf{R}^{k+1} ,

$$G(f) = \{(u, f(u)) : u \in D\}.$$

Note that f is a 1-surface in \mathbf{R} , while the surface that captures the idea of the graph of f as a k -surface in \mathbf{R}^{k+1} is not f itself but rather

$$\Phi : D \rightarrow \mathbf{R}^{k+1}, \quad \Phi(u) = (u, f(u)).$$

Derive a formula for the k -dimensional volume of Φ . In particular, show that when $k = 2$ the formula is

$$\text{area}(\Phi) = \int_D \sqrt{1 + (D_1 f)^2 + (D_2 f)^2}.$$

(b) What is the area of the graph of the function $f : D \rightarrow \mathbf{R}$ (where D is the unit disk in the plane) given by $f(x, y) = x^2 + y^2$?

8.2 Flow and Flux Integrals

Let A be an open subset of \mathbf{R}^n . A mapping $F : A \rightarrow \mathbf{R}^n$ is also called a *vector field on A* . (The usage of “field” here is unrelated to the field axioms.) If $\gamma : I \rightarrow A$ is a curve in A and u is a point of I , then the *flow of F along γ at u* is the (scalar) component of $F(\gamma(u))$ tangent to γ at $\gamma(u)$. If $\Phi : D \rightarrow A$ is an $(n - 1)$ -surface in A and u is a point of D , then the *flux of F through Φ at u* is the (scalar) component of F normal to Φ at $\Phi(u)$. Surface integrals involving the flow or the flux of a vector field arise naturally. If F is viewed as a force field then its flow integrals (also called *line integrals*) measure the work of moving along curves γ in A . If F is viewed as a velocity field describing the motion of some fluid then its flux integrals measure the rate at which fluid passes through permeable membranes Φ in A . Each of the classical theorems of vector integral calculus to be proved at the end of the chapter involves a flow integral or a flux integral.

Flow and flux integrals have a more convenient form than the general integral of a function over a surface, in that the k -volume factor from Definition 8.1.3 (an unpleasant square root) cancels, and what remains is naturally expressed in terms of determinants of the derivatives of the component functions of Φ . These formulas rapidly become complicated, so the point of this section is only to see what form they take.

Working first in two dimensions, consider a vector field,

$$F = (F_1, F_2) : \mathbf{R}^2 \rightarrow \mathbf{R}^2,$$

and a curve,

$$\gamma = (\gamma_1, \gamma_2) : [a, b] \rightarrow \mathbf{R}^2.$$

Assuming that the derivative γ' is always nonzero, the unit tangent vector to γ at the point $\gamma(u)$, pointing in the direction of the traversal, is

$$\widehat{T}(\gamma(u)) = \frac{\gamma'(u)}{|\gamma'(u)|}.$$

Note that the denominator is the length factor in Definition 8.1.3. The parallel component of $F(\gamma(u))$ along $\widehat{T}(\gamma(u))$ has magnitude $(F \cdot \widehat{T})(\gamma(u))$. (See exercise 2.2.15.) Therefore the net flow of F along γ in the direction of traversal is $\int_{\gamma} F \cdot \widehat{T}$. By Definition 8.1.3 this **flow integral** is

$$\int_{\gamma} F \cdot \widehat{T} = \int_{u=a}^b F(\gamma(u)) \cdot \frac{\gamma'(u)}{|\gamma'(u)|} |\gamma'(u)| = \int_{u=a}^b F(\gamma(u)) \cdot \gamma'(u), \quad (8.3)$$

and the length factor has canceled. In coordinates, the flow integral is

$$\int_{\gamma} F \cdot \widehat{T} = \int_{u=a}^b ((F_1 \circ \gamma)\gamma'_1 + (F_2 \circ \gamma)\gamma'_2)(u). \quad (8.4)$$

On the other hand, for any vector $(x, y) \in \mathbf{R}^2$, define $(x, y)^\times = (-y, x)$. (This seemingly ad hoc procedure of negating one of the vector entries and then exchanging them will be revisited soon as a particular manifestation of a general idea.) Then the unit normal vector to the curve γ at the point $\gamma(u)$, at angle $\pi/2$ counterclockwise from $\widehat{T}(\gamma(u))$, is

$$\widehat{N}(\gamma(u)) = \frac{\gamma'(u)^\times}{|\gamma'(u)|}.$$

Therefore the net flux of F through γ counterclockwise to the direction of traversal is the **flux integral**

$$\int_{\gamma} F \cdot \widehat{N} = \int_{u=a}^b F(\gamma(u)) \cdot \gamma'(u)^\times, \quad (8.5)$$

or, in coordinates,

$$\int_{\gamma} F \cdot \widehat{N} = \int_{u=a}^b ((F_2 \circ \gamma)\gamma'_1 - (F_1 \circ \gamma)\gamma'_2)(u). \quad (8.6)$$

Next let $n = 3$ and modify the vector field F suitably to

$$F = (F_1, F_2, F_3) : \mathbf{R}^3 \longrightarrow \mathbf{R}^3.$$

The intrinsic expression (8.3) for the flow integral of F along a curve γ remains unchanged in \mathbf{R}^3 , making the 3-dimensional counterpart of (8.4) in coordinates obvious,

$$\int_{\gamma} F \cdot \widehat{T} = \int_{u=a}^b ((F_1 \circ \gamma)\gamma'_1 + (F_2 \circ \gamma)\gamma'_2 + (F_3 \circ \gamma)\gamma'_3)(u).$$

As for the flux integral, consider a 2-surface in \mathbf{R}^3 ,

$$\Phi = (\Phi_1, \Phi_2, \Phi_3) : D \longrightarrow \mathbf{R}^3.$$

Assuming that the two columns $D_1\Phi$ and $D_2\Phi$ of the derivative matrix Φ' are always linearly independent, a unit normal to the surface Φ at the point $\Phi(u)$ (where now $u = (u_1, u_2)$), is obtained from their cross product,

$$\widehat{N}(\Phi(u)) = \frac{D_1\Phi(u) \times D_2\Phi(u)}{|D_1\Phi(u) \times D_2\Phi(u)|}.$$

By property CP6 of the cross product, the denominator in this expression is the area of the parallelogram spanned by $D_1\Phi(u)$ and $D_2\Phi(u)$, and this is the area factor in Definition 8.1.3 of the surface integral. Therefore this factor cancels in the flux integral of F through Φ in the \widehat{N} -direction,

$$\int_{\Phi} F \cdot \widehat{N} = \int_{u \in D} F(\Phi(u)) \cdot (D_1\Phi(u) \times D_2\Phi(u)), \quad (8.7)$$

or, in coordinates,

$$\int_{\Phi} F \cdot \widehat{N} = \int_{u \in D} \begin{pmatrix} (F_1 \circ \Phi)(D_1\Phi_2 D_2\Phi_3 - D_1\Phi_3 D_2\Phi_2) \\ + (F_2 \circ \Phi)(D_1\Phi_3 D_2\Phi_1 - D_1\Phi_1 D_2\Phi_3) \\ + (F_3 \circ \Phi)(D_1\Phi_1 D_2\Phi_2 - D_1\Phi_2 D_2\Phi_1) \end{pmatrix} (u). \quad (8.8)$$

Whereas the 2-dimensional flow and flux integrands and the 3-dimensional flow integrand involved derivatives γ'_j of the 1-surface γ , the integrand here contains the determinants of all 2-by-2 subblocks of the 3-by-2 derivative matrix of the 2-surface Φ ,

$$\Phi' = \begin{bmatrix} D_1\Phi_1 & D_2\Phi_1 \\ D_1\Phi_2 & D_2\Phi_2 \\ D_1\Phi_3 & D_2\Phi_3 \end{bmatrix}.$$

This gives a hint about the general picture. Nonetheless, (8.8) is foreboding enough that we should pause and think before trying to compute more formulas.

For general n , formula (8.3) for the flow integral of a vector field along a curve generalizes transparently,

$$\int_{\gamma} F \cdot \widehat{T} = \int_{u=a}^b ((F \circ \gamma) \cdot \gamma')(u) = \int_{u=a}^b \left(\sum_{i=1}^n (F_i \circ \gamma) \gamma'_i \right) (u). \quad (8.9)$$

But the generalization of formulas (8.5) through (8.8) to a formula for the flux integral of a vector field in \mathbf{R}^n through an $(n-1)$ -surface is not so obvious. Based on (8.7) the intrinsic formula should be

$$\int_{\Phi} F \cdot \widehat{N} = \int_{u \in D} ((F \circ \Phi) \cdot (D_1\Phi \times \cdots \times D_{n-1}\Phi))(u), \quad (8.10)$$

where the $(n-1)$ -fold cross product on \mathbf{R}^n is analogous to the 2-fold cross product on \mathbf{R}^3 from section 3.10. That is, the cross product should be orthogonal to each of the multiplicand-vectors, its length should be their $(n-1)$ -dimensional volume, and when the multiplicands are linearly independent, they should combine with their cross product to form a positive basis of \mathbf{R}^n .

Such a cross product exists by methods virtually identical to section 3.10. What is special to three dimensions is that the cross product is binary, i.e., it is a twofold product. In coordinates, a mnemonic formula for the cross product in \mathbf{R}^3 (viewing the vectors as columns) is

$$v_1 \times v_2 = \det \left[\begin{array}{c|c|c} v_1 & v_2 & \begin{bmatrix} e_1 \\ e_2 \\ e_3 \end{bmatrix} \end{array} \right].$$

This formula appeared in row form in section 3.10, and it makes the corresponding formula for the cross product of $n-1$ vectors in \mathbf{R}^n inevitable,

$$v_1 \times \cdots \times v_{n-1} = \det \left[\begin{array}{c|c|c|c} v_1 & \cdots & v_{n-1} & \begin{array}{c} e_1 \\ \vdots \\ e_n \end{array} \end{array} \right]. \quad (8.11)$$

For example, in \mathbf{R}^2 a single vector $v = (x, y)$ has a sort of cross product,

$$v^\times = \det \left[\begin{array}{c|c} x & e_1 \\ y & e_2 \end{array} \right] = (-y, x)$$

This is the formula that appeared with no explanation as part of the flux integral in \mathbf{R}^2 . That is, the generalization (8.10) of the 3-dimensional flux integral to higher dimensions also subsumes the 2-dimensional case. Returning to \mathbf{R}^n , the cross product of the vectors $D_1\Phi(u), \dots, D_{n-1}\Phi(u)$ is

$$(D_1\Phi \times \cdots \times D_{n-1}\Phi)(u) = \det \left[\begin{array}{c|c|c|c} D_1\Phi(u) & \cdots & D_{n-1}\Phi(u) & \begin{array}{c} e_1 \\ \vdots \\ e_n \end{array} \end{array} \right].$$

This can be understood better by considering the data in the matrix as rows. Recall that for $i = 1, \dots, n$, the i th row of the n -by- $(n-1)$ derivative matrix Φ' is the derivative matrix of the i th component function of Φ ,

$$\Phi'_i(u) = [D_1\Phi_i(u) \cdots D_{n-1}\Phi_i(u)].$$

In terms of these component function derivatives, the general cross product is

$$\begin{aligned} (D_1\Phi \times \cdots \times D_{n-1}\Phi)(u) &= \det \left[\begin{array}{c|c} \Phi'_1(u) & e_1 \\ \vdots & \vdots \\ \Phi'_n(u) & e_n \end{array} \right] = (-1)^{n-1} \det \left[\begin{array}{c|c} e_1 & \Phi'_1(u) \\ \vdots & \vdots \\ e_n & \Phi'_n(u) \end{array} \right] \\ &= (-1)^{n-1} \left(\det \begin{bmatrix} \Phi'_2(u) \\ \Phi'_3(u) \\ \Phi'_4(u) \\ \vdots \\ \Phi'_n(u) \end{bmatrix} e_1 - \det \begin{bmatrix} \Phi'_1(u) \\ \Phi'_3(u) \\ \Phi'_4(u) \\ \vdots \\ \Phi'_n(u) \end{bmatrix} e_2 + \det \begin{bmatrix} \Phi'_1(u) \\ \Phi'_2(u) \\ \Phi'_4(u) \\ \vdots \\ \Phi'_n(u) \end{bmatrix} e_3 + \cdots \right) \\ &= (-1)^{n-1} \sum_{i=1}^n (-1)^{i-1} \det \begin{bmatrix} \Phi'_1(u) \\ \vdots \\ \Phi'_{i-1}(u) \\ \Phi'_{i+1}(u) \\ \vdots \\ \Phi'_n(u) \end{bmatrix} e_i. \end{aligned}$$

Thus finally, the general flux integral in coordinates is

$$\int_{\Phi} F \cdot \widehat{N} = (-1)^{n-1} \int_{u \in D} \left(\sum_{i=1}^n (-1)^{i-1} (F_i \circ \Phi) \det \begin{bmatrix} \Phi'_1 \\ \vdots \\ \Phi'_{i-1} \\ \Phi'_{i+1} \\ \vdots \\ \Phi'_n \end{bmatrix} \right) (u). \quad (8.12)$$

The integrand here contains the determinants of all $(n-1)$ -by- $(n-1)$ subblocks of the n -by- $(n-1)$ derivative matrix of the $(n-1)$ -surface Φ . The best way to understand the notation of (8.12) is to derive (8.6) and (8.8) from it by setting $n = 2$ and then $n = 3$.

We end the section by mentioning one more integral. Let $k = 2$ and let $n = 4$, and consider a 2-surface in \mathbf{R}^4 ,

$$\Phi = (\Phi_1, \Phi_2, \Phi_3, \Phi_4) : D \rightarrow \mathbf{R}^4.$$

Note that Φ' is a 4-by-2 matrix,

$$\Phi' = \begin{bmatrix} \Phi'_1 \\ \Phi'_2 \\ \Phi'_3 \\ \Phi'_4 \end{bmatrix} = \begin{bmatrix} D_1\Phi_1 & D_2\Phi_1 \\ D_1\Phi_2 & D_2\Phi_2 \\ D_1\Phi_3 & D_2\Phi_3 \\ D_1\Phi_4 & D_2\Phi_4 \end{bmatrix},$$

so that any two of its rows form a square matrix. Consider also any six smooth functions

$$F_{1,2}, F_{1,3}, F_{1,4}, F_{2,3}, F_{2,4}, F_{3,4} : \mathbf{R}^4 \rightarrow \mathbf{R}.$$

Then we can define an integral,

$$\int_{u \in D} \left(\begin{aligned} &(F_{1,2} \circ \Phi) \det \begin{bmatrix} \Phi'_1 \\ \Phi'_2 \end{bmatrix} + (F_{1,3} \circ \Phi) \det \begin{bmatrix} \Phi'_1 \\ \Phi'_3 \end{bmatrix} + (F_{1,4} \circ \Phi) \det \begin{bmatrix} \Phi'_1 \\ \Phi'_4 \end{bmatrix} \\ &+ (F_{2,3} \circ \Phi) \det \begin{bmatrix} \Phi'_2 \\ \Phi'_3 \end{bmatrix} + (F_{2,4} \circ \Phi) \det \begin{bmatrix} \Phi'_2 \\ \Phi'_4 \end{bmatrix} + (F_{3,4} \circ \Phi) \det \begin{bmatrix} \Phi'_3 \\ \Phi'_4 \end{bmatrix} \end{aligned} \right) (u). \quad (8.13)$$

Since the surface Φ is not 1-dimensional, this is not a flow integral. And since Φ is not $(n-1)$ -dimensional, it is not a flux integral either. Nonetheless, since the integrand contains the determinants of all 2-by-2 subblocks of the 4-by-2 derivative matrix of the 2-surface Φ , it is clearly cut from the same cloth as the flow and flux integrands of this section. The ideas of this chapter will encompass this integral and many others in the same vein.

As promised at the beginning of the section, the k -volume factor has canceled in flow and flux integrals, and the remaining integrand features determinants of the derivatives of the component functions of the surface of integration. Rather than analyze such cluttered integrals, the method of this chapter is to abstract their key properties into symbol patterns, and then work with

the patterns algebraically instead. An analysis tracking all the details of the original setup would be excruciating to follow, not to mention being unimaginable to recreate ourselves. Instead, we will work insightfully, economy of ideas leading to ease of execution. Since the definitions to follow do indeed distill the essence of vector integration, they will enable us to think fluently about the phenomena that we encounter. This is real progress in methodology, much less laborious than the classical approach. Indeed, having seen the modern argument, it is unimaginable to *want* to recreate the older one.

Exercises

8.2.1. Show that the n -dimensional cross product defined by a formula in (8.11) satisfies the property

$$\langle v_1 \times \cdots \times v_{n-1}, w \rangle = \det(v_1, \dots, v_{n-1}, w) \quad \text{for all } w \in \mathbf{R}^n.$$

As in section 3.10, this property characterizes the cross product uniquely. Are there significant differences between deriving the properties of the cross product from its characterization (cf. Proposition 3.10.2) in n dimensions rather than in 3?

8.2.2. Derive equations (8.6) and (8.8) from equation (8.12).

8.3 Differential Forms Syntactically and Operationally

We need objects to integrate over surfaces, objects whose integrals encompass at least the general flow integral (8.9) and flux integral (8.12) of the previous section. Let A be an open subset of \mathbf{R}^n . The objects are called **differential forms of order k on A** or simply **k -forms on A** . Thus a k -form ω is some sort of mapping

$$\omega : \{k\text{-surfaces in } A\} \longrightarrow \mathbf{R}.$$

Naturally the value $\omega(\Phi)$ will be denoted $\int_{\Phi} \omega$. The definition of a k -form will come in two parts. The first is syntactic: it doesn't say what a k -form *is* as a function of k -surfaces, only what kind of *name* a k -form can have. This definition requires some preliminary vocabulary: An **ordered k -tuple from $\{1, \dots, n\}$** is a vector

$$(i_1, \dots, i_k) \text{ with each } i_j \in \{1, \dots, n\}.$$

For example, the ordered 3-tuples from $\{1, 2\}$ are

$$(1, 1, 1), (1, 1, 2), (1, 2, 1), (1, 2, 2), (2, 1, 1), (2, 1, 2), (2, 2, 1), (2, 2, 2).$$

A sum over the ordered k -tuples from $\{1, \dots, n\}$ means simply a sum of terms with each term corresponding to a distinct k -tuple. Thus we may think of an

ordered k -tuple (i_1, \dots, i_k) as a sort of multiple index or multiple subscript, and for this reason we often will abbreviate it to I . These multiple subscripts will figure prominently throughout the chapter, so you should get comfortable with them. Exercise 8.3.1 provides some practice.

Definition 8.3.1 (Syntax of Differential Forms). *Let A be an open subset of \mathbf{R}^n . A 0-form on A is a smooth function $f : A \rightarrow \mathbf{R}$. For $k \geq 1$, a k -form on A is an element of the form*

$$\sum_{i_1, \dots, i_k=1}^n f_{(i_1, \dots, i_k)} dx_{i_1} \wedge \cdots \wedge dx_{i_k},$$

or

$$\sum_I f_I dx_I,$$

where each $I = (i_1, \dots, i_k)$ is an ordered k -tuple from $\{1, \dots, n\}$ and each f_I is a smooth function $f_I : A \rightarrow \mathbf{R}$.

Make the convention that the empty set $I = \emptyset$ is the only ordered 0-tuple from $\{1, \dots, n\}$, and that the corresponding empty product dx_\emptyset is 1. Then the definition of a k -form for $k \geq 1$ in Definition 8.3.1 also makes sense for $k = 0$, and it subsumes the special definition that was given for $k = 0$.

For example, a differential form for $n = 3$ and $k = 1$ is

$$e^{x+y+z} dx + \sin(yz) dy + x^2 z dz,$$

and a differential form for $n = 2$ and $k = 2$ is

$$y dx \wedge dx + e^x dx \wedge dy + y \cos x dy \wedge dx.$$

The expression

$$\frac{1}{x} dx$$

is a 1-form on the open subset $A = \{x \in \mathbf{R} : x \neq 0\}$ of \mathbf{R} , but it is not a 1-form on all of \mathbf{R} . The hybrid expression

$$z dx \wedge dy + e^x dz$$

is not a differential form because it mixes an order 2 term and an order 1 term.

Before completing the definition of differential form, we need one more piece of terminology. If M is an n -by- k matrix and $I = (i_1, \dots, i_k)$ is an ordered k -tuple from $\{1, \dots, n\}$, then M_I denotes the square k -by- k matrix comprised of the I th rows of M . For example, if

$$M = \begin{bmatrix} 1 & 2 \\ 3 & 4 \\ 5 & 6 \end{bmatrix},$$

and if $I = (3, 1)$, then

$$M_I = \begin{bmatrix} 5 & 6 \\ 1 & 2 \end{bmatrix}.$$

The second part of the definition of a k -form explains how to integrate it over a k -surface. In this definition, a differential form in the sense of Definition 8.3.1 is called a syntactic differential form.

Definition 8.3.2 (Integration of Differential Forms). *Let A be an open subset of \mathbf{R}^n . For $k = 0$, a syntactic 0-form $\omega = f$ on A gives rise to a function of 0-surfaces in A , also called ω ,*

$$\omega : \{0\text{-surfaces in } A\} \longrightarrow \mathbf{R},$$

defined by the rule that for any point $p \in A$,

$$\omega(\Phi_p) = f(p).$$

That is, integrating ω over a one-point surface consists simply of evaluating f at the point. For $k \geq 1$, a syntactic k -form $\omega = \sum_I f_I dx_I$ on A gives rise to a function of k -surfaces in A , also called ω ,

$$\omega : \{k\text{-surfaces in } A\} \longrightarrow \mathbf{R},$$

defined by the rule that for any k -surface $\Phi : D \rightarrow A$,

$$\omega(\Phi) = \int_D \sum_I (f_I \circ \Phi) \det \Phi'_I. \quad (8.14)$$

For all k , the integral of ω over Φ is defined to be $\omega(\Phi)$,

$$\int_{\Phi} \omega = \omega(\Phi).$$

Formula (8.14), defining $\omega(\Phi)$, is the key for everything to follow in this chapter. It defines an integral over the image $\Phi(D)$, which may have volume zero in \mathbf{R}^n , by “pulling back”—this term will later be defined precisely—to an integral over the parameter domain D , which is a full-dimensional set in \mathbf{R}^k and hence has positive k -dimensional volume.

Under Definition 8.3.2, the integral of a differential form over a surface depends on the surface as a mapping, i.e., as a parameterization. However, it is a straightforward exercise to show that the Multivariable Change of Variable Theorem implies that the integral is unaffected by reasonable changes of parameterization.

Returning to formula (8.14): despite looking like the flux integral (8.12), it may initially be impenetrable to the reader who (like the author) does not assimilate notation quickly. The next two sections will illustrate the formula

in specific instances, after which its general workings should be clear. Before long, you will have an operational understanding of the definition.

Operational understanding should be complemented by structural understanding. The fact that the formal consequences of Definitions 8.3.1 and 8.3.2 subsume the main results of classical integral vector calculus still doesn't explain these *ad hoc* definitions conceptually. For everything to play out so nicely, the definitions must somehow be natural rather than merely clever, and a structural sense of why they work so well might let us extend the ideas to other contexts rather than simply tracking them. Indeed, differential forms fit into a mathematical structure called a *cotangent bundle*, with each differential form being a *section* of the bundle. The construction of the cotangent bundle involves the *dual space* of the *alternation* of a *tensor product*, all of these formidable-sounding technologies being utterly Platonic mathematical objects. However, understanding this language requires an investment in ideas and abstraction, and in the author's judgment the startup cost is much higher without some experience first. Hence the focus of the chapter is purely operational. Since formula (8.14) may be opaque to the reader for now, the first order of business is to render it transparent by working easy concrete examples.

Exercises

8.3.1. Write out all ordered k -tuples from $\{1, \dots, n\}$ in the cases $n = 4, k = 1$; $n = 3, k = 2$. In general, how many ordered k -tuples $I = (i_1, \dots, i_k)$ from $\{1, \dots, n\}$ are there? How many of these are **increasing**, meaning that $i_1 < \dots < i_k$? Write out all increasing k -tuples from $\{1, 2, 3, 4\}$ for $k = 1, 2, 3, 4$.

8.3.2. An expression $\omega = \sum_I f_I dx_I$ where the sum is over only increasing k -tuples from $\{1, \dots, n\}$ is called a **standard presentation** of ω . Write out explicitly what a standard presentation for a k -form on \mathbf{R}^4 looks like for $k = 0, 1, 2, 3, 4$.

8.4 Examples: 1-forms

A k -form is a function of k -surfaces. That is, one can think of a k -form ω as a set of instructions: given a k -surface Φ , ω carries out some procedure on Φ to produce a real number, $\int_{\Phi} \omega$.

For example, let

$$\omega = x dy \quad \text{and} \quad \lambda = y dz,$$

both 1-forms on \mathbf{R}^3 . A 1-surface in \mathbf{R}^3 is a curve,

$$\gamma = (\gamma_1, \gamma_2, \gamma_3) : [a, b] \longrightarrow \mathbf{R}^3,$$

with 3-by-1 derivative matrix

$$\gamma' = \begin{bmatrix} \gamma'_1 \\ \gamma'_2 \\ \gamma'_3 \end{bmatrix}.$$

For any such curve, ω is the instructions, “integrate $\gamma_1\gamma'_2$ over the parameter domain $[a, b]$,” and similarly λ instructs to integrate $\gamma_2\gamma'_3$. You should work through applying formula (8.14) to ω and λ to see how it produces these directions. Note that x and y are being treated as functions on \mathbf{R}^3 —for example,

$$x(a, b, c) = a \quad \text{for all } (a, b, c),$$

so that $x \circ \gamma = \gamma_1$.

To see ω and λ work on a specific curve, consider the helix

$$H : [0, 2\pi] \rightarrow \mathbf{R}^3, \quad H(t) = (a \cos t, a \sin t, bt).$$

Its derivative matrix is

$$H'(t) = \begin{bmatrix} -a \sin t \\ a \cos t \\ b \end{bmatrix} \quad \text{for all } t \in [0, 2\pi].$$

Thus by (8.14),

$$\int_H \omega = \int_{t=0}^{2\pi} a \cos t \cdot a \cos t = \pi a^2 \quad \text{and} \quad \int_H \lambda = \int_{t=0}^{2\pi} a \sin t \cdot b = 0.$$

Looking at the projections of the helix in the (x, y) -plane and the (y, z) -plane suggests that these are the right values for $\int_H x dy$ and $\int_H y dz$ if we interpret the symbols $x dy$ and $y dz$ as in one-variable calculus. (See figure 8.4.)

For another example, let

$$\omega = dx,$$

a 1-form on \mathbf{R}^3 , and consider any curve

$$\gamma : [a, b] \rightarrow \mathbf{R}^3, \quad \gamma(t) = (\gamma_1(t), \gamma_2(t), \gamma_3(t)).$$

Then

$$\int_\gamma \omega = \int_a^b (1 \circ \gamma) \cdot \gamma'_1 = \int_a^b \gamma'_1 = \gamma_1(b) - \gamma_1(a).$$

A change of notation makes this example more telling. Rewrite the component functions of the curve as x , y , and z ,

$$\gamma : [a, b] \rightarrow \mathbf{R}^3, \quad \gamma(t) = (x(t), y(t), z(t)).$$

So now x is not a function on \mathbf{R}^3 as in the previous example, but a function on $[a, b]$. The integral rewrites as

$$\int_\gamma \omega = \int_a^b x' = x(b) - x(a).$$




Figure 8.4. Integrating 1-forms over a helix

This shows that *the form dx does indeed measure change in x along curves*. As a set of instructions it simply says to evaluate the x -coordinate difference from the initial point on the curve to the final point. Think of dx as a Pac-Man. You give it a curve, it runs along the curve and gives you back your score: the net change in x along the curve. Returning to the helix H , it is now clear with no further work that

$$\int_H dx = 0, \quad \int_H dy = 0, \quad \int_H dz = 2\pi b.$$

It is a good practice with formula (8.14) to confirm these values.

To generalize the previous example, let A be an open subset of \mathbf{R}^n , let $f : A \rightarrow \mathbf{R}$ be any smooth function, and associate a 1-form ω to f ,

$$\omega = D_1 f dx_1 + \cdots + D_n f dx_n.$$

Then for any curve $\gamma : [a, b] \rightarrow A$,

$$\begin{aligned} \int_\gamma \omega &= \int_a^b (D_1 f \circ \gamma) \gamma'_1 + \cdots + (D_n f \circ \gamma) \gamma'_n \\ &= \int_a^b (f \circ \gamma)' \quad \text{by the chain rule in coordinates} \\ &= (f \circ \gamma) \Big|_a^b \\ &= f(\gamma(b)) - f(\gamma(a)). \end{aligned}$$

That is, *the form ω measures change in f along curves*. This makes it tempting to give ω the name df , i.e., to define

$$df = D_1 f dx_1 + \cdots + D_n f dx_n.$$

Soon we will do so as part of a more general definition.

(Recall the chain rule: If $A \subset \mathbf{R}^n$ is open, then for any smooth $\gamma : [a, b] \rightarrow A$ and $f : A \rightarrow \mathbf{R}$,

$$\begin{aligned} (f \circ \gamma)'(t) &= f'(\gamma(t))\gamma'(t) \\ &= [D_1 f(\gamma(t)) \ \cdots \ D_n f(\gamma(t))] \begin{bmatrix} \gamma'_1(t) \\ \vdots \\ \gamma'_n(t) \end{bmatrix} \\ &= \sum_{i=1}^n D_i f(\gamma(t))\gamma'_i(t) \\ &= \left[\sum_{i=1}^n (D_i f \circ \gamma)\gamma'_i \right] (t), \end{aligned}$$

so indeed $(f \circ \gamma)' = \sum_{i=1}^n (D_i f \circ \gamma)\gamma'_i$.

Continuing to generalize, consider now a 1-form that does not necessarily arise from differentiation,

$$\omega = F_1 dx_1 + \cdots + F_n dx_n.$$

For any curve $\gamma : [a, b] \rightarrow \mathbf{R}^n$ the integral of ω over γ is

$$\int_{\gamma} \omega = \int_{u=a}^b \left(\sum_{i=1}^n (F_i \circ \gamma)\gamma'_i \right)(u),$$

and this is the general flow integral (8.9) of the vector field (F_1, \dots, F_n) along γ . That is, the flow integrals from section 8.2 are precisely the integrals of 1-forms.

Exercises

8.4.1. Let $\omega = x dy - y dx$, a 1-form on \mathbf{R}^2 . Evaluate $\int_{\gamma} \omega$ for the following curves.

- (a) $\gamma : [-1, 1] \rightarrow \mathbf{R}^2$, $\gamma(t) = (t^2 - 1, t^3 - t)$;
- (b) $\gamma : [0, 2] \rightarrow \mathbf{R}^2$, $\gamma(t) = (t, t^2)$.

8.4.2. Let $\omega = z dx + x^2 dy + y dz$, a 1-form on \mathbf{R}^3 . Evaluate $\int_{\gamma} \omega$ for the following two curves.

- (a) $\gamma : [-1, 1] \rightarrow \mathbf{R}^3$, $\gamma(t) = (t, at^2, bt^3)$;
- (b) $\gamma : [0, 2\pi] \rightarrow \mathbf{R}^3$, $\gamma(t) = (a \cos t, a \sin t, bt)$.

8.4.3. (a) Let $\omega = f dy$ where $f : \mathbf{R}^2 \rightarrow \mathbf{R}$ depends only on y . That is, $f(x, y) = \varphi(y)$ for some $\varphi : \mathbf{R} \rightarrow \mathbf{R}$. Show that for any curve $\gamma = (\gamma_1, \gamma_2) : [a, b] \rightarrow \mathbf{R}^2$,

$$\int_{\gamma} \omega = \int_{\gamma_2(a)}^{\gamma_2(b)} \varphi.$$

(b) Let $\omega = f dx + g dy$ where f depends only on x and g depends only on y . Show that $\int_{\gamma} \omega = 0$ whenever $\gamma : [a, b] \rightarrow \mathbf{R}^2$ is a *closed* curve, meaning that $\gamma(b) = \gamma(a)$.

8.5 Examples: 2-forms on \mathbf{R}^3

To get a more complete sense of what formula (8.14) is doing, we need to study a case with $k > 1$, i.e., integration on surfaces of more than one dimension. Fortunately, the case $n = 3, k = 2$ is rich enough in geometry to understand in general how k -forms on n -space work.

Consider figure 8.5. The figure shows a 2-surface in \mathbf{R}^3 ,

$$\Phi = (\Phi_1, \Phi_2, \Phi_3) : D \rightarrow \mathbf{R}^3.$$

The parameter domain D has been partitioned into subrectangles, and the image $\Phi(D)$ has been divided up into subpatches by mapping the grid lines in D over to it via Φ . The subrectangle J of D maps to the subpatch B of $\Phi(D)$, which in turn has been projected down to its shadow $B_{(1,2)}$ in the (x, y) -plane. The point (u_J, v_J) resides in J , and its image under Φ is $\Phi(u_J, v_J) = (x_B, y_B, z_B)$.

Note that $B_{(1,2)} = (\Phi_1, \Phi_2)(J)$. Rewrite this as

$$B_{(1,2)} = \Phi_{(1,2)}(J).$$

That is, $B_{(1,2)}$ is the image of J under the $(1, 2)$ component functions of Φ . If J is small then results on determinants give

$$\text{area}(B_{(1,2)}) \approx |\det \Phi'_{(1,2)}(u_J, v_J)| \text{area}(J).$$

Thus, *the magnification factor between subrectangles of D and (x, y) -projected subpatches of $\Phi(D)$ is (up to sign) the factor $\det \Phi'_I$ from formula (8.14) for $I = (1, 2)$. The sign is somehow keeping track of the orientation of the projected patch, which would be reversed under projection onto the (y, x) -plane. (See figure 8.6.)*

Let $\omega = f dx \wedge dy$, a 2-form on \mathbf{R}^3 , where $f : \mathbf{R}^3 \rightarrow \mathbf{R}$ is a smooth function. By (8.14) and Riemann sum approximation,




Figure 8.5. 2-surface in 3-space




Figure 8.6. Projected patch and its reversal

$$\begin{aligned}
 \int_{\Phi} \omega &= \int_D (f \circ \Phi) \det \Phi'_{(1,2)} \\
 &\approx \sum_J (f \circ \Phi)(u_J, v_J) \det \Phi'_{(1,2)}(u_J, v_J) \text{area}(J) \\
 &\approx \sum_B f(x_B, y_B, z_B) (\pm \text{area}(B_{(1,2)})).
 \end{aligned}$$

This gives a geometric interpretation of what it means to integrate $f dx \wedge dy$ over Φ : to evaluate $\int_{\Phi} f dx \wedge dy$, traverse the set $\Phi(D)$ and measure projected, oriented area in the (x, y) -plane, weighted by the density function f . And similarly for forms with $dy \wedge dz$ and so on.

For an illustrative example, consider the forms $dx \wedge dy$, $dz \wedge dx$, and $dy \wedge dz$ integrated over the arch surface

$$\Phi : [-1, 1] \times [0, 1] \longrightarrow \mathbf{R}^3, \quad \Phi(u, v) = (u, v, 1 - u^2).$$

(See figure 8.7.) The (x, y) -shadows of B_1, B_2 have the same areas as J_1, J_2 and positive orientation, so $\int_{\Phi} dx \wedge dy$ should be equal to $\text{area}(D)$, i.e., 2. (See the left half of figure 8.8.) The (z, x) -shadows of B_1, B_2 have area zero, so $\int_{\Phi} dz \wedge dx$ should be an emphatic 0. (See the right half of figure 8.8.) The (y, z) -shadows of B_1, B_2 have the same area but opposite orientations, so $\int_{\Phi} dy \wedge dz$ should be 0 by some cancellation on opposite sides of the (y, z) -plane or equivalently, cancellation in the u -direction of the parameter domain. (See figure 8.9.)




Figure 8.7. An arch

Integrating with formula (8.14) confirms this intuition. Since

$$\Phi'(u, v) = \begin{bmatrix} 1 & 0 \\ 0 & 1 \\ -2u & 0 \end{bmatrix},$$

we have




Figure 8.8. (x, y) -shadows and (z, x) -shadows




Figure 8.9. (y, z) -shadows

$$\int_{\Phi} dx \wedge dy = \int_D \det \Phi'_{(1,2)} = \int_{v=0}^1 \int_{u=-1}^1 \det \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} = 2,$$

and similarly

$$\int_{\Phi} dz \wedge dx = \int_D \det \Phi'_{(3,1)} = \int_{v=0}^1 \int_{u=-1}^1 \det \begin{bmatrix} -2u & 0 \\ 1 & 0 \end{bmatrix} = \int_v \int_u 0 = 0,$$

$$\int_{\Phi} dy \wedge dz = \int_D \det \Phi'_{(2,3)} = \int_{v=0}^1 \int_{u=-1}^1 \det \begin{bmatrix} 0 & 1 \\ -2u & 0 \end{bmatrix} = \int_v \int_u 2u = 0.$$

Note how the first integral reduces to integrating 1 over the parameter domain, the second integral vanishes because its integrand is zero, and the third integral vanishes because of cancellation in the u -direction. All three of these behaviors confirm our geometric insight into how forms should behave.

Since the differential form $dx \wedge dy$ measures projected area in the (x, y) -plane, the integral

$$\int_{\Phi} z \, dx \wedge dy$$

should give the volume under the arch. And indeed formula (8.14) gives

$$\int_{\Phi} z \, dx \wedge dy = \int_{(u,v) \in D} (1 - u^2) \cdot 1,$$

which is the volume. Specifically, the integral is

$$\int_{\Phi} z \, dx \wedge dy = \int_{v=0}^1 \int_{u=-1}^1 (1 - u^2) = 1 \cdot (2 - u^3/3) \Big|_{-1}^1 = 4/3.$$

Similarly, since $dy \wedge dz$ measures oriented projected area in the (y, z) -plane, integrating the differential form $x \, dy \wedge dz$ should also give the volume under the arch. Here the interesting feature is that for $x > 0$ the form will multiply the positive distance from the (y, z) -plane to the arch by positive (y, z) -area, while for $x < 0$ the form will multiply the negative distance from the plane to the arch by negative (y, z) -area, again measuring a positive quantity. To see explicitly that the integral is again $4/3$, compute:

$$\int_{\Phi} x \, dy \wedge dz = \int_{v=0}^1 \int_{u=-1}^1 u \cdot 2u = 1 \cdot (2/3)u^3 \Big|_{-1}^1 = 4/3.$$

With these examples, the meaning of a k -form $\omega = f dx_I$ on n -space is fairly clear:

Integrating ω over a surface $\Phi : D \rightarrow \mathbf{R}^n$ means traversing the set $\Phi(D)$ and measuring oriented, k -dimensional volume of $\Phi(D)$ projected into k -space \mathbf{R}^k weighted by the density function f .

This interpretation explains the results from integrating various 1-forms over the helix in the previous section. Those integrals deserve reviewing in light of this interpretation.

As the last example of this section, consider a 2-form on \mathbf{R}^3 ,

$$\omega = F_1 \, dx_2 \wedge dx_3 + F_2 \, dx_3 \wedge dx_1 + F_3 \, dx_1 \wedge dx_2.$$

For any 2-surface $\Phi : D \rightarrow \mathbf{R}^3$ the integral of ω over Φ is

$$\int_{\Phi} \omega = \int_{u \in D} \left(\begin{array}{l} (F_1 \circ \Phi)(D_1\Phi_2 D_2\Phi_3 - D_1\Phi_3 D_2\Phi_2) \\ + (F_2 \circ \Phi)(D_1\Phi_3 D_2\Phi_1 - D_1\Phi_1 D_2\Phi_3) \\ + (F_3 \circ \Phi)(D_1\Phi_1 D_2\Phi_2 - D_1\Phi_2 D_2\Phi_1) \end{array} \right) (u),$$

and this is the flux integral (8.8) of the vector field (F_1, F_2, F_3) through Φ . A straightforward generalization of this example shows that the general integral of an $(n - 1)$ -form over an $(n - 1)$ -surface in \mathbf{R}^n is the general flux

integral (8.12). That is, the flux integrals from section 8.2 are precisely the integrals of $(n - 1)$ -forms.

Along with the last example of the previous section, this raises the question: Why bother with k -forms for values of k other than 1 and $n - 1$, and maybe also 0 and n ? The answer is that the amalgamation of k -forms for all values of k has a coherent algebraic structure, making the whole easier to study than its parts. The remainder of the chapter is largely an elaboration of this point.

After this discussion of the mechanics and meaning of integrating forms, you should be ready to prove a result that has already been mentioned: integration of forms reduces to ordinary integration when $k = n$, and integration of forms is unaffected by reasonable changes of parameterization. These points are covered in the next set of exercises.

Exercises

8.5.1. Let a be a positive number. Consider a 2-surface in \mathbf{R}^3 ,

$$\Phi : [0, a] \times [0, \pi] \longrightarrow \mathbf{R}^3, \quad \Phi(r, \theta) = (r \cos \theta, r \sin \theta, r^2).$$

Sketch this surface, noting that θ varies from 0 to π , not from 0 to 2π . Try to determine $\int_{\Phi} dx \wedge dy$ by geometrical reasoning, and then check your answer by using (8.14) to evaluate the integral. Do the same for $dy \wedge dz$ and $dz \wedge dx$.

8.5.2. Let $\omega = x dy \wedge dz + y dx \wedge dy$, a 2-form on \mathbf{R}^3 . Evaluate $\int_{\Phi} \omega$ when Φ is the 2-surface (a) $\Phi : [0, 1] \times [0, 1] \longrightarrow \mathbf{R}^3$, $\Phi(u, v) = (u + v, u^2 - v^2, uv)$; (b) $\Phi : [0, 2\pi] \times [0, 1] \longrightarrow \mathbf{R}^3$, $\Phi(u, v) = (v \cos u, v \sin u, u)$.

8.5.3. Consider a 2-form on \mathbf{R}^4 ,

$$\begin{aligned} \omega = & F_{1,2} dx_1 \wedge dx_2 + F_{1,3} dx_1 \wedge dx_3 + F_{1,4} dx_1 \wedge dx_4 \\ & + F_{2,3} dx_2 \wedge dx_3 + F_{2,4} dx_2 \wedge dx_4 + F_{3,4} dx_3 \wedge dx_4. \end{aligned}$$

Show that for any 2-surface $\Phi : D \longrightarrow \mathbf{R}^4$, the integral of ω over Φ is given by formula (8.13) from near the end of section 8.2.

8.5.4. This exercise proves that integration of k -forms on \mathbf{R}^n reduces to standard integration when $k = n$

Let $D \subset \mathbf{R}^n$ be compact and connected. Define the corresponding natural parameterization, $\Delta : D \longrightarrow \mathbf{R}^n$, by $\Delta(u_1, \dots, u_n) = (u_1, \dots, u_n)$. (This is how to turn a set in \mathbf{R}^n , where we can integrate functions, into the corresponding surface, where we can integrate n -forms.) Let $\omega = f dx_1 \wedge \dots \wedge dx_n$, an n -form on \mathbf{R}^n . Use (8.14) to show that

$$\int_{\Delta} \omega = \int_D f.$$

In particular if $f = 1$, then $\omega = dx_1 \wedge \cdots \wedge dx_n$ and $\int_{\Delta} \omega = \text{vol}(D)$, explaining why in this case ω is called the **volume form**.

Thus in \mathbf{R}^n , we may from now on blur the distinction between integrating the function f over a set and integrating the n -form $\omega = f dx_I$ over a surface, provided that $I = (1, \dots, n)$ (i.e., the dx_i factors appear in canonical order), and provided that the surface is parameterized trivially.

8.5.5. This exercise proves that because of the Change of Variable Theorem, the integration of differential forms is invariant under orientation-preserving reparameterizations of a surface.

Let A be an open subset of \mathbf{R}^n . Let $\Phi : D \rightarrow A$ and $\Psi : \tilde{D} \rightarrow A$ be k -surfaces in A . Suppose that there exists a smoothly invertible mapping $T : D \rightarrow \tilde{D}$ such that $\Psi \circ T = \Phi$. In other words, T is smooth, T is invertible, its inverse is also smooth, and the following diagram commutes:

$$\begin{array}{ccc} D & & \\ \downarrow T & \searrow \Phi & \\ & & A \\ & \nearrow \Psi & \\ \tilde{D} & & \end{array}$$

If $\det T' > 0$ on D then the surface Ψ is called an **orientation-preserving reparameterization** of Φ , whereas if $\det T' < 0$ on D then Ψ is an **orientation-reversing reparameterization** of Φ .

(a) Let T be a reparameterization as just defined. Let $S = T^{-1} : \tilde{D} \rightarrow D$, a smooth mapping. Starting from the relation $(S \circ T)(u) = id(u)$ for all $u \in D$ (where id is the identity mapping on \tilde{D}), differentiate, use the chain rule, and take determinants to show that $\det T'(u) \neq 0$ for all $u \in D$.

(b) Assume now that the reparameterization T is orientation-preserving. For any n -by- k matrix M and any ordered k -tuple I from $\{1, \dots, n\}$, recall that M_I denotes the k -by- k matrix comprised of the I th rows of M . If N is a k -by- k matrix, prove the equality

$$(MN)_I = M_I N.$$

In words, this says that

the I th rows of $(M \text{ times } N)$ are (the I th rows of M) times N .

(Suggestion: Do it first for the case $I = i$, that is, I denotes a single row.)

(c) Use the chain rule and part (b) to show that for any I ,

$$\det \Phi'_I(u) = \det \Psi'_I(T(u)) \det T'(u) \quad \text{for all } u \in D.$$

(d) Let $\omega = f(x) dx_I$, a k -form on A . Show that

$$\int_{\Psi} \omega = \int_{T(D)} (f \circ \Psi) \det \Psi'_I.$$

Explain why the Change of Variable Theorem shows that

$$\int_{\Psi} \omega = \int_D ((f \circ \Psi) \det \Psi'_I) \circ T \cdot \det T'.$$

Explain why this shows that

$$\int_{\Psi} \omega = \int_{\Phi} \omega.$$

What would the conclusion be for orientation-reversing Ψ ?

(e) Do the results from (d) remain valid if ω has the more general form $\omega = \sum_I f_I dx_I$?

8.6 Algebra of Forms: Basic Properties

One advantage of forms over earlier set-ups of vector integral calculus is that one can do much of the necessary work with them algebraically. That is, crucial properties will follow from purely rule-driven symbolic manipulation rather than geometric intuition or close analysis.

Let A be an open subset of \mathbf{R}^n . Since k -forms on A are functions (functions of k -surfaces), they come with an inherent notion of equality. The meaning of

$$\omega_1 = \omega_2$$

is that $\omega_1(\Phi) = \omega_2(\Phi)$ for all k -surfaces Φ in A . In particular, the meaning of $\omega = 0$ is that $\omega(\Phi) = 0$ for all Φ , where the first 0 is a form while the second is a real number. Addition of k -forms is defined naturally,

$$(\omega_1 + \omega_2)(\Phi) = \omega_1(\Phi) + \omega_2(\Phi) \quad \text{for all } \omega_1, \omega_2, \Phi,$$

where the first “+” is between forms, the second between real numbers. Similarly for scalar multiplication,

$$(c\omega)(\Phi) = c(\omega(\Phi)) \quad \text{for all } c, \omega, \Phi.$$

The addition of forms here is compatible with the twofold use of summation in the definition of forms and how they integrate. Addition and scalar multiplication of forms inherit all the vector space properties from corresponding properties of addition and multiplication in the real numbers, showing that the set of all k -forms on A forms a vector space. Proving familiar looking facts about addition and scalar multiplication of forms reduces quickly to citing the analogous facts in \mathbf{R} . For example, $(-1)\omega = -\omega$ for any k -form ω (where the second minus sign denotes additive inverse), because for any k -surface Φ ,

$$(\omega + (-1)\omega)(\Phi) = \omega(\Phi) + ((-1)\omega)(\Phi) = \omega(\Phi) + (-1)\omega(\Phi) = 0,$$

the last equality holding since $(-1)x = -x$ for all real numbers x .

Forms have other algebraic properties that are less familiar. For example, on \mathbf{R}^2 , $dy \wedge dx = -dx \wedge dy$. This is because of the skew symmetry of the determinant: For any 2-surface $\Phi : D \rightarrow \mathbf{R}^2$,

$$(dy \wedge dx)(\Phi) = \int_D \det \Phi'_{(2,1)} = - \int_D \det \Phi'_{(1,2)} = -(dx \wedge dy)(\Phi).$$

More generally, given two k -tuples I and J from $\{1, \dots, n\}$, $dx_J = -dx_I$ if J is obtained from I by an odd number of transpositions. Thus for example,

$$dz \wedge dy \wedge dx = -dx \wedge dy \wedge dz$$

since $(3, 2, 1)$ is obtained from $(1, 2, 3)$ by swapping the first and third entries. Showing this reduces again to the skew symmetry of the determinant. As a special case, $dx_I = 0$ whenever the k -tuple I has two matching entries. This is because exchanging those matching entries has no effect on I but negates dx_I , and so $dx_I = -dx_I$, forcing $dx_I = 0$. One can also verify directly that $dx_I = 0$ if I has matching entries by referring back to the fact that the determinant of a matrix with matching rows vanishes.

Using these rules ($dy \wedge dx = -dx \wedge dy$, $dx \wedge dx = 0$ and their generalizations), one quickly convinces oneself that every k -form can be written

$$\omega = \sum_I f_I dx_I$$

where the sum is over only increasing k -tuples $I = (i_1, \dots, i_k)$, meaning those k -tuples such that $i_1 < \dots < i_k$, as mentioned in exercise 8.3.1. This is the standard presentation for ω mentioned in Exercise 8.3.2. It is not hard to show that the standard presentation for ω is unique. In particular, $\omega = 0$ as a function of surfaces if and only if ω has standard presentation 0.

The next few sections will define certain operations on forms and develop rules of algebra for manipulating the forms under these operations. Like other rules of algebra, they will be unfamiliar at first and deserve to be scrutinized critically, but eventually they should become second nature and you should find yourself skipping steps fluently.

Exercise

8.6.1. Show that if ω is a k -form on \mathbf{R}^n that satisfies $\omega = -\omega$, then $\omega = 0$.

8.7 Algebra of Forms: Multiplication

Given a k -tuple $I = (i_1, \dots, i_k)$ and an l -tuple $J = (j_1, \dots, j_l)$ both from $\{1, \dots, n\}$, define their concatenation (I, J) , a $(k+l)$ -tuple from $\{1, \dots, n\}$, in

the obvious way: $(I, J) = (i_1, \dots, i_k, j_1, \dots, j_\ell)$. Also, if f and g are functions on an open subset A of \mathbf{R}^n then their product fg is the function defined by the formula $(fg)(x) = f(x)g(x)$.

Definition 8.7.1 (Wedge Product). *Let A be an open subset of \mathbf{R}^n . If $\omega = \sum_I f_I dx_I$ and $\lambda = \sum_J g_J dx_J$ are respectively a k -form and an ℓ -form on A , then their **wedge product** $\omega \wedge \lambda$ is a $(k + \ell)$ -form on A ,*

$$\omega \wedge \lambda = \sum_{I, J} f_I g_J dx_{(I, J)}.$$

For convenient notation, let $\Lambda^k(A)$ denote the vector space of k -forms on A . Thus the wedge product is a mapping,

$$\wedge : \Lambda^k(A) \times \Lambda^\ell(A) \longrightarrow \Lambda^{k+\ell}(A).$$

For example, a wedge product of a 1-form and a 2-form on \mathbf{R}^3 is

$$\begin{aligned} (f_1 dx + f_2 dy + f_3 dz) \wedge (g_1 dy \wedge dz + g_2 dz \wedge dx + g_3 dx \wedge dy) \\ &= f_1 g_1 dx \wedge dy \wedge dz + f_1 g_2 dx \wedge dz \wedge dx + f_1 g_3 dx \wedge dx \wedge dy \\ &\quad + f_2 g_1 dy \wedge dy \wedge dz + f_2 g_2 dy \wedge dz \wedge dx + f_2 g_3 dy \wedge dx \wedge dy \\ &\quad + f_3 g_1 dz \wedge dy \wedge dz + f_3 g_2 dz \wedge dz \wedge dx + f_3 g_3 dz \wedge dx \wedge dy \\ &= (f_1 g_1 + f_2 g_2 + f_3 g_3) dx \wedge dy \wedge dz. \end{aligned}$$

This shows that the wedge product automatically encodes the inner product in \mathbf{R}^3 , and the idea generalizes easily to \mathbf{R}^n . For another example, a wedge product of two 1-forms on \mathbf{R}^3 is

$$\begin{aligned} (x_u dx + y_u dy + z_u dz) \wedge (x_v dx + y_v dy + z_v dz) \\ &= (y_u z_v - z_u y_v) dy \wedge dz \\ &\quad + (z_u x_v - x_u z_v) dz \wedge dx \\ &\quad + (x_u y_v - y_u x_v) dx \wedge dy. \end{aligned}$$

Comparing this to the formula for the cross product in section 3.10 shows that the wedge product automatically encodes the cross product. Similarly, a wedge product of two 1-forms on \mathbf{R}^2 is

$$(a dx + b dy) \wedge (c dx + d dy) = (ad - bc) dx \wedge dy,$$

showing that the wedge product encodes the 2-by-2 determinant as well. Lemma 8.9.2 to follow will show that it encodes the general n -by- n determinant.

Naturally the wedge in Definition 8.7.1 is the same as the one in Definition 8.3.1. There is no conflict in now saying that the two wedges are the same, since each wedge in the earlier definition sits between two 1-forms and

the definition attached no meaning to the wedge symbol. Definition 8.3.1 also juxtaposes functions (0-forms) and dx_I terms (k -forms) without putting a wedge between them, and it is still unclear what sort of multiplication that juxtaposition connotes. In fact, it is also a wedge product, but when we wedge-multiply a 0-form and a k -form we usually suppress the wedge. A basic property of the wedge, its skew symmetry, will explain why in a moment.

Proposition 8.7.2 (Properties of the Wedge Product). *Let A be an open subset of \mathbf{R}^n . The wedge product has the following properties.*

- (1) *The wedge product distributes over form addition: for all $\omega \in \Lambda^k(A)$ and $\lambda_1, \lambda_2 \in \Lambda^\ell(A)$,*

$$\omega \wedge (\lambda_1 + \lambda_2) = \omega \wedge \lambda_1 + \omega \wedge \lambda_2.$$

- (2) *The wedge product is associative: for all $\omega \in \Lambda^k(A)$, $\lambda \in \Lambda^\ell(A)$, and $\mu \in \Lambda^m(A)$,*

$$(\omega \wedge \lambda) \wedge \mu = \omega \wedge (\lambda \wedge \mu).$$

- (3) *The wedge product is skew symmetric: for all $\omega \in \Lambda^k(A)$ and $\lambda \in \Lambda^\ell(A)$,*

$$\lambda \wedge \omega = (-1)^{k\ell} \omega \wedge \lambda.$$

The proof is an exercise. The unfamiliar (and hence interesting) property is the third one. The essence of its proof is to show that for any k -tuple I and any ℓ -tuple J ,

$$dx_J \wedge dx_I = (-1)^{k\ell} dx_I \wedge dx_J.$$

This follows from counting transpositions.

Note that the skew symmetry of the wedge product reduces to symmetry (i.e., commutativity) when either of the forms being multiplied is a 0-form. This is why one generally doesn't bother writing the wedge when a 0-form is involved. In fact, the wedge symbol is unnecessary in all cases, and typically in multivariable calculus one sees, for example,

$$dx \, dy \, dz \quad \text{rather than} \quad dx \wedge dy \wedge dz.$$

Indeed, we could use mere juxtaposition to denote form-multiplication, but because this new multiplication obeys unfamiliar rules, giving it a new symbol reminds us of its novel properties as we study it.

Also, the special case of multiplying a constant function c and a k -form ω is consistent with scalar multiplication of c (viewed now as a real number) and ω . Thus all of our notions of multiplication are compatible.

Exercises

8.7.1. Find a wedge product of two differential forms that encodes the inner product of \mathbf{R}^4 .

8.7.2. Find a wedge product of three differential forms that encodes the 3-by-3 determinant.

8.7.3. Prove the properties of the wedge product.

8.7.4. Prove that $(\omega_1 + \omega_2) \wedge \lambda = \omega_1 \wedge \lambda + \omega_2 \wedge \lambda$ for all $\omega_1, \omega_2 \in \Lambda^k(A)$ and $\lambda \in \Lambda^\ell(A)$. (Use skew symmetry, distributivity, and skew symmetry again.)

8.8 Algebra of Forms: Differentiation

Definition 8.8.1 (Derivative of a Differential Form). *Let A be an open subset of \mathbf{R}^n . For each integer $k \geq 0$ define the derivative mapping,*

$$d: \Lambda^k(A) \longrightarrow \Lambda^{k+1}(A),$$

by the rules

$$\begin{aligned} df &= \sum_{i=1}^n D_i f dx_i && \text{for a 0-form } f, \\ d\omega &= \sum_I df_I \wedge dx_I && \text{for a } k\text{-form } \omega = \sum_I f_I dx_I. \end{aligned}$$

For example, we saw in section 8.4 that for a function f , the 1-form

$$df = D_1 f dx_1 + \cdots + D_n f dx_n$$

is the form that measures change in f along curves. To practice this new kind of function-differentiation in a specific case, define the function

$$\pi_1: \mathbf{R}^3 \longrightarrow \mathbf{R}$$

to be projection onto the first coordinate,

$$\pi_1(x, y, z) = x \quad \text{for all } (x, y, z) \in \mathbf{R}^3.$$

Then by the definition of the derivative,

$$d\pi_1 = D_1 \pi_1 dx + D_2 \pi_1 dy + D_3 \pi_1 dz = dx. \quad (8.15)$$

This is purely routine. In practice, however, one often blurs the distinction between the name of a function and its output, for instance speaking of “the function x^2 ” rather than “the function $f: \mathbf{R} \rightarrow \mathbf{R}$ where $f(x) = x^2$ ” or “the squaring function on \mathbf{R} .” Such sloppiness is usually harmless enough and indeed downright essential in any explicit calculation where we compute using a function’s values. But in this case, if we blur the distinction between π_1 and its output x then the calculation of $d\pi_1$ in (8.15) rewrites as

$$dx = dx. \quad (!)$$

The two sides here mean different things. The left side is the operator d acting on the projection function x , while the right side is a single entity, the 1-form denoted dx . The equation is better written

$$d(x) = dx.$$

However it is written, this equality ensures that there is no possible conflict between naming the differential operator d and using this same letter as part of the definition of differential form.

Similarly, for a function of one variable $f : \mathbf{R} \rightarrow \mathbf{R}$, the definition of d immediately says that

$$df = \frac{df}{dx} dx,$$

where the single, indivisible symbol df/dx is the Leibniz notation for the derivative of f . This relation, which is sometimes presented in first semester calculus with nebulous meanings attached to df and dx , and which can *NOT* be proved by cancellation, is now a relation between 1-forms that follows from the definition of d . The moral is that the operator d has been so named to make such vague, undefined formulas into definitions and theorems. For more examples of differentiation, if

$$\omega = x dy - y dx$$

then according to Definition 8.8.1,

$$d\omega = (D_1x dx + D_2x dy) \wedge dy - (D_1y dx + D_2y dy) \wedge dx = 2 dx \wedge dy.$$

And if

$$\omega = x dy \wedge dz + y dz \wedge dx + z dx \wedge dy$$

then

$$d\omega = 3 dx \wedge dy \wedge dz.$$

The differentiation operator d commutes with sums and scalar multiples. That is, if ω_1, ω_2 are k -forms and c is a constant then

$$d(c\omega_1 + \omega_2) = c d\omega_1 + d\omega_2.$$

More interesting are the following two results about form differentiation.

Theorem 8.8.2 (Product Rule for Differential Forms). *Let A be an open subset of \mathbf{R}^n . Let ω and λ be respectively a k -form and an ℓ -form on A . Then*

$$d(\omega \wedge \lambda) = d\omega \wedge \lambda + (-1)^k \omega \wedge d\lambda.$$

Proof. Start with the case of 0-forms f and g . Then

$$\begin{aligned}
 d(fg) &= \sum_{i=1}^n D_i(fg) dx_i \\
 &= \sum_{i=1}^n (D_i f g + f D_i g) dx_i \\
 &= \left(\sum_{i=1}^n D_i f dx_i \right) g + f \left(\sum_{i=1}^n D_i g dx_i \right) \\
 &= df g + f dg.
 \end{aligned}$$

Next consider a k -form and an ℓ -form with one term each, $f_I dx_I$ and $g_J dx_J$. Then

$$\begin{aligned}
 d(f_I dx_I \wedge g_J dx_J) &= d(f_I g_J dx_I \wedge dx_J) && \text{by definition of multiplication} \\
 &= d(f_I g_J) \wedge dx_I \wedge dx_J && \text{by definition of } d \\
 &= (df_I g_J + f_I dg_J) \wedge dx_I \wedge dx_J && \text{by the result just shown} \\
 & && \text{for 0-forms} \\
 &= df_I (g_J \wedge dx_I) \wedge dx_J && \text{by distributivity and} \\
 &\quad + f_I (dg_J \wedge dx_I) \wedge dx_J && \text{associativity of } \wedge \\
 &= df_I \wedge (-1)^{0 \cdot k} (dx_I \wedge g_J) \wedge dx_J && \text{by skew symmetry} \\
 &\quad + f_I (-1)^{1 \cdot k} (dx_I \wedge dg_J) \wedge dx_J \\
 &= (df_I \wedge dx_I) \wedge g_J dx_J && \text{by associativity and symmetry.} \\
 &\quad + (-1)^k f_I dx_I \wedge (dg_J \wedge dx_J)
 \end{aligned}$$

Finally in the general case $\omega = \sum_I f_I dx_I$, $\lambda = \sum_J g_J dx_J$,

$$\begin{aligned}
 d(\omega \wedge \lambda) &= d \left(\sum_I f_I dx_I \wedge \sum_J g_J dx_J \right) = \sum_{I,J} d(f_I dx_I \wedge g_J dx_J) \\
 &= \sum_{I,J} (df_I \wedge dx_I) \wedge g_J dx_J + (-1)^k f_I dx_I \wedge (dg_J \wedge dx_J) \\
 &= \left(\sum_I df_I \wedge dx_I \right) \wedge \left(\sum_J g_J dx_J \right) \\
 &\quad + (-1)^k \left(\sum_I f_I dx_I \right) \wedge \left(\sum_J dg_J \wedge dx_J \right) \\
 &= d\omega \wedge \lambda + (-1)^k \omega \wedge d\lambda.
 \end{aligned}$$

□

As the last step in this proof consisted of tediously pushing sums through the other operations, it will typically be omitted from now on, and proofs will be carried out for the case of one-term forms.

Theorem 8.8.3 (Nilpotence of d). *Let A be an open subset of \mathbf{R}^n . Then $d^2\omega = 0$ for any form $\omega \in \Lambda^k(A)$, where d^2 means $d \circ d$. In other words,*

$$d^2 = 0.$$

Proof. For a 0-form f ,

$$df = \sum_{i=1}^n D_i f dx_i,$$

and so

$$d^2 f = d(df) = \sum_{i=1}^n d(D_i f) \wedge dx_i = \sum_{i,j} D_{ij} f dx_j \wedge dx_i.$$

All terms with $i = j$ cancel since $dx_i \wedge dx_i = 0$, and the rest of the terms cancel pairwise since for $i \neq j$, $D_{ji} f = D_{ij} f$ (equality of mixed partial derivatives) and $dx_i \wedge dx_j = -dx_j \wedge dx_i$ (skew symmetry of the wedge product). Thus

$$d^2 f = 0.$$

Also, for a k -form dx_I with constant coefficient function 1,

$$d(dx_I) = d(1dx_I) = (d1) \wedge dx_I = 0.$$

Next, for a one-term k -form $\omega = f dx_I$,

$$d\omega = df \wedge dx_I$$

and so by the first two calculations,

$$d^2\omega = d(df \wedge dx_I) = d^2 f \wedge dx_I + (-1)^1 df \wedge d(dx_I) = 0 + 0 = 0.$$

For a general k -form, pass sums and d^2 s through each other. \square

Exercises

8.8.1. Let $\omega = f dx + g dy + h dz$. Show that

$$d\omega = (D_2 h - D_3 g) dy \wedge dz + (D_3 f - D_1 h) dz \wedge dx + (D_1 g - D_2 f) dx \wedge dy.$$

8.8.2. Let $\omega = f dy \wedge dz + g dz \wedge dx + h dx \wedge dy$. Evaluate $d\omega$.

8.8.3. Differential forms of orders 0, 1, 2, 3 on \mathbf{R}^3 are written

$$\begin{aligned}\omega_0 &= \phi, \\ \omega_1 &= f_1 dx + f_2 dy + f_3 dz, \\ \omega_2 &= g_1 dy \wedge dz + g_2 dz \wedge dx + g_3 dx \wedge dy, \\ \omega_3 &= h dx \wedge dy \wedge dz.\end{aligned}$$

- (a) For a 0-form ϕ , what are the coefficients f_i of $d\phi$ in terms of ϕ ?
 (b) For a 1-form ω_1 , what are the coefficients g_i of $d\omega_1$ in terms of the coefficients f_i of ω_1 ?
 (c) For a 2-form ω_2 , what is the coefficient h of $d\omega_2$ in terms of the coefficients g_i of ω_2 ?

8.8.4. Classical vector analysis features the operator

$$\nabla = (D_1, D_2, D_3),$$

where the D_i are familiar partial derivative operators. Thus, for a function $\phi : \mathbf{R}^3 \rightarrow \mathbf{R}$,

$$\nabla\phi = (D_1\phi, D_2\phi, D_3\phi).$$

Similarly, for a mapping $F = (f_1, f_2, f_3) : \mathbf{R}^3 \rightarrow \mathbf{R}^3$, $\nabla \times F$ is defined in the symbolically appropriate way, and for a mapping $G = (g_1, g_2, g_3) : \mathbf{R}^3 \rightarrow \mathbf{R}^3$, so is $\langle \nabla, G \rangle$. Write down explicitly the vector-valued mapping $\nabla \times F$ and the function $\langle \nabla, G \rangle$ for F and G as just described. The vector-valued mapping $\nabla\phi$ is the **gradient** of ϕ from section 4.6,

$$\text{grad } \phi = \nabla\phi.$$

The vector-valued mapping $\nabla \times F$ is the **curl** of F ,

$$\text{curl } F = \nabla \times F.$$

And the scalar-valued function $\langle \nabla, G \rangle$ is the **divergence** of G ,

$$\text{div } G = \langle \nabla, G \rangle.$$

8.8.5. Continuing with the notation of the previous two problems, introduce correspondences between the classical scalar-vector environment and the environment of differential forms, as follows. Let

$$\begin{aligned}\vec{ds} &= (dx, dy, dz), \\ \vec{dn} &= (dy \wedge dz, dz \wedge dx, dx \wedge dy), \\ dV &= dx \wedge dy \wedge dz.\end{aligned}$$

Let id be the mapping that takes each function $\phi : \mathbf{R}^3 \rightarrow \mathbf{R}$ to itself. Let $\cdot\vec{ds}$ be the mapping that takes each vector-valued mapping $F = (f_1, f_2, f_3)$ to the 1-form

$$F \cdot \vec{ds} = f_1 dx + f_2 dy + f_3 dz.$$

Let $\cdot \vec{dn}$ be the mapping that takes each vector-valued mapping $G = (g_1, g_2, g_3)$ to the 2-form

$$G \cdot \vec{dn} = g_1 dy \wedge dz + g_2 dz \wedge dx + g_3 dx \wedge dy.$$

And let dV be the mapping that takes each function h to the 3-form

$$h dV = h dx \wedge dy \wedge dz.$$

Combine the previous problems to verify that the following diagram commutes, meaning that either path around each square yields the same result.

$$\begin{array}{ccccccc}
 \phi & \xrightarrow{\text{grad}} & (f_1, f_2, f_3) & \xrightarrow{\text{curl}} & (g_1, g_2, g_3) & \xrightarrow{\text{div}} & h \\
 \text{id} \downarrow & & \cdot \vec{ds} \downarrow & & \downarrow \cdot \vec{dn} & & \downarrow dV \\
 \phi & \xrightarrow{d} & f_1 dx & \xrightarrow{d} & g_1 dy \wedge dz & \xrightarrow{d} & h dx \wedge dy \wedge dz \\
 & & + f_2 dy & & + g_2 dz \wedge dx & & \\
 & & + f_3 dz & & + g_3 dx \wedge dy & &
 \end{array}$$

Thus the form-differentiation operator d , specialized to three dimensions, unifies the classical gradient, divergence, and curl operators.

8.8.6. Two of these operators are zero:

$$\text{curl} \circ \text{grad}, \quad \text{div} \circ \text{curl}, \quad \text{div} \circ \text{grad}.$$

Explain, using the diagram from the preceding exercise and the nilpotence of d . For a function $\phi : \mathbf{R}^3 \rightarrow \mathbf{R}$, write out the **harmonic equation** (or **Laplace's equation**),

$$\text{div}(\text{grad } \phi) = 0.$$

8.8.7. A form ω is called **closed** if $d\omega = 0$, and **exact** if $\omega = d\lambda$ for some form λ .

(a) Prove that any exact form is closed.

(b) Whether the converse holds is a much deeper question, depending on the topology of the domain A on which the forms are defined. Here is a special case of showing that a closed form is exact: A function $f : \mathbf{R}^3 \rightarrow \mathbf{R}$ is called **homogeneous of degree k** if

$$f(tx, ty, tz) = t^k f(x, y, z) \quad \text{for all } t \in \mathbf{R} \text{ and } (x, y, z) \in \mathbf{R}^3.$$

Such a function must satisfy **Euler's identity**,

$$xD_1f + yD_2f + zD_3f = kf.$$

Suppose that $\omega = f_1 dx + f_2 dy + f_3 dz$ is a closed 1-form whose coefficient functions are all homogeneous of degree k where $k \geq 0$. Show that $\omega = d\phi$ where

$$\phi = \frac{1}{k+1}(xf_1 + yf_2 + zf_3).$$

(c) Here is a closed form that is not exact: Let

$$\omega = \frac{x dy - y dx}{x^2 + y^2},$$

a 1-form on the punctured plane $A = \mathbf{R}^2 - \{(0, 0)\}$. Show that ω is closed. Integrate ω around the counterclockwise unit circle,

$$\gamma : [0, 2\pi] \longrightarrow A, \quad \gamma(t) = (\cos t, \sin t),$$

to show that there is no 0-form (i.e., function) θ on the punctured plane such that $\omega = d\theta$.

8.9 Algebra of Forms: the Pullback

Recall the Change of Variable Theorem from chapter 6: Given a change of variable mapping now called T (rather than Φ as in chapter 6), and given a function f on the range space of T , the appropriate function to integrate over the domain is obtained by composing with T and multiplying by an absolute determinant factor,

$$\int_{T(D)} f = \int_D (f \circ T) \cdot |\det T'|.$$

A generalization to forms of the notion of composing with T lets us similarly transfer forms—rather than functions—from the range space of a mapping T to the domain. This generalization will naturally include a determinant factor that is no longer encumbered by absolute value signs. The next section will show that integration of differential forms is inherently invariant under change of variable.

We start with some examples. The familiar polar coordinate mapping from (r, θ) -space to (x, y) -space is

$$T(r, \theta) = (r \cos \theta, r \sin \theta) \stackrel{\text{call}}{=} (x, y).$$

Using this formula, and thinking of T as mapping from (r, θ) -space *forward* to (x, y) -space, any form on (x, y) -space can naturally be converted *back* into a form on (r, θ) -space, simply by substituting $r \cos \theta$ for x and $r \sin \theta$ for y . If the form on (x, y) -space is named λ then the form on (r, θ) -space is denoted $T^*\lambda$. For example, the 2-form that gives area on (x, y) -space,

$$\lambda = dx \wedge dy,$$

has a naturally corresponding 2-form on (r, θ) -space,

$$T^*\lambda = d(r \cos \theta) \wedge d(r \sin \theta).$$

Working out the derivatives and then the wedge shows that

$$\begin{aligned} T^*\lambda &= (\cos \theta dr - r \sin \theta d\theta) \wedge (\sin \theta dr + r \cos \theta d\theta) \\ &= r dr \wedge d\theta. \end{aligned}$$

Thus (now dropping the wedges from the notation), this process has converted $dx dy$ into $r dr d\theta$ as required by the Change of Variable Theorem.

For another example, continue to let T denote the polar coordinate mapping, and consider a 1-form on (x, y) -space (for $(x, y) \neq (0, 0)$),

$$\omega = \frac{x dy - y dx}{x^2 + y^2}.$$

We studied this form in exercise 8.8.7(c). The corresponding 1-form on (r, θ) space (for $r \neq 0$) is

$$T^*\omega = \frac{r \cos \theta d(r \sin \theta) - r \sin \theta d(r \cos \theta)}{(r \cos \theta)^2 + (r \sin \theta)^2}.$$

Here the differentiations give

$$d(r \sin \theta) = \sin \theta dr + r \cos \theta d\theta, \quad d(r \cos \theta) = \cos \theta dr - r \sin \theta d\theta,$$

and so the form on (r, θ) space is

$$\begin{aligned} T^*\omega &= \frac{r \cos \theta (\sin \theta dr + r \cos \theta d\theta) - r \sin \theta (\cos \theta dr - r \sin \theta d\theta)}{r^2} \\ &= d\theta. \end{aligned}$$

That is, despite the conclusion of exercise 8.8.7(c), apparently the form ω is $d\theta$ in some sense after all.

For a third example, again start with the 1-form

$$\omega = \frac{x dy - y dx}{x^2 + y^2},$$

but this time consider a different change of variable mapping,

$$T(u, v) = (u^2 - v^2, 2uv) \stackrel{\text{call}}{=} (x, y).$$

The 1-form on (u, v) -space (for $(u, v) \neq (0, 0)$) corresponding to ω is now

$$T^*\omega = \frac{(u^2 - v^2) d(2uv) - 2uv d(u^2 - v^2)}{(u^2 - v^2)^2 + (2uv)^2}.$$

The derivatives are

$$d(2uv) = 2(v du + u dv), \quad d(u^2 - v^2) = 2(u du - v dv),$$

and so

$$\begin{aligned} T^*\omega &= 2 \frac{(u^2 - v^2)(v du + u dv) - 2uv(u du - v dv)}{(u^2 + v^2)^2} \\ &= 2 \frac{((u^2 - v^2)v - 2u^2v) du + ((u^2 - v^2)u + 2uv^2) dv}{(u^2 + v^2)^2} \\ &= 2 \frac{u dv - v du}{u^2 + v^2}. \end{aligned}$$

This is essentially the original form, except that it is doubled and now it is a form on (u, v) space. What's happening here is that T is the complex square mapping, which doubles angles, and our calculation somehow reflects this. The original form ω , which measures change of angle in (x, y) -space, has transformed back to the form that measures twice the change of angle in (u, v) -space.

Given a mapping, the natural process of changing variables in a differential form on the range of the mapping to produce a differential form on the domain of the mapping is called *pulling the differential form back* through the mapping. The general definition is as follows.

Definition 8.9.1 (Pullback of a Differential Form). *Let k be a nonnegative integer. Let A be an open subset of \mathbf{R}^n , and let B be an open subset of \mathbf{R}^m . Let*

$$T = (T_1, \dots, T_m) : A \longrightarrow B$$

*be a smooth mapping. Then T gives rise to a **pullback mapping** of k -forms in the other direction,*

$$T^* : \Lambda^k(B) \longrightarrow \Lambda^k(A).$$

Let the coordinates on \mathbf{R}^n be (x_1, \dots, x_n) , and let the coordinates on \mathbf{R}^m be (y_1, \dots, y_m) . For each k -tuple $I = (i_1, \dots, i_k)$ from $\{1, \dots, m\}$, let dT_I denote $dT_{i_1} \wedge \dots \wedge dT_{i_k}$. Then the pullback of any k -form on B ,

$$\omega = \sum f_I dy_I,$$

is

$$T^*\omega = \sum_I (f_I \circ T) dT_I.$$

Since each T_{i_j} is a function on A , each dT_{i_j} is a 1-form on A and the definition makes sense. As usual, when $k = 0$, the empty products dy_I and dT_I are interpreted as 1, and the pullback is simply composition,

$$T^*f = f \circ T.$$

As the examples before the definition have shown, computing pullbacks is easy and purely mechanical: given a form ω in terms of y 's and dy 's, its pullback $T^*\omega$ comes from replacing each y_i in ω by the expression $T_i(x_1, \dots, x_n)$ and then working out the resulting d 's and wedges.

The fact that pulling the form $dx \wedge dy$ back through the polar coordinate mapping produced the factor r from the Change of Variable Theorem is no coincidence.

Lemma 8.9.2 (Wedge–Determinant Lemma). *Define an n -form valued function Δ on n -tuples of n -vectors as follows. For any n vectors in \mathbf{R}^n ,*

$$\begin{aligned} a_1 &= (a_{11}, a_{12}, \dots, a_{1n}), \\ a_2 &= (a_{21}, a_{22}, \dots, a_{2n}), \\ &\vdots \\ a_n &= (a_{n1}, a_{n2}, \dots, a_{nn}), \end{aligned}$$

create the corresponding 1-forms,

$$\begin{aligned} \omega_1 &= a_{11} dx_1 + a_{12} dx_2 + \dots + a_{1n} dx_n, \\ \omega_2 &= a_{21} dx_1 + a_{22} dx_2 + \dots + a_{2n} dx_n, \\ &\vdots \\ \omega_n &= a_{n1} dx_1 + a_{n2} dx_2 + \dots + a_{nn} dx_n, \end{aligned}$$

and then define

$$\Delta(a_1, a_2, \dots, a_n) = \omega_1 \wedge \omega_2 \wedge \dots \wedge \omega_n.$$

Then

$$\Delta(a_1, a_2, \dots, a_n) = \det(a_1, a_2, \dots, a_n) dx_1 \wedge \dots \wedge dx_n.$$

That is, $\Delta = \det \cdot dx_{(1, \dots, n)}$.

We have already seen this result for $n = 2$ in section 8.7 and for $n = 3$ in exercise 8.7.2.

Proof. The only increasing n -tuple from $\{1, \dots, n\}$ is $(1, \dots, n)$. As a product of n 1-forms on \mathbf{R}^n , $\Delta(a_1, a_2, \dots, a_n)$ is an n -form on \mathbf{R}^n , and therefore it is a scalar-valued function $\delta(a_1, a_2, \dots, a_n)$ times $dx_{(1, \dots, n)}$. Various properties of the wedge product show that

- the function δ is linear in each of its vector variables, e.g.,

$$\delta(a_1, a_2 + \tilde{a}_2, \dots, a_n) = \delta(a_1, a_2, \dots, a_n) + \delta(a_1, \tilde{a}_2, \dots, a_n)$$

and

$$\delta(a_1, ca_2, \dots, a_n) = c \delta(a_1, a_2, \dots, a_n),$$

- the function δ is skew symmetric, i.e., transposing two of its vector variables changes its sign,
- and the function δ is normalized, i.e., $\delta(e_1, e_2, \dots, e_n) = 1$.

The determinant is the unique function satisfying these three conditions, so $\delta = \det$. \square

Theorem 8.9.3 (Pullback–Determinant Theorem). *Let A be an open subset of \mathbf{R}^n , and let B be an open subset of \mathbf{R}^m . Let $T : A \rightarrow B$ be a smooth mapping. Let \mathbf{R}^n have coordinates (x_1, \dots, x_n) , and let \mathbf{R}^m have coordinates (y_1, \dots, y_m) . Let $I = (i_1, \dots, i_n)$ be an n -tuple from $\{1, \dots, m\}$. Then*

$$T^*dy_I = \det T'_I dx_1 \wedge \cdots \wedge dx_n.$$

Proof. By definition,

$$\begin{aligned} T^*dy_I &= dT_I = dT_{i_1} \wedge \cdots \wedge dT_{i_n} \\ &= (D_1T_{i_1} dx_1 + \cdots + D_nT_{i_1} dx_n) \\ &\quad \wedge (D_1T_{i_2} dx_1 + \cdots + D_nT_{i_2} dx_n) \\ &\quad \vdots \\ &\quad \wedge (D_1T_{i_n} dx_1 + \cdots + D_nT_{i_n} dx_n). \end{aligned}$$

This is precisely $\Delta(T'_{i_1}, T'_{i_2}, \dots, T'_{i_n})$, so the lemma completes the proof. \square

In particular, when $m = n$ and $I = (1, \dots, n)$, the theorem says that

$$T^*(dy_1 \wedge \cdots \wedge dy_n) = \det T' dx_1 \wedge \cdots \wedge dx_n,$$

confirming the polar coordinate example early in this section. Similarly, if T is the spherical coordinate mapping,

$$T(\rho, \theta, \phi) = (\rho \cos \theta \sin \phi, \rho \sin \theta \sin \phi, \rho \cos \phi),$$

then the theorem tells us that

$$T^*(dx \wedge dy \wedge dz) = -\rho^2 \sin \phi d\rho \wedge d\theta \wedge d\phi.$$

You may want to verify this directly to get a better feel for the pullback and the lemma. In general, the Pullback–Determinant Theorem can be a big time-saver for computing pullbacks when the degree of the form equals the dimension of the domain space. Instead of multiplying out lots of wedge products, simply compute the relevant subdeterminant of a derivative matrix.

What makes the integration of differential forms invariant under change of variable is that the pullback operator commutes with everything else in sight.

Theorem 8.9.4 (Properties of the Pullback). *Let A be an open subset of \mathbf{R}^n , and let B be an open subset of \mathbf{R}^m . Let $T = (T_1, \dots, T_m) : A \rightarrow B$ be a smooth mapping. Then:*

(1) For all $\omega_1, \omega_2, \omega \in \Lambda^k(B)$ and $c \in \mathbf{R}$,

$$\begin{aligned} T^*(\omega_1 + \omega_2) &= T^*\omega_1 + T^*\omega_2, \\ T^*(c\omega) &= cT^*\omega. \end{aligned}$$

(2) For all $\omega \in \Lambda^k(B)$ and $\lambda \in \Lambda^\ell(B)$,

$$T^*(\omega \wedge \lambda) = (T^*\omega) \wedge (T^*\lambda).$$

(3) For all $\omega \in \Lambda^k(B)$,

$$T^*(d\omega) = d(T^*\omega).$$

That is, the pullback is linear, the pullback is multiplicative (meaning that it preserves products), and the pullback of the derivative is the derivative of the pullback. The results in the theorem can be expressed in commutative diagrams, as in exercise 8.8.5. Part (2) says that the following diagram commutes:

$$\begin{array}{ccc} \Lambda^k(B) \times \Lambda^\ell(B) & \xrightarrow{(T^*, T^*)} & \Lambda^k(A) \times \Lambda^\ell(A) \\ \wedge \downarrow & & \downarrow \wedge \\ \Lambda^{k+\ell}(B) & \xrightarrow{T^*} & \Lambda^{k+\ell}(A), \end{array}$$

and part (3) says that the following diagram commutes:

$$\begin{array}{ccc} \Lambda^k(B) & \xrightarrow{T^*} & \Lambda^k(A) \\ d \downarrow & & \downarrow d \\ \Lambda^{k+1}(B) & \xrightarrow{T^*} & \Lambda^{k+1}(A). \end{array}$$

All of this is especially gratifying because the pullback itself is entirely natural. Furthermore, the proofs are straightforward: all we need to do is compute, apply definitions, and recognize definitions. The only obstacle is that the process requires patience.

Proof. (1) Is immediate from the definition.

(2) For one-term forms $f dy_I$ and $g dy_J$,

$$\begin{aligned} T^*(f dy_I \wedge g dy_J) &= T^*(fg dy_{(I,J)}) && \text{by definition of multiplication} \\ &= (fg) \circ T dT_{(I,J)} && \text{by definition of the pullback} \\ &= f \circ T dT_I \wedge g \circ T dT_J && \text{since } (fg) \circ T = (f \circ T)(g \circ T) \\ &= T^*(f dy_I) \wedge T^*(g dy_J) && \text{by definition of the pullback.} \end{aligned}$$

The result on multi-term forms follows from this and (1).

(3) For a 0-form $f : \mathbf{R}^m \rightarrow \mathbf{R}$, $df = \sum_{i=1}^m D_i f dy_i$, so by the chain rule,

$$\begin{aligned}
 T^*(df) &= T^*\left(\sum_{i=1}^m D_i f dy_i\right) && \text{applying the definition of } d \\
 &= \sum_{i=1}^m (D_i f \circ T) dT_i && \text{applying the definition} \\
 & && \text{of the pullback} \\
 &= \sum_{i=1}^m D_i f \circ T \cdot \sum_{j=1}^n D_j T_i dx_j && \text{applying the definition of } d \\
 &= \sum_{j=1}^n \left[\sum_{i=1}^m D_i f \circ T \cdot D_j T_i \right] dx_j && \text{interchanging the sums} \\
 &= \sum_{j=1}^n D_j (f \circ T) dx_j && \text{recognizing the chain rule} \\
 &= d(f \circ T) && \text{recognizing the definition of } d \\
 &= d(T^* f) && \text{recognizing the pullback.}
 \end{aligned}$$

For a one-term k -form $f dy_I$, $d(f dy_I) = df \wedge dy_I$, so by (2) and the result for 0-forms,

$$\begin{aligned}
 T^*(d(f dy_I)) &= T^*(df \wedge dy_I) && \text{applying the definition of } d \\
 &= T^* df \wedge T^* dy_I && \text{since pullback and wedge commute} \\
 &= d(T^* f) \wedge T^* dy_I && \text{by the just-established result} \\
 &= d(f \circ T) \wedge dT_I && \text{by definition of the pullback, twice} \\
 &= d(f \circ T dT_I) && \text{recognizing the definition of } d \\
 &= d(T^*(f dy_I)) && \text{recognizing the pullback.}
 \end{aligned}$$

The multi-term result follows from this and (1). □

The pullback also behaves naturally with respect to composition.

Theorem 8.9.5 (Contravariance of the Pullback). *Let A be an open subset of \mathbf{R}^n , let B be an open subset of \mathbf{R}^m , and let C be an open subset of \mathbf{R}^ℓ . Let $T : A \rightarrow B$ and $S : B \rightarrow C$ be smooth mappings. Then for any form $\omega \in \Lambda^k(C)$,*

$$(S \circ T)^* \omega = (T^* \circ S^*) \omega.$$

This peculiar-looking result—that the pullback of a composition is the composition of the pullbacks, but in reverse order—is grammatically inevitable. Again, a commutative diagram expresses the idea:

$$\begin{array}{ccccc}
 \Lambda^k(C) & \xrightarrow{S^*} & \Lambda^k(B) & \xrightarrow{T^*} & \Lambda^k(A). \\
 & & \searrow & \nearrow & \\
 & & & (S \circ T)^* &
 \end{array}$$

Proof. For a 0-form $f : C \rightarrow \mathbf{R}$, the result is simply the associativity of composition,

$$(S \circ T)^* f = f \circ (S \circ T) = (f \circ S) \circ T = T^*(S^* f) = (T^* \circ S^*) f.$$

Let (x_1, \dots, x_n) be coordinates on \mathbf{R}^n , and similarly for (y_1, \dots, y_m) on \mathbf{R}^m and (z_1, \dots, z_ℓ) on \mathbf{R}^ℓ . For a one-term 1-form dz_q (where q is an integer from $\{1, \dots, \ell\}$),

$(S \circ T)^* dz_q = d(S \circ T)_q$	applying the definition of the pullback
$= d(S_q \circ T)$	since $(S \circ T)_q = S_q \circ T$
$= \sum_{i=1}^n D_i(S_q \circ T) dx_i$	applying the definition of d
$= \sum_{i=1}^n \left[\sum_{j=1}^m D_j S_q \circ T \cdot D_i T_j \right] dx_i$	by the chain rule
$= \sum_{j=1}^m D_j S_q \circ T \cdot \sum_{i=1}^n D_i T_j dx_i$	interchanging the sums
$= \sum_{j=1}^m (D_j S_q \circ T) dT_j$	recognizing the definition of d
$= T^* \left(\sum_{j=1}^m D_j S_q dy_j \right)$	recognizing the pullback
$= T^*(dS_q)$	recognizing the definition of d
$= T^*(S^* dz_q)$	recognizing the pullback
$= (T^* \circ S^*) dz_q$	by definition of composition.

Since any k -form is a sum of wedge products of 0-forms and 1-forms, and since the pullback passes through sums and products, the general case follows. \square

Recapitulating the section: To pull a differential form back through a map is to change variables in the form naturally. Because the wedge product has the determinant wired into it, so does the pullback. Because the pullback is natural, it commutes with addition, scalar multiplication, wedge multiplication, and differentiation of forms, and it anticommutes with composition of forms. That is, everything that we are doing is preserved under change of variables.

The results of this section are the technical heart of the chapter. The reader is encouraged to contrast their systematic algebraic proofs with the tricky analytic estimates in the main proofs of chapter 6. The work of this section will allow the pending proof of the Generalized Fundamental Theorem

of Integral Calculus to be carried out by algebra, an improvement over hand-waving geometry or tortuous analysis. The classical integration theorems of the nineteenth century will follow without recourse to the classical procedure of cutting a big curvy object into many pieces and then approximating each small piece by a straight piece instead. The classical procedure is either imprecise or Byzantine, but for those willing to think algebraically, the modern procedure is accurate and clear.

Exercises

8.9.1. Recall the two forms from the beginning of the section,

$$\lambda = dx \wedge dy, \quad \omega = \frac{x dy - y dx}{x^2 + y^2}.$$

Consider a mapping from the nonzero points of (u, v) -space to nonzero points of (x, y) -space.

$$T(u, v) = \left(\frac{u}{u^2 + v^2}, \frac{-v}{u^2 + v^2} \right) \stackrel{\text{call}}{=} (x, y).$$

Compute the pullbacks $T^*\lambda$ and $T^*\omega$, and explain why your answers make sense in light of the fact that T is the complex reciprocal mapping. Or better yet, in light of the fact that T is the complex reciprocal mapping, determine what $T^*\lambda$ and $T^*\omega$ need to be, and then confirm your answers by computing them.

8.9.2. Define $S : \mathbf{R}^2 \rightarrow \mathbf{R}^2$ by $S(u, v) = (u + v, uv) \stackrel{\text{call}}{=} (x, y)$. Let $\omega = x^2 dy + y^2 dx$ and $\lambda = xy dx$, forms on (x, y) -space.

(a) Compute $\omega \wedge \lambda$, $S^*(\omega \wedge \lambda)$, and $S^*(\omega \wedge \lambda)$. (Use the Pullback–Determinant Theorem.)

(b) Compute $S^*\omega$, $S^*\lambda$, and $S^*\omega \wedge S^*\lambda$. How do you check the last of these? Which of the three commutative diagrams from the section is relevant here?

(c) Compute $d\omega$ and $S^*(d\omega)$.

(d) Compute $d(S^*\omega)$. How do you check this? Which commutative diagram is relevant?

(e) Define $T : \mathbf{R}^2 \rightarrow \mathbf{R}^2$ by $T(s, t) = (s - t, se^t) \stackrel{\text{call}}{=} (u, v)$. Compute $T^*(S^*\lambda)$.

(f) What is the composite mapping $S \circ T$? Compute $(S \circ T)^*\lambda$. How do you check this, and which commutative diagram is relevant?

8.9.3. Let r be a fixed positive real number. Consider a 2-surface in \mathbf{R}^3 ,

$$\Phi : [0, 2\pi] \times [0, \pi] \rightarrow \mathbf{R}^3, \quad \Phi(\theta, \varphi) = (r \cos \theta \sin \varphi, r \sin \theta \sin \varphi, r \cos \varphi).$$

Consider also a 2-form on \mathbf{R}^3 ,

$$\omega = -(x/r) dy \wedge dz - (y/r) dz \wedge dx - (z/r) dx \wedge dy.$$

Compute the derivative matrix $\Phi'(\theta, \varphi)$, and use the Pullback-Determinant Theorem three times to compute the pullback $\Phi^*\omega$. Compare your answer to the integrand of the surface integral near the end of section 8.1 used to compute the volume of the sphere of radius r . (It follows that ω is the area-form for the particular surface Φ in this exercise, but *not* that ω is a general area-form for all surfaces.)

8.10 Change of Variable for Differential Forms

The definition of integration and the algebra of forms combine to make a Change of Variable Theorem for differential forms a triviality. First, a theorem of independent interest allows us to replace any integral of a differential form over a parameterized surface with an integral over the trivial parameterization of the surface's parameter domain.

Theorem 8.10.1 (Pullback Theorem). *Let A be an open subset of \mathbf{R}^n . Let ω be a k -form on A and let $\Phi : D \rightarrow A$ be a k -surface in A . Define a k -surface in \mathbf{R}^k ,*

$$\Delta^D : D \rightarrow \mathbf{R}^k, \quad \Delta^D(u) = u \text{ for all } u \in D.$$

Then

$$\int_{\Phi} \omega = \int_{\Delta^D} \Phi^* \omega.$$

Proof. As usual, just do the case of a one-term form, $\omega = f dx_I$. Then

$$\begin{aligned} \int_{\Phi} f dx_I &= \int_D (f \circ \Phi) \det \Phi'_I && \text{by definition, per (8.14)} \\ &= \int_{\Delta^D} (f \circ \Phi) \det \Phi'_I du_1 \wedge \cdots \wedge du_k && \text{by Exercise 8.5.4} \\ &= \int_{\Delta^D} (f \circ \Phi) \Phi^* dx_I && \text{by Theorem 8.9.3} \\ &= \int_{\Delta^D} \Phi^*(f dx_I) && \text{by definition of pullback.} \end{aligned}$$

□

The general Change of Variable Theorem for differential forms follows immediately from the Pullback Theorem and the contravariance of the pullback.

Theorem 8.10.2 (Change of Variable for Differential Forms). *Let A be an open subset of \mathbf{R}^n , and let B be an open subset of \mathbf{R}^m . Let $T : A \rightarrow B$*

be a smooth mapping. For any k -surface in A , $\Phi : D \rightarrow A$, the composition $T \circ \Phi : D \rightarrow B$ is thus a k -surface in B . Let ω be a k -form on B . Then

$$\int_{T \circ \Phi} \omega = \int_{\Phi} T^* \omega.$$

Proof. Let $\Delta : D \rightarrow \mathbf{R}^k$ be as above. Then

$$\int_{T \circ \Phi} \omega = \int_{\Delta} (T \circ \Phi)^* \omega = \int_{\Delta} \Phi^* (T^* \omega) = \int_{\Phi} T^* \omega.$$

□

The Pullback Theorem is essentially equivalent to the definition of integration once one has the Pullback–Determinant Theorem. Thus, a logically equivalent route to ours through this material is to define integration of a k -form in k -space as ordinary integration, and integration of a k -form in n -space for $k < n$ via the pullback. Doing so would have been a little tidier (there would not be two notions of integration when $k = n$ whose compatibility needs to be verified), but the approach here has the advantage that one can start integrating immediately before developing all the algebra.

Exercise

8.10.1. Let $T : \mathbf{R}^2 \rightarrow \mathbf{R}^2$ be given by $T(x_1, x_2) = (x_1^2 - x_2^2, 2x_1x_2) \stackrel{\text{call}}{=} (y_1, y_2)$. Let γ be the curve $\gamma : [0, 1] \rightarrow \mathbf{R}^2$ given by $\gamma(t) = (1, t)$ mapping the unit interval into (x_1, x_2) -space, and let $T \circ \gamma$ be the corresponding curve mapping into (y_1, y_2) -space. Let $\omega = y_1 dy_2$, a 1-form on (y_1, y_2) -space.

- Compute $\int_{T \circ \gamma} \omega$.
- Compute $T^* \omega$, the pullback of ω by T .
- Compute $\int_{\gamma} T^* \omega$. What theorem says that the answer here is the same as (a)?
- Let $\lambda = dy_1 \wedge dy_2$, the area form on (y_1, y_2) -space. Compute $T^* \lambda$.
- A rectangle in the first quadrant of (x_1, x_2) -space,

$$R = \{(x_1, x_2) : a_1 \leq x_1 \leq b_1, a_2 \leq x_2 \leq b_2\},$$

gets taken to some indeterminate patch $B = T(R)$ by T . Find the area of B , $\int_B \lambda$, using (d). (This exercise abuses notation slightly, and identifying R with its natural parameterization and B with the corresponding surface $T \circ R$.)

(f) Why does this exercise require that R live in the first quadrant? Can the restriction be weakened?

8.11 Cubes and Chains

Sections 8.7 through 8.9 introduced algebraic operators on differential forms: the wedge product, the derivative, and the pullback. The next section will

introduce a geometric operator on surfaces. The first thing to do is specialize the definition of a surface a bit. As usual, let $[0, 1]$ denote the unit interval. For $k \geq 0$ the unit k -cube is the cartesian product

$$[0, 1]^k = [0, 1] \times \cdots \times [0, 1] = \{(u_1, \dots, u_k) : u_i \in [0, 1] \text{ for } i = 1, \dots, k\}.$$

As mentioned in section 8.3, when $k = 0$ this means the one-point set whose point is $()$.

Definition 8.11.1 (Singular Cube, Standard Cube). *Let A be an open subset of \mathbf{R}^n . A singular k -cube in A is a k -surface whose parameter domain is the unit box,*

$$\Phi : [0, 1]^k \longrightarrow A.$$

In particular, the standard k -cube is

$$\Delta^k : [0, 1]^k \longrightarrow \mathbf{R}^k, \quad \Delta^k(u) = u \text{ for all } u \in [0, 1]^k.$$

As with Definition 8.1.1 of a surface, now a cube is by definition a mapping, and in particular a 0-cube is the parameterization of a point. In practice, we often blur the distinction between a mapping and its image, and under this blurring the word “cube” now encompasses noncubical objects such as a torus-surface or a solid sphere. The next definition allows us to consider more than one cube at a time. The purpose is to integrate over several cubes in succession, integrating over each of them a prescribed number of times.

Definition 8.11.2 (Chain). *Let A be an open subset of \mathbf{R}^n . A k -chain in A is a finite formal linear combination*

$$\mathcal{C} = \sum_s \nu_s \Phi_{(s)},$$

where each ν_s is an integer and each $\Phi_{(s)}$ is a singular k -cube in A . (The surface subscript is in parentheses only to distinguish it from a component function subscript.)

For example, if Φ , Ψ and Γ are singular k -cubes in \mathbf{R}^n then

$$2\Phi - 3\Psi + 23\Gamma$$

is a k -chain in \mathbf{R}^n . This is *not* the singular k -cube mapping points u to $2\Phi(u) - 3\Psi(u) + 23\Gamma(u)$ in \mathbf{R}^n . The term “formal linear combination” in the definition means that we don’t actually carry out any additions and scalings. Rather, the coefficients ν_s are to be interpreted as integration multiplicities. A k -chain, like a k -form, is a set of instructions.

Definition 8.11.3 (Integral of a k -form over a k -chain in n -space). *Let A be an open subset of \mathbf{R}^n . Let*

$$\mathcal{C} = \sum_s \nu_s \Phi_{(s)}$$

be a k -chain in A , and let ω be a k -form on A . Then the integral of ω over \mathcal{C} is

$$\int_{\mathcal{C}} \omega = \sum_s \nu_s \int_{\Phi_{(s)}} \omega,$$

This definition can be written more suggestively as

$$\int_{\sum \nu_s \Phi_{(s)}} \omega = \sum_s \nu_s \int_{\Phi_{(s)}} \omega.$$

Although \mathcal{C} is a formal linear combination, the operations on the right of the equality are literal addition and multiplication in \mathbf{R} . For example, let a and b be points in \mathbf{R}^n , and let Φ_a and Φ_b be the corresponding 0-cubes. Then for any 0-form on \mathbf{R}^n , $\omega = f : \mathbf{R}^n \rightarrow \mathbf{R}$,

$$\int_{\Phi_b - \Phi_a} \omega = f(b) - f(a).$$

One can define predictable rules for addition and scalar multiplication (integer scalars) of chains, all of which will pass through the integral sign tautologically. Especially, the Change of Variable Theorem for differential forms extends from integrals over surfaces to integrals over chains,

$$\int_{T \circ \mathcal{C}} \omega = \int_{\mathcal{C}} T^* \omega.$$

We will quote this formula in the proof of the General FTIC.

Also, one can compose a suitable function g with a chain to get another chain. That is,

$$\text{if } \mathcal{C} = \sum_s \nu_s \Phi_{(s)} \text{ then } g \circ \mathcal{C} = \sum_s \nu_s (g \circ \Phi_{(s)}),$$

assuming that everything here is sensible.

Exercise

8.11.1. Consider the “inner product” (actually, it is called a *pairing*)

$$\langle , \rangle : \{k\text{-forms on } \mathbf{R}^n\} \times \{k\text{-chains in } \mathbf{R}^n\} \rightarrow \mathbf{R}$$

defined by the rule $\langle \omega, \mathcal{C} \rangle = \int_{\mathcal{C}} \omega$. Show that this inner product is bilinear, meaning $\langle \sum_i c_i \omega_i, \mathcal{C} \rangle = \sum_i c_i \langle \omega_i, \mathcal{C} \rangle$ and $\langle \omega, \sum_i c_i \mathcal{C}_i \rangle = \sum_i c_i \langle \omega, \mathcal{C}_i \rangle$.

It makes no sense to speak of symmetry of this pairing since the arguments may not be exchanged.

Do you think the pairing is *nondegenerate*, meaning that for any fixed chain \mathcal{C} , if $\langle \omega, \mathcal{C} \rangle = 0$ for all ω then \mathcal{C} must be 0, and for any fixed ω , if $\langle \omega, \mathcal{C} \rangle = 0$ for all \mathcal{C} then ω must be 0?

8.12 Geometry of Chains: the Boundary Operator

This section defines an operator that takes k -chains to $(k - 1)$ -chains. The idea is to traverse the edge of each singular k -cube in the chain, with suitable multiplicity and orientation. The following definition gives three rules that say how to do so. The first rule reduces taking the boundary of a k -chain to taking the boundary of its constituent singular k -cubes. The second rule reduces taking the boundary of a singular k -cube to taking the boundary of the standard k -cube. The third rule, giving the procedure for taking the boundary of the standard k -cube, is the substance of the definition. It is best understood by working through specific cases.

Definition 8.12.1 (Boundary). *Let A be an open subset of \mathbf{R}^n . For each $k \geq 1$, define the **boundary mapping***

$$\partial : \{k\text{-chains in } A\} \longrightarrow \{(k - 1)\text{-chains in } A\}$$

by the properties:

(1) For any k -chain $\sum \nu_s \Phi_{(s)}$,

$$\partial \left(\sum \nu_s \Phi_{(s)} \right) = \sum \nu_s \partial \Phi_{(s)}.$$

(2) For any singular k -cube Φ ,

$$\partial \Phi = \Phi \circ \partial \Delta^k.$$

(The composition here is of the sort defined at the end of the previous section.)

(3) Define mappings from the standard $(k - 1)$ -cube to the faces of the standard k -cube as follows: for any $i \in \{1, \dots, n\}$ and $j \in \{0, 1\}$, the mapping to the face where the i th coordinate equals α is

$$\Delta_{i,\alpha}^k : [0, 1]^{k-1} \longrightarrow [0, 1]^k,$$

given by

$$\Delta_{i,\alpha}^k(u_1, \dots, u_{k-1}) = (u_1, \dots, u_{i-1}, \alpha, u_i, \dots, u_{k-1}).$$

Then

$$\partial \Delta^k = \sum_{i=1}^k \sum_{\alpha=0}^1 (-1)^{i+\alpha} \Delta_{i,\alpha}^k. \quad (8.16)$$

In property (2) the composition symbol “ \circ ” has been generalized a little from its ordinary usage. Since $\partial \Delta^k$ is a chain $\sum \mu_s \Psi_{(s)}$, the composition $\Phi \circ \partial \Delta^k$ is defined as the corresponding chain $\sum \mu_s \Phi \circ \Psi_{(s)}$. The compositions in the sum make sense since by property (3), each $\Psi_{(s)}$ maps $[0, 1]^{k-1}$ into $[0, 1]^k$.

To remember the definition of $\Delta_{i,\alpha}^k$ in (8.16), read its name as “set the i th of k variables to α ” or just “set the i th variable to α .” The idea of formula (8.16) is that for each of the directions in k -space ($i = 1, \dots, k$), the standard k -cube has two faces with normal vectors in the i th direction ($\alpha = 0, 1$), and we should take these two faces with opposite orientations in order to make both normal vectors point outward. Unlike differentiation, which increments the degree of the form it acts on, the boundary operator decrements chain dimension.

For example, the boundary of the standard 1-cube is given by (8.16)

$$\partial\Delta^1 = -\Delta_{1,0}^1 + \Delta_{1,1}^1.$$

This is the right endpoint of $[0, 1]$ with a plus and the left endpoint with a minus. (See figure 8.10. The figures for this sections show the *images* of the various mappings involved, with symbols added as a reminder that the images are being traversed by the mappings.) One consequence of this is that the familiar formula from the one-variable Fundamental Theorem of Integral Calculus,

$$\int_0^1 f' = f(1) - f(0),$$

is now expressed suggestively in the notation of differential forms as

$$\int_{\Delta^1} df = \int_{\partial\Delta^1} f.$$

As for the boundary of a singular 1-cube $\Phi : [0, 1] \rightarrow \mathbf{R}^n$ (i.e., a curve in space) with $\Phi(0) = a$ and $\Phi(1) = b$, property (2) of the boundary definition gives

$$\partial\Phi = \Phi \circ \partial\Delta^1 = -\Phi \circ \Delta_{1,0}^1 + \Phi \circ \Delta_{1,1}^1.$$

This is the curve’s endpoint b with a plus and the start-point a with a minus. The last example of section 8.4 now also takes on a more suggestive expression,

$$\int_{\Phi} df = \int_{\partial\Phi} f.$$



Figure 8.10. Standard 1-cube and its boundary

The boundary of the standard 2-cube is again given by (8.16)

$$\partial\Delta^2 = -\Delta_{1,0}^2 + \Delta_{1,1}^2 + \Delta_{2,0}^2 - \Delta_{2,1}^2.$$

This traverses the boundary square of $[0, 1]^2$ once counterclockwise. (See figure 8.11.) Next consider a singular 2-cube that parameterizes the unit disk,

$$\Phi : [0, 1]^2 \longrightarrow \mathbf{R}^2, \quad \Phi(r, \theta) = (r \cos 2\pi\theta, r \sin 2\pi\theta).$$

By property (2), $\partial\Phi = \Phi \circ \partial\Delta^2$. This traverses the boundary circle once counterclockwise, two radial traversals cancel, and there is a degenerate mapping to the centerpoint. (See figure 8.12.) Changing to $\Phi(r, \theta) = (r \cos 2\pi\theta, -r \sin 2\pi\theta)$ also parameterizes the unit disk, but now $\partial\Phi$ traverses the boundary circle clockwise.




Figure 8.11. Standard 2-cube and its boundary




Figure 8.12. Boundary of a singular 2-cube

The boundary of the standard 3-cube is, by (8.16),

$$\partial\Delta^3 = -\Delta_{1,0}^3 + \Delta_{1,1}^3 + \Delta_{2,0}^3 - \Delta_{2,1}^3 - \Delta_{3,0}^3 + \Delta_{3,1}^3.$$

This traverses the faces of $[0, 1]^3$, oriented positively if we look at them from outside the solid cube. (See figure 8.13.)

The second boundary of the standard 2-cube works out by cancellation to




Figure 8.13. Boundary of the standard 3-cube

$$\partial^2 \Delta^2 = 0.$$

(See the left side of figure 8.14.) And the second boundary of the standard 3-cube similarly is

$$\partial^2 \Delta^3 = 0.$$

(See the right side of figure 8.14.) These two examples suggest that the notational counterpart to the nilpotence of d is also true,

$$\partial^2 = 0.$$

This is indeed a theorem, and it is readily shown by a double sum calculation in which terms cancel pairwise. But it will also follow immediately from the main theorem of the chapter, the Generalized FTIC, which states that in a precise sense the differentiation operator d and the boundary operator ∂ are complementary. This is why they are notated so similarly.

Since integration is invariant under reparameterization, you needn't be too formal in computing boundaries once you understand how they work on standard cubes. The boundary of the unit square (the 2-cube), for example, is adequately described as its edge traversed counterclockwise at unit speed, and so the boundary of any singular 2-cube Φ from the unit square into \mathbf{R}^n is simply the restriction of Φ to the edge of the square with appropriate traversal, or (once we finally prove the Change of Variable Theorem from the previous chapter) any orientation-preserving reparameterization thereof. In particular,




Figure 8.14. Second boundaries

any rectangle in \mathbf{R}^2 can be obtained by scaling and translating the unit square in an orientation-preserving fashion, so the boundary of such a rectangle is, as one would hope, its edge, counterclockwise. More generally, a singular 2-cube in \mathbf{R}^3 is a sort of parameterized membrane floating in space, and its boundary is just its edge, traversed in the direction inherited from the parameterization, as we saw for the disk. Without the parameterization, neither direction of traversing the membrane's edge in \mathbf{R}^n for any $n > 2$ is naturally preferable to the other. Similarly in \mathbf{R}^3 , the boundary of the unit cube is its six faces, oriented to look positive from outside the cube. In other words, an acceptable coordinate system for a boundary face of the cube is two orthonormal vectors whose cross product is an outward unit normal to the cube. The boundary of any singular 3-cube $\Phi : [0, 1]^3 \rightarrow \mathbf{R}^3$ is the restriction of Φ to the boundary faces of $[0, 1]^3$.

Exercises

8.12.1. Define a singular k -cube called the **simplex**, $\Phi : [0, 1]^k \rightarrow \mathbf{R}^k$, by

$$\Phi(u_1, \dots, u_k) = (u_1, (1 - u_1)u_2, (1 - u_1)(1 - u_2)u_3, \dots, \prod_{i=1}^{k-1} (1 - u_i)u_k).$$

(a) Show that if $(x_1, \dots, x_k) = \Phi(u_1, \dots, u_k)$ then $\sum_{i=1}^k x_i = 1 - \prod_{i=1}^k (1 - u_i)$.

(b) Show that the image of Φ lies in the set (also called the simplex)

$$S = \{(x_1, \dots, x_k) : x_1 \geq 0, \dots, x_k \geq 0, \sum_{i=1}^k x_i \leq 1\}.$$

(In fact, the image is all of the simplex, but showing this would take us too far afield.)

(c) For each of the values $k = 1, 2, 3$, do the following. Calculate $\partial\Phi$ (the result is a $(k - 1)$ -chain). Graph $\partial\Phi$ by graphing each $(k - 1)$ -cube in the chain and indicating its coefficient (+1 or -1) beneath the graph. Each graph should show $[0, 1]^{k-1}$ and \mathbf{R}^k .

8.12.2. Describe the boundary of the hemispherical shell $H : D \rightarrow \mathbf{R}^3$ where D is the unit disk in \mathbf{R}^2 and $H(x, y) = (x, y, \sqrt{1 - x^2 - y^2})$.

8.12.3. Describe the boundary of the solid unit upper hemisphere

$$\mathbf{H} = \{(x, y, z) \in \mathbf{R}^3 : x^2 + y^2 + z^2 \leq 1, z \geq 0\}.$$

8.12.4. Describe the boundary of the paraboloid $\Phi : \{(u, v) \in \mathbf{R}^2 : u^2 + v^2 \leq 1\} \rightarrow \mathbf{R}^3$ where

$$\Phi(u, v) = (u, v, u^2 + v^2).$$

8.12.5. Describe the boundary of $\Phi : [0, 2\pi] \times [0, \pi] \rightarrow \mathbf{R}^3$ where

$$\Phi(\theta, \phi) = (\cos \theta \sin \phi, \sin \theta \sin \phi, \cos \phi).$$

8.12.6. Describe the boundary of $\Phi : [0, 1] \times [0, 2\pi] \times [0, \pi] \rightarrow \mathbf{R}^3$ where

$$\Phi(\rho, \theta, \phi) = (\rho \cos \theta \sin \phi, \rho \sin \theta \sin \phi, \rho \cos \phi).$$

8.12.7. Fix constants $0 < a < b$. Describe the boundary of $\Phi : [0, 2\pi] \times [0, 2\pi] \times [0, 1] \rightarrow \mathbf{R}^3$ where $\Phi(u, v, t) = (\cos u(b + at \cos v), \sin u(b + at \cos v), at \sin v)$.

8.13 The Generalized Fundamental Theorem of Integral Calculus

As mentioned in the previous section, the algebraic encoding d of the derivative (an analytic operator) and the algebraic encoding ∂ of the boundary (a geometric operator) are complementary with respect to integration:

Theorem 8.13.1 (Generalized FTIC). *Let A be an open subset of \mathbf{R}^n . Let \mathcal{C} be a k -chain in A , and let ω be a $(k - 1)$ -form on A . Then*

$$\int_{\mathcal{C}} d\omega = \int_{\partial\mathcal{C}} \omega. \quad (8.17)$$

Before proving the theorem, we study two examples. First, suppose that $k = n = 1$, and that the 1-chain \mathcal{C} is a singular 1-cube $\Phi : [0, 1] \rightarrow \mathbf{R}$ taking 0 and 1 to some points a and b . Then the theorem says that for any suitable smooth function f ,

$$\int_a^b f'(x) dx = f(b) - f(a).$$

This is the one-variable Fundamental Theorem of Integral Calculus. Thus, whatever else we are doing, we are indeed generalizing it.

Second, to study a simple case involving more than one variable, suppose that $\mathcal{C} = \Delta^2$ (the standard 2-cube) and $\omega = f dy$ for some smooth function $f : [0, 1]^2 \rightarrow \mathbf{R}$, i.e., f is a function of x and y . The derivative in the left side of (8.17) works out to

$$d\omega = D_1 f dx \wedge dy,$$

Exercise 8.5.4 says that we may drop the wedges from the integral of this 2-form over the full-dimensional surface Δ^2 in 2-space to obtain a chapter 6 function-integral, and so the left side of (8.17) works out to

$$\int_{\Delta^2} d\omega = \int_{\Delta^2} D_1 f dx \wedge dy = \int_{[0,1]^2} D_1 f.$$

Meanwhile, on the right side of (8.17), the boundary $\partial\Delta^2$ has four pieces, but on the two horizontal pieces dy is zero since y is constant. Thus only the integrals over the two vertical pieces contribute, giving

$$\int_{\partial\Delta^2} \omega = \int_{u=0}^1 f(1, u) - \int_{u=0}^1 f(0, u) = \int_{u=0}^1 f(1, u) - f(0, u).$$

By the one-variable Fundamental Theorem, the integrand is

$$f(1, u) - f(0, u) = \int_{t=0}^1 D_1 f(t, u),$$

and so by Fubini's Theorem, the integral is

$$\int_{u=0}^1 \int_{t=0}^1 D_1 f(t, u) = \int_{[0,1]^2} D_1 f.$$

Thus both sides of (8.17) work out to $\int_{[0,1]^2} D_1 f$, making them equal as desired, and the Generalized FTIC holds in this case. The first step of its proof is essentially the same process as in this example.

Proof. Recall that we want to establish formula (8.17), $\int_{\mathcal{C}} d\omega = \int_{\partial\mathcal{C}} \omega$, where \mathcal{C} is a k -chain and ω is a $(k-1)$ -form. Begin with the special case where \mathcal{C} is the standard k -cube,

$$\mathcal{C} = \Delta^k,$$

and ω takes the form $\omega = f dx_1 \wedge \cdots \wedge \widehat{dx_j} \wedge \cdots \wedge dx_k$, where the $\widehat{\quad}$ means to omit the term. Thus

$$\omega = f dx_J \quad \text{where } J = (1, \dots, \widehat{j}, \dots, k).$$

To evaluate the left side $\int_C d\omega$ of (8.17), we need to compute $d\omega$. In this special case

$$d\omega = D_j f dx_j \wedge dx_J = (-1)^{j-1} D_j f dx_{(1, \dots, k)},$$

and so by exercise 8.5.4, the left side reduces to the function-integral of the j th partial derivative over the unit box,

$$\int_{\Delta^k} d\omega = (-1)^{j-1} \int_{\Delta^k} D_j f dx_{(1, \dots, k)} = (-1)^{j-1} \int_{[0,1]^k} D_j f. \quad (8.18)$$

To evaluate the right side $\int_{\partial C} \omega$ of (8.17), we need to examine the boundary

$$\partial \Delta^k = \sum_{i=1}^k \sum_{\alpha=0}^1 (-1)^{i+\alpha} \Delta_{i,\alpha}^k,$$

where $\Delta_{i,\alpha}^k(u_1, \dots, u_{k-1}) = (u_1, \dots, u_{i-1}, \alpha, u_i, \dots, u_{k-1})$. Note that

$$(\Delta_{i,\alpha}^k)' = \left[\begin{array}{cccc|cccc} 1 & 0 & \cdots & 0 & 0 & \cdots & 0 & 0 \\ 0 & 1 & \cdots & 0 & 0 & \cdots & 0 & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots & & & \vdots \\ 0 & 0 & \cdots & 1 & 0 & \cdots & 0 & 0 \\ \hline 0 & 0 & \cdots & 0 & 0 & \cdots & 0 & 0 \\ 0 & 0 & \cdots & 0 & 1 & \cdots & 0 & 0 \\ \hline \vdots & \vdots & & \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & \cdots & 0 & 0 & \cdots & 1 & 0 \end{array} \right].$$

This derivative matrix is k -by- $(k-1)$, consisting of the identity matrix except that zeros have been inserted at the i th row, displacing everything from there downwards. On the other hand, recall that $J = (1, \dots, \hat{j}, \dots, k)$. It follows that the determinant of the J th rows of the matrix is

$$\det(\Delta_{i,\alpha}^k)'_J = \begin{cases} 1 & \text{if } i = j, \\ 0 & \text{if } i \neq j. \end{cases}$$

That is, the integral of $\omega = f dx_J$ can be nonzero only for the two terms in the boundary chain $\partial \Delta^k$ with $i = j$, parameterizing the two boundary faces whose normal vectors point in the direction missing from dx_J :

$$\begin{aligned} \int_{\partial \Delta^k} f dx_J &= \int_{(-1)^{j+1}(\Delta_{j,1}^k - \Delta_{j,0}^k)} f dx_J \\ &= (-1)^{j+1} \int_{[0,1]^{k-1}} (f \circ \Delta_{j,1}^k) \cdot 1 - (f \circ \Delta_{j,0}^k) \cdot 1. \end{aligned}$$

Here the last equality follows from the definition of integration over chains and the defining formula (8.14). For any point $u = (u_1, \dots, u_{k-1}) \in [0, 1]^{k-1}$,

the integrand can be rewritten as an integral of the j th partial derivative by the one-variable Fundamental Theorem of Integral Calculus,

$$\begin{aligned} (f \circ \Delta_{j,1}^k - f \circ \Delta_{j,0}^k)(u) &= f(u_1, \dots, u_{j-1}, 1, u_j, \dots, u_{k-1}) - f(u_1, \dots, u_{j-1}, 0, u_j, \dots, u_{k-1}) \\ &= \int_{t \in [0,1]} D_j f(u_1, \dots, u_{j-1}, t, u_j, \dots, u_{k-1}). \end{aligned}$$

Therefore, the right side of (8.17) is

$$\int_{\partial \Delta^k} \omega = (-1)^{j+1} \int_{u \in [0,1]^{k-1}} \int_{t \in [0,1]} D_j f(u_1, \dots, u_{j-1}, t, u_j, \dots, u_{k-1}).$$

By Fubini's Theorem this is equal to the right side of (8.18), and so the Generalized FTIC is proved in the special case.

The rest of the proof is handled effortlessly by the machinery of forms and chains. A general $(k-1)$ -form on A is

$$\omega = \sum_{j=1}^k f_j dx_1 \wedge \cdots \wedge \widehat{dx}_j \wedge \cdots \wedge dx_k = \sum_{j=1}^k \omega_j$$

where each ω_j is a form of the type covered by the special case. Also, $d\omega = \sum_j d\omega_j$. So continuing to integrate over the standard k -cube,

$$\begin{aligned} \int_{\Delta^k} d\omega &= \int_{\Delta^k} \sum_j d\omega_j = \sum_j \int_{\Delta^k} d\omega_j \\ &= \sum_j \int_{\partial \Delta^k} \omega_j = \int_{\partial \Delta^k} \sum_j \omega_j = \int_{\partial \Delta^k} \omega, \end{aligned}$$

where the crucial equality in the middle is the special case just shown. Thus the theorem holds for a general form when $C = \Delta^k$.

For a singular k -cube Φ in A and for any $(k-1)$ -form ω on A , we now have

$$\begin{aligned} \int_{\Phi} d\omega &= \int_{\Delta^k} \Phi^*(d\omega) && \text{by the Pullback Theorem} \\ &= \int_{\Delta^k} d(\Phi^*\omega) && \text{since derivative commutes with pullback} \\ &= \int_{\partial \Delta^k} \Phi^*\omega && \text{since the result holds on } \Delta^k \\ &= \int_{\Phi \circ \partial \Delta^k} \omega && \text{by the Change of Variable Theorem for} \\ & && \text{differential forms, extended to chains} \\ &= \int_{\partial \Phi} \omega && \text{by definition of boundary.} \end{aligned}$$

So the result holds for singular cubes.

Finally, for a k -chain $\mathcal{C} = \sum_s \nu_s \Phi_{(s)}$ in A and for any $(k-1)$ -form ω on A ,

$$\int_{\mathcal{C}} d\omega = \int_{\sum_s \nu_s \Phi_{(s)}} d\omega = \sum_s \nu_s \int_{\Phi_{(s)}} d\omega = \sum_s \nu_s \int_{\partial\Phi_{(s)}} \omega,$$

with the third equality due to the result for singular cubes, and the calculation continues

$$\sum_s \nu_s \int_{\partial\Phi_{(s)}} \omega = \int_{\sum_s \nu_s \partial\Phi_{(s)}} \omega = \int_{\partial(\sum_s \nu_s \Phi_{(s)})} \omega = \int_{\partial\mathcal{C}} \omega.$$

This completes the proof. \square

The beauty of this argument is that the only analytic results that it uses are the one-variable FTIC and Fubini's Theorem, and the only geometry that it uses is the definition of the boundary of a standard k -cube. All the complicated twisting and turning of k -surfaces in n -space is filtered out automatically by the algebra of differential forms.

Computationally, the Generalized FTIC will sometimes give you a choice between evaluating two integrals, one of which may be easier to work. Note that the integral of lower dimension may not be the preferable one, however; for example, integrating over a solid 3-cube may be quicker than integrating over the six faces of its boundary.

Conceptually the Generalized FTIC is exciting because it allows the possibility of evaluating an integral over a region by antidifferentiating and then integrating only over the boundary of the region.

Exercises

8.13.1. Similarly to the second example before the proof of the Generalized FTIC, show that the theorem holds when $\mathcal{C} = \Delta^3$ and $\omega = f dz \wedge dx$.

8.13.2. Prove as a corollary to the Generalized FTIC that $\partial^2 = 0$, in the sense that $\int_{\partial^2\mathcal{C}} \omega = 0$ for all forms ω .

8.13.3. Let \mathcal{C} be a k -chain in \mathbf{R}^n , $f : \mathbf{R}^n \rightarrow \mathbf{R}$ a function, and ω a $(k-1)$ -form on \mathbf{R}^n . Use the Generalized FTIC to prove a generalization of the formula for integration by parts,

$$\int_{\mathcal{C}} f d\omega = \int_{\partial\mathcal{C}} f\omega - \int_{\mathcal{C}} df \wedge \omega.$$

8.13.4. Let Φ is a 4-chain in \mathbf{R}^4 with boundary $\partial\Phi$. Prove the identity,

$$\begin{aligned} \int_{\partial\Phi} f_1 dy \wedge dz \wedge dw + f_2 dz \wedge dw \wedge dx + f_3 dw \wedge dx \wedge dy + f_4 dx \wedge dy \wedge dz \\ = \int_{\Phi} (D_1 f_1 - D_2 f_2 + D_3 f_3 - D_4 f_4) dx \wedge dy \wedge dz \wedge dw. \end{aligned}$$

Here the order of the variables is (x, y, z, w) .

8.14 The Classical Theorems

The classical integration theorems of vector calculus arise from specializing n and k in the General FTIC. As already noted, the values $n = k = 1$ give the one-variable FTIC,

$$\int_a^b \frac{df}{dx} dx = f(b) - f(a).$$

If $k = 1$ but n is left arbitrary then the result is familiar from section 8.4. For any curve $\gamma : [0, 1] \rightarrow \mathbf{R}^n$, let $a = \gamma(0)$ and $b = \gamma(1)$. Then

$$\int_{\gamma} \frac{\partial f}{\partial x_1} dx_1 + \cdots + \frac{\partial f}{\partial x_n} dx_n = f(b) - f(a).$$

Setting $n = 2$, $k = 2$ gives **Green's Theorem**: Let A be an open subset of \mathbf{R}^2 . For any singular 2-cube Φ in A and functions $f, g : A \rightarrow \mathbf{R}$,

$$\iint_{\Phi} \left(\frac{\partial g}{\partial x} - \frac{\partial f}{\partial y} \right) dx \wedge dy = \int_{\partial\Phi} f dx + g dy.$$

The double integral sign is used on the left side of Green's Theorem to emphasize the two-dimensional integral. Naturally the classical statement doesn't refer to a singular cube or include a wedge. Instead, the idea is to view Φ as a set in the plane and require a traversal of $\partial\Phi$ (also viewed as a set) such that Φ is always to the left as one moves along $\partial\Phi$. Other than this, the boundary integral is independent of how the boundary is traversed because the whole theory is to be invariant under orientation-preserving reparameterization. (See figure 8.15.)




Figure 8.15. Traversing the boundary in Green's Theorem

Green's Theorem has two geometric interpretations. To understand them, first let $A \subset \mathbf{R}^2$ be open and think of a vector-valued mapping $\vec{F} : A \rightarrow \mathbf{R}^2$ as defining a fluid flow in A . Define two related scalar-valued functions on A ,

$$\operatorname{curl} \vec{F} = D_1 F_2 - D_2 F_1 \quad \text{and} \quad \operatorname{div} \vec{F} = D_1 F_1 + D_2 F_2.$$

These are two-dimensional versions of the quantities from exercises 8.8.4 and 8.8.5. Now consider a point p in A . Note that $\operatorname{curl} \vec{F}(p)$ and $\operatorname{div} \vec{F}(p)$ depend only on the derivatives of \vec{F} at p , not on $\vec{F}(p)$ itself. So replacing \vec{F} by $\vec{F} - \vec{F}(p)$, we may assume that $\vec{F}(p) = \mathbf{0}$, i.e., the fluid flow is stationary at p . Recall that $D_1 F_2$ is the rate of change of the vertical component of F with respect to change in the horizontal component of its input, and $D_2 F_1$ is the rate of change of the horizontal component of F with respect to change in the vertical component of its input. Thus, a scenario where the two terms $D_1 F_2$ and $-D_2 F_1$ of $(\operatorname{curl} \vec{F})(p)$ are positive looks like the left side of figure 8.16. This explains why $\operatorname{curl} \vec{F}$ is interpreted as measuring the vorticity of \vec{F} at p , its tendency to rotate a paddle-wheel at p counterclockwise. Similarly, $D_1 F_1$ is the rate of change of the horizontal component of F with respect to change in the horizontal component of its input, and $D_2 F_2$ is the rate of change of the vertical component of F with respect to change in the vertical component of its input. A scenario where the terms of $(\operatorname{div} \vec{F})(p)$ are positive looks like the right side of figure 8.16. Thus $\operatorname{div} \vec{F}$ is viewed as measuring the extent that fluid is spreading out from p , i.e., how much fluid is being pumped into or drained out of the system at the point. Specifically, the left side of the figure shows the vector field

$$\vec{F}(x, y) = (-y, x)$$

whose curl and divergence at the origin are

$$(\operatorname{curl} \vec{F})(\mathbf{0}) = 2, \quad (\operatorname{div} \vec{F})(\mathbf{0}) = 0,$$

and the right side shows (with some artistic license taken to make the figure legible rather than accurate) the vector field

$$\vec{F}(x, y) = (x, y)$$

whose curl and divergence at the origin are

$$(\operatorname{curl} \vec{F})(\mathbf{0}) = 0, \quad (\operatorname{div} \vec{F})(\mathbf{0}) = 2.$$

For the two geometric interpretations of Green's Theorem, introduce the notation

$$dA = dx \wedge dy, \quad \vec{ds} = (dx, dy), \quad \vec{dn} = (dy, -dx).$$

The form-vectors \vec{ds} and \vec{dn} on $\partial\Phi$ are viewed respectively as differential increment around the boundary and differential outward normal (see the exercises), while dA is differential area. Then setting $\vec{F} = (f, g)$ and $\vec{F} = (g, -f)$ respectively shows that Green's Theorem says

$$\iint_{\Phi} \operatorname{curl} \vec{F} dA = \int_{\partial\Phi} \vec{F} \cdot \vec{ds} \quad \text{and} \quad \iint_{\Phi} \operatorname{div} \vec{F} dA = \int_{\partial\Phi} \vec{F} \cdot \vec{dn}.$$

The resulting two interpretations are




Figure 8.16. Positive curl and positive divergence

the net counterclockwise vorticity of \vec{F} throughout Φ equals the net flow of \vec{F} counterclockwise around $\partial\Phi$

and

the net positive creation of fluid by \vec{F} throughout Φ equals the net flux of \vec{F} outward through $\partial\Phi$.

These appeal strongly to physical intuition.

Setting $n = 3$, $k = 2$ gives **Stokes's Theorem**: Let A be an open subset of \mathbf{R}^3 . For a singular 2-cube Φ in A and functions $f, g, h : A \rightarrow \mathbf{R}$,

$$\begin{aligned} \iint_{\Phi} \left(\frac{\partial h}{\partial y} - \frac{\partial g}{\partial z} \right) dy \wedge dz + \left(\frac{\partial f}{\partial z} - \frac{\partial h}{\partial x} \right) dz \wedge dx + \left(\frac{\partial g}{\partial x} - \frac{\partial f}{\partial y} \right) dx \wedge dy \\ = \int_{\partial\Phi} f dx + g dy + h dz. \end{aligned}$$

Introduce the notation

$$\vec{ds} = (dx, dy, dz) \quad \text{and} \quad \vec{dn} = (dy \wedge dz, dz \wedge dx, dx \wedge dy),$$

and for a vector-valued mapping $\vec{F} : \mathbf{R}^3 \rightarrow \mathbf{R}^3$ define

$$\text{curl } \vec{F} = (D_2 F_3 - D_3 F_2, D_3 F_1 - D_1 F_3, D_1 F_2 - D_2 F_1).$$

Then setting $\vec{F} = (f, g, h)$ shows that Stokes's Theorem is

$$\iint_{\Phi} \text{curl } \vec{F} \cdot \vec{dn} = \int_{\partial\Phi} \vec{F} \cdot \vec{ds}$$

As with Green's Theorem, the classical statement doesn't refer to a singular cube or include a wedge. Instead, Φ is an orientable two-dimensional set in

space, and its boundary $\partial\Phi$ is traversed counterclockwise about its normal vectors. The integrals on both sides of the equality are independent of how Φ and $\partial\Phi$ are parameterized, provided that the geometry is just described.

For the interpretation of Stokes's Theorem, think of a mapping $\vec{F} : \mathbf{R}^3 \rightarrow \mathbf{R}^3$ as describing a fluid flow in space. The mapping $\text{curl } \vec{F}$ is interpreted as measuring the local vorticity of \vec{F} around each positive coordinate direction. The form-vector $\vec{d}\vec{n}$ on Φ is viewed as differential outward normal, while $\vec{d}\vec{s}$ on $\partial\Phi$ is viewed as differential increment around the boundary. Thus the interpretation of Stokes's Theorem is a 3-dimensional version of the first interpretation of Green's Theorem,

*the net tangent vorticity of \vec{F} throughout Φ
equals the net flow of \vec{F} around $\partial\Phi$.*

Setting $n = 3$, $k = 3$ gives the **Divergence Theorem** (or **Gauss's Theorem**): Let A be an open subset of \mathbf{R}^3 . For a singular 3-cube Φ in A and functions $f, g, h : A \rightarrow \mathbf{R}$,

$$\iiint_{\Phi} \left(\frac{\partial f}{\partial x} + \frac{\partial g}{\partial y} + \frac{\partial h}{\partial z} \right) dx \wedge dy \wedge dz = \iint_{\partial\Phi} f dy \wedge dz + g dz \wedge dx + h dx \wedge dy.$$

Introduce the notation

$$dV = dx \wedge dy \wedge dz,$$

and for a vector-valued mapping $\vec{F} : \mathbf{R}^3 \rightarrow \mathbf{R}^3$ define

$$\text{div } \vec{F} = D_1 F_1 + D_2 F_2 + D_3 F_3.$$

Then setting $\vec{F} = (f, g, h)$ shows that the Divergence Theorem is

$$\iiint_{\Phi} \text{div } \vec{F} dV = \iint_{\partial\Phi} \vec{F} \cdot \vec{d}\vec{n}$$

Thus the interpretation of the Divergence Theorem is a 3-dimensional version of the second interpretation of Green's Theorem,

*the net positive creation of fluid by \vec{F} throughout Φ
equals the net flux of \vec{F} outward through $\partial\Phi$.*

Again, the classical theorem views Φ and $\partial\Phi$ as sets, so long as whatever parameterization of $\partial\Phi$ is used to compute the right-side integral has the same orientation as the boundary of the parameterization of Φ used to compute the left-side integral.

Exercises

8.14.1. (a) Let $\gamma : [0, 1] \rightarrow \mathbf{R}^2$, $t \mapsto \gamma(t)$ be a curve, and recall the form-vectors on \mathbf{R}^2 $\vec{d}\vec{s} = (dx, dy)$, $\vec{d}\vec{n} = (dy, -dx)$. Compute the pullbacks $\gamma^*(\vec{d}\vec{s})$

and $\gamma^*(\vec{d}\mathbf{n})$ and explain why these are interpreted as differential tangent and normal vectors to γ .

(b) Let $\gamma : [0, 1] \rightarrow \mathbf{R}^3$, $t \mapsto \gamma(t)$ be a curve and $\Phi : [0, 1]^2 \rightarrow \mathbf{R}^3$, $(u, v) \mapsto \Phi(u, v)$ a surface, and recall the form-vectors on \mathbf{R}^3 $\vec{d}s = (dx, dy, dz)$, $\vec{d}\mathbf{n} = (dy \wedge dz, dz \wedge dx, dx \wedge dy)$. Compute the pullbacks $\gamma^*(\vec{d}s)$ and $\Phi^*(\vec{d}\mathbf{n})$ and explain why these are interpreted respectively as differential tangent vector to γ and differential normal vector to Φ .

8.14.2. Use Green's Theorem to show that for a planar region Φ ,

$$\text{area}(\Phi) = \int_{\partial\Phi} x \, dy = - \int_{\partial\Phi} y \, dx.$$

Thus one can measure the area of a planar set by traversing its boundary. (This principle was used to construct ingenious area-measuring machines before Green's Theorem was ever written down.)

8.14.3. Let H be the upper unit hemispherical shell,

$$H = \{(x, y, z) \in \mathbf{R}^3 : x^2 + y^2 + z^2 = 1, z \geq 0\}.$$

Define a vector-valued function on \mathbf{R}^3 ,

$$F(x, y, z) = (x + y + z, xy + yz + zx, xyz).$$

Use Stokes's Theorem to calculate $\iint_H \text{curl } F \cdot \vec{d}\mathbf{n}$.

8.14.4. Use the Divergence Theorem to evaluate

$$\int_{\partial\mathbf{H}} x^2 \, dy \wedge dz + y^2 \, dz \wedge dx + z^2 \, dx \wedge dy,$$

where $\partial\mathbf{H}$ is the boundary of the solid unit hemisphere

$$\mathbf{H} = \{(x, y, z) \in \mathbf{R}^3 : x^2 + y^2 + z^2 \leq 1, z \geq 0\}.$$

(Thus $\partial\mathbf{H}$ is the union of the unit disk in the (x, y) -plane and the unit upper hemispherical shell.) Feel free to cancel terms by citing symmetry if you're confident of what you're doing.

8.14.5. Let g and h be functions on \mathbf{R}^3 . Recall the operator $\nabla = (D_1, D_2, D_3)$, which takes scalar-valued functions to vector-valued functions. As usual, define the Laplacian operator to be $\Delta = D_{11} + D_{22} + D_{33}$. From an earlier exercise, $\Delta = \text{div} \circ \text{grad}$.

(a) Prove that $\text{div}(g \nabla h) = g \Delta h + \nabla g \cdot \nabla h$.

(b) If D is a closed compact subset of \mathbf{R}^3 with positively oriented boundary ∂D , prove that

$$\iiint_D (g \Delta h + \nabla g \cdot \nabla h) \, dV = \iint_{\partial D} g \nabla h \cdot \vec{d}\mathbf{n}.$$

(Here n is the unit outward normal to D and $\nabla h \cdot n$ is the directional derivative of h in the direction of n .) Interchange g and h and subtract the resulting formula from the first one to get

$$\iiint_D (g \Delta h - h \Delta g) dV = \iint_{\partial D} (g \nabla h - h \nabla g) \cdot d\vec{n}.$$

These two formulas are **Green's identities**.

(c) Assume that h is harmonic, meaning that it satisfies the harmonic equation $\Delta h = 0$.

Take $g = h$ and use Green's first identity to conclude that if $h = 0$ on the boundary ∂D then $h = 0$ on all of D .

Take $g = 1$ and use Green's second identity to show that

$$\iint_{\partial D} \nabla h \cdot d\vec{n} = 0.$$

What does this say about harmonic functions and flux?

8.15 Divergence and Curl in Polar Coordinates

The picture-explanations given in the previous section to interpret the divergence and the curl are not entirely satisfying. Working with the polar coordinate system further quantifies the ideas and makes them more coherent by applying to both operators in the same way.

Rather than study the divergence and the curl of a vector field \tilde{F} at a general point p , we may study the divergence and the curl of the modified vector field

$$F(x) = \tilde{F}(x + p) - \tilde{F}(p)$$

at the convenient particular point $\mathbf{0}$, at which the value of F is $\mathbf{0}$ as well. That is, we may normalize the point p to be $\mathbf{0}$ by prepending a translation of the domain, and we also may normalize $F(\mathbf{0})$ to $\mathbf{0}$ by postpending a translation of the range. With this in mind, let $A \subset \mathbf{R}^2$ be an open set that contains the origin, and let F be a continuous vector field on A that is stationary at the origin,

$$F = (F_1, F_2) : A \longrightarrow \mathbf{R}^2, \quad F(\mathbf{0}) = \mathbf{0}.$$

At any point other than the origin, F resolves into a radial component and an angular component. Specifically,

$$F = F_r + F_\theta,$$

where

$$\begin{aligned} F_r &= f_r \hat{r}, & f_r &= F \cdot \hat{r}, & \hat{r} &= (\cos \theta, \sin \theta) = (x, y)/|(x, y)|, \\ F_\theta &= f_\theta \hat{\theta}, & f_\theta &= F \cdot \hat{\theta}, & \hat{\theta} &= \hat{r}^\times = (-\sin \theta, \cos \theta) = (-y, x)/|(x, y)|. \end{aligned}$$

(Recall that the unary cross product $(x, y)^\times = (-y, x)$ in \mathbf{R}^2 rotates vectors 90 degrees counterclockwise.) Here f_r is positive if F_r points outward and negative if F_r points inward, and f_θ is positive if F_θ points counterclockwise and negative if F_θ points clockwise. Since $F(\mathbf{0}) = \mathbf{0}$, the resolution of F into radial and angular components extends continuously to the origin, $f_r(\mathbf{0}) = f_\theta(\mathbf{0}) = 0$, so that $F_r(\mathbf{0}) = F_\theta(\mathbf{0}) = \mathbf{0}$ even though \hat{r} and $\hat{\theta}$ are undefined at the origin.

The goal of this section is to express the divergence and the curl of F at the origin in terms of the polar coordinate system derivatives that seem naturally suited to describe them, the radial derivative of the (scalar) radial component of F ,

$$D_r f_r(\mathbf{0}) = \lim_{r \rightarrow 0^+} \frac{f_r(r \cos \theta, r \sin \theta)}{r},$$

and the radial derivative of the (scalar) angular component of F ,

$$D_r f_\theta(\mathbf{0}) = \lim_{r \rightarrow 0^+} \frac{f_\theta(r \cos \theta, r \sin \theta)}{r}.$$

However, matters aren't as simple here as one might hope. If the (vector) radial and angular components F_r and F_θ are differentiable at the origin then so is their sum F , but the converse is not true. So first we need sufficient conditions for the converse, i.e., sufficient conditions for the components to be differentiable. Necessary conditions are always easier to find, so Proposition 8.15.1 will do so, and then Proposition 8.15.2 will show that the necessary conditions are also sufficient. The conditions in question are the **Cauchy–Riemann equations**,

$$\begin{aligned} D_1 F_1(\mathbf{0}) &= D_2 F_2(\mathbf{0}), \\ D_1 F_2(\mathbf{0}) &= -D_2 F_1(\mathbf{0}). \end{aligned}$$

When the Cauchy–Riemann equations hold, we can describe the divergence and the curl of F at the origin in polar terms, as desired. This will be the content of Theorem 8.15.3.

Before we proceed to the details, a brief geometric discussion of the Cauchy–Riemann equations may be helpful. The equation $D_1 F_1 = D_2 F_2$ describes the left side of figure 8.17, in which the radial component of F on the horizontal axis is growing at the same rate as the radial component on the vertical axis. Similarly, the equation $D_2 F_1 = -D_1 F_2$ describes the right side of the figure, in which the angular component on the vertical axis is growing at the same rate as the angular component on the horizontal axis. Combined with differentiability at the origin, these two conditions will imply that moving outward in any direction, the radial component of F is growing at the same rate as it is on the axes, and similarly for the angular component. Thus the two limits that define the radial derivatives of the radial and angular components

of F at $\mathbf{0}$ (these were displayed in the previous paragraph) are independent of θ . An example of this situation, with radial and angular components both present, is shown in figure 8.18.




Figure 8.17. Geometry of the Cauchy–Riemann equations individually




Figure 8.18. Geometry of the Cauchy–Riemann equations together

As mentioned, the necessity of the Cauchy–Riemann equations is the natural starting point.

Proposition 8.15.1 (Polar Differentiability Implies Differentiability and the Cauchy–Riemann Equations). *Let $A \subset \mathbb{R}^2$ be an open set that contains the origin, and let F be a continuous vector field on A that is stationary at the origin,*

$$F = (F_1, F_2) : A \longrightarrow \mathbf{R}^2, \quad F(\mathbf{0}) = \mathbf{0}.$$

Assume that the radial and angular components F_r and F_θ of F are differentiable at the origin. Then F is differentiable at the origin, and the Cauchy–Riemann equations hold at the origin.

For example, the vector field $F(x, y) = (x, 0)$ is differentiable at the origin, but since $D_1 F_1(\mathbf{0}) = 1$ and $D_2 F_2(\mathbf{0}) = 0$, it does not satisfy the Cauchy–Riemann equations, and so the derivatives of the radial and angular components of F at the origin do not exist.

Proof. As already noted, the differentiability of F at the origin is immediate. To begin establishing the Cauchy–Riemann equations, consider the radial component of F ,

$$\begin{aligned} F_r(x, y) &= \begin{cases} f_r(x, y) \frac{(x, y)}{|(x, y)|} & \text{if } (x, y) \neq \mathbf{0}, \\ \mathbf{0} & \text{if } (x, y) = \mathbf{0} \end{cases} \\ &= \begin{cases} g_r(x, y)(x, y) & \text{if } (x, y) \neq \mathbf{0}, \\ \mathbf{0} & \text{if } (x, y) = \mathbf{0}, \end{cases} \end{aligned}$$

where

$$g_r(x, y) = \frac{f_r(x, y)}{|(x, y)|} \quad \text{for } (x, y) \neq \mathbf{0}.$$

The first component function of F_r vanishes at $\mathbf{0}$, i.e., $F_{r,1}(\mathbf{0}) = 0$. Also the first component function of F_r vanishes on the y -axis away from the origin because F_r is radial, and so $D_2 F_{r,1}(\mathbf{0}) = 0$ as well. Thus the condition that $F_{r,1}$ is differentiable at $\mathbf{0}$ is

$$\lim_{(h,k) \rightarrow \mathbf{0}} \frac{|F_{r,1}(h, k) - h D_1 F_{r,1}(\mathbf{0})|}{|(h, k)|} = 0,$$

But $F_{\theta,1}$ vanishes on the x -axis since F_θ is angular, and so $D_1 F_{\theta,1}(\mathbf{0}) = 0$, giving $D_1 F_{r,1}(\mathbf{0}) = D_1 F_1(\mathbf{0})$. Also $F_{r,1}(h, k) = h g_r(h, k)$ away from the origin, so that the previous condition becomes

$$\lim_{(h,k) \rightarrow \mathbf{0}} \frac{|h| |g_r(h, k) - D_1 F_1(\mathbf{0})|}{|(h, k)|} = 0,$$

A similar argument using the second component function shows that

$$\lim_{(h,k) \rightarrow \mathbf{0}} \frac{|k| |g_r(h, k) - D_2 F_2(\mathbf{0})|}{|(h, k)|} = 0.$$

Let $(h, k) \rightarrow \mathbf{0}$ along the line $h = k$ to see that

$$D_1 F_1(\mathbf{0}) = D_2 F_2(\mathbf{0})$$

since both are $\lim_{h \rightarrow 0} g_r(h, h)$. This is the first Cauchy–Riemann equation at the origin, but we show a bit more, to be used in the proof of Theorem 8.15.3. Add the previous two displayed limits to get

$$\lim_{(h,k) \rightarrow \mathbf{0}} \frac{(|h| + |k|) |g_r(h, k) - D_1 F_r(\mathbf{0})|}{|(h, k)|} = 0,$$

and since $|(h, k)| \leq |h| + |k|$, it follows that

$$\lim_{(h,k) \rightarrow \mathbf{0}} |g_r(h, k) - D_1 F_r(\mathbf{0})| = 0.$$

That is,

$$\lim_{(h,k) \rightarrow \mathbf{0}} g_r(h, k) = D_1 F_r(\mathbf{0}) = D_2 F_{r,2}(\mathbf{0}).$$

Next consider the radial component of the vector field $-F^\times = f_\theta \hat{r} - f_r \hat{\theta}$,

$$(-F^\times)_r(x, y) = \begin{cases} g_\theta(x, y)(x, y) & \text{if } (x, y) \neq \mathbf{0}, \\ \mathbf{0} & \text{if } (x, y) = \mathbf{0}, \end{cases}$$

where

$$g_\theta(x, y) = \frac{f_\theta(x, y)}{|(x, y)|} \quad \text{for } (x, y) \neq \mathbf{0}.$$

This radial component is differentiable at the origin since it is a rotation of the angular component of the original F , so as just argued,

$$\lim_{(x,y) \rightarrow \mathbf{0}} g_\theta(x, y) = -D_1 F_1^\times(\mathbf{0}) = -D_2 F_2^\times(\mathbf{0}).$$

But

$$-D_1 F_1^\times(\mathbf{0}) = D_1 F_2(\mathbf{0}) \quad \text{and} \quad -D_2 F_2^\times(\mathbf{0}) = -D_2 F_1(\mathbf{0}),$$

and this gives the second Cauchy–Riemann equation at the origin. \square

Also as mentioned, the converse to Proposition 8.15.1 holds too.

Proposition 8.15.2 (Differentiability and the Cauchy–Riemann Equations Imply Polar Differentiability). *Let $A \subset \mathbf{R}^2$ be an open set that contains the origin, and let F be a continuous vector field on A that is stationary at the origin,*

$$F = (F_1, F_2) : A \longrightarrow \mathbf{R}^2, \quad F(\mathbf{0}) = \mathbf{0}.$$

Assume that F is differentiable at the origin, and assume that the Cauchy–Riemann equations hold at the origin. Then the radial and angular components F_r and F_θ are differentiable at the origin.

Proof. Let $a = D_1F_1(\mathbf{0})$ and let $b = D_1F_2(\mathbf{0})$. By the Cauchy–Riemann equations, also $a = D_2F_2(\mathbf{0})$ and $b = -D_2F_1(\mathbf{0})$, so that the Jacobian matrix of F at $\mathbf{0}$ is

$$F'(\mathbf{0}) = \begin{bmatrix} a & -b \\ b & a \end{bmatrix}.$$

The condition that F is differentiable at $\mathbf{0}$ is

$$\lim_{(h,k) \rightarrow \mathbf{0}} \frac{|F(h,k) - (ah - bk, bh + ak)|}{|(h,k)|} = 0.$$

Decompose the quantity whose absolute value is the numerator into radial and angular components,

$$F(h,k) - (ah - bk, bh + ak) = (F_r(h,k) - (ah, ak)) + (F_\theta(h,k) - (-bk, bh))$$

Since the direction vectors $\hat{r} = (h,k)/|(h,k)|$ and $\hat{\theta} = (-k,h)/|(h,k)|$ are orthogonal, and

$$F_r(h,k) - (ah, ak) \parallel \hat{r} \quad \text{and} \quad F_\theta(h,k) - (-bk, bh) \parallel \hat{\theta},$$

it follows that

$$|F_r(h,k) - (ah, ak)| \leq |F(h,k) - (ah - bk, bh + ak)|$$

and

$$|F_\theta(h,k) - (-bk, bh)| \leq |F(h,k) - (ah - bk, bh + ak)|.$$

Therefore,

$$\lim_{(h,k) \rightarrow \mathbf{0}} \frac{|F_r(h,k) - (ah, ak)|}{|(h,k)|} = 0$$

and

$$\lim_{(h,k) \rightarrow \mathbf{0}} \frac{|F_\theta(h,k) - (-bk, bh)|}{|(h,k)|} = 0.$$

That is, F_r and F_θ are differentiable at the origin with respective Jacobian matrices

$$F'_r(\mathbf{0}) = \begin{bmatrix} a & 0 \\ 0 & a \end{bmatrix} \quad \text{and} \quad F'_\theta(\mathbf{0}) = \begin{bmatrix} 0 & -b \\ b & 0 \end{bmatrix}.$$

This completes the proof. \square

Now we can return to the divergence and the curl.

Theorem 8.15.3 (Divergence and Curl in Polar Coordinates). *Let $A \subset \mathbf{R}^2$ be a region of \mathbf{R}^2 containing the origin, and let F be a continuous vector field on A that is stationary at the origin,*

$$F = (F_1, F_2) : A \longrightarrow \mathbf{R}^2, \quad F(\mathbf{0}) = \mathbf{0}.$$

Assume that F is differentiable at the origin and that the Cauchy–Riemann equations hold at the origin. Then the radial derivatives of the radial and angular components of F at the origin,

$$D_r f_r(\mathbf{0}) = \lim_{r \rightarrow 0^+} \frac{f_r(r \cos \theta, r \sin \theta)}{r}$$

and

$$D_r f_\theta(\mathbf{0}) = \lim_{r \rightarrow 0^+} \frac{f_\theta(r \cos \theta, r \sin \theta)}{r},$$

both exist independently of how θ behaves as r shrinks to 0. Furthermore, the divergence of F at the origin is twice the radial derivative of the radial component,

$$(\operatorname{div} F)(\mathbf{0}) = 2D_r f_r(\mathbf{0}),$$

and the curl of F at the origin is twice the radial derivative of the angular component,

$$(\operatorname{curl} F)(\mathbf{0}) = 2D_r f_\theta(\mathbf{0}).$$

Proof. By Proposition 8.15.2, the angular and radial components of F are differentiable at the origin, so that the hypotheses of Proposition 8.15.1 are met. The first limit in the statement of the theorem was calculated in the proof of Proposition 8.15.1,

$$D_r f_r(\mathbf{0}) = \lim_{(x,y) \rightarrow \mathbf{0}} \frac{f_r(x,y)}{|(x,y)|} = \lim_{(x,y) \rightarrow \mathbf{0}} g_r(x,y) = D_1 F_1(\mathbf{0}) = D_2 F_2(\mathbf{0}).$$

This makes the formula for the divergence immediate,

$$(\operatorname{div} F)(\mathbf{0}) = D_1 F_1(\mathbf{0}) + D_2 F_2(\mathbf{0}) = 2D_r f_r(\mathbf{0}).$$

Similarly,

$$D_r f_\theta(\mathbf{0}) = \lim_{(x,y) \rightarrow \mathbf{0}} \frac{f_\theta(x,y)}{|(x,y)|} = \lim_{(x,y) \rightarrow \mathbf{0}} g_\theta(x,y) = D_1 F_2(\mathbf{0}) = -D_2 F_1(\mathbf{0}),$$

so that

$$(\operatorname{curl} F)(\mathbf{0}) = D_1 F_2(\mathbf{0}) - D_2 F_1(\mathbf{0}) = 2D_r f_\theta(\mathbf{0}).$$

□

If F is a velocity field then the limit in the formula

$$(\operatorname{curl} F)(\mathbf{0}) = 2 \lim_{r \rightarrow 0^+} \frac{f_\theta(r \cos \theta, r \sin \theta)}{r}$$

has the interpretation of the angular velocity of F at the origin. That is:

When the Cauchy–Riemann equations hold, the curl is twice the angular velocity.

Indeed, the angular velocity ω away from the origin is by definition the rate of increase of the polar angle θ with the motion of F . This is not the counter-clockwise component f_θ , but rather $\omega = f_\theta/r$, i.e., ω is the function called g_θ in the proof of Proposition 8.15.1. To understand this, think of a uniformly spinning disk such as a record on a turntable. At each point except the center, the angular velocity is the same. But the speed of motion is not constant over the disk, it is the angular velocity times the distance from the center. That is, the angular velocity is the speed divided by the radius, as claimed. In these terms, the proof showed that the angular velocity ω extends continuously to $\mathbf{0}$, and that $(\text{curl } F)(\mathbf{0})$ is twice the extended value $\omega(\mathbf{0})$.

Also, if F is a velocity field then the right side of the formula

$$(\text{div } F)(\mathbf{0}) = 2 \lim_{r \rightarrow 0^+} \frac{f_r(r \cos \theta, r \sin \theta)}{r}$$

has the interpretation of the flux density of F at the origin. That is:

When the Cauchy–Riemann equations hold, the divergence is the flux density.

To understand this, think of a planar region of incompressible fluid about the origin, and let r be a positive number small enough that the fluid fills the area inside the circle of radius r . Suppose that new fluid being added throughout the interior of the circle, at rate c per unit of area. Thus fluid is being added to the area inside the circle at total rate $\pi r^2 c$. Here c is called the flux density over the circle and it is measured in reciprocal time units, while the units of $\pi r^2 c$ are area over time. Since the fluid is incompressible, $\pi r^2 c$ is also the rate at which fluid is passing normally outward through the circle. And since the circle has circumference $2\pi r$, fluid is therefore passing normally outward through each point of the circle with radial velocity

$$f_r(r \cos \theta, r \sin \theta) = \frac{\pi r^2 c}{2\pi r} = \frac{rc}{2}.$$

Consequently,

$$2 \frac{f_r(r \cos \theta, r \sin \theta)}{r} = c.$$

Now let r shrink to 0. The left side of the display goes to the divergence of F at $\mathbf{0}$, and the right side becomes the continuous extension to radius 0 of the flux density over the circle of radius r . That is, the divergence is the flux density when fluid is being added at a single point.

Exercises

8.15.1. Put \mathbf{R}^2 into correspondence with the complex number field \mathbf{C} as follows:

$$\begin{bmatrix} x \\ y \end{bmatrix} \longleftrightarrow x + iy.$$

Show that the correspondence extends to

$$\begin{bmatrix} a & -b \\ b & a \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix} \longleftrightarrow (a + ib)(x + iy).$$

Show also that the correspondence preserves absolute value, i.e.,

$$\left| \begin{bmatrix} x \\ y \end{bmatrix} \right| = |x + iy|,$$

where the first absolute value is on \mathbf{R}^2 and the second one on \mathbf{C} .

8.15.2. Let $A \subset \mathbf{R}^2$ be an open set that contains the origin, and let $F : A \rightarrow \mathbf{R}^2$ be a vector field on A that is stationary at the origin. Define a complex-valued function of a complex variable corresponding to F ,

$$f(x + iy) = F_1(x, y) + iF_2(x, y), \quad (x, y) \in A.$$

Then f is called **complex-differentiable at 0** if the following limit exists:

$$\lim_{\Delta z \rightarrow 0} \frac{f(z + \Delta z) - f(z)}{\Delta z}.$$

The limit is denoted $f'(z)$.

(a) Suppose that f is complex-differentiable at 0. Compute $f'(z)$ first by letting Δz go to 0 along the x -axis, and again by letting Δz go to 0 along the y -axis. Explain how your calculation shows that the Cauchy–Riemann equations hold at 0.

(b) Show also that if f is complex differentiable at 0 then F is vector differentiable at $\mathbf{0}$, meaning differentiable in the usual sense. Suppose that f is complex-differentiable at 0, and that $f'(0) = re^{i\theta}$. Show that

$$(\operatorname{div} F)(\mathbf{0}) = 2r \cos \theta, \quad (\operatorname{curl} F)(\mathbf{0}) = 2r \sin \theta.$$

(c) Suppose that F is vector-differentiable at $\mathbf{0}$ and that the Cauchy–Riemann equations hold at $\mathbf{0}$. Show that f is complex-differentiable at 0.

8.16 Summary

The bulk of the ideas introduced in this chapter are algebraic. Even so, the General FTIC eases the proof of the classical Change of Variable Theorem, and it subsumes the three classical integration theorems of vector calculus.

Index

- absolute value of a vector, 25
- absolute value on $C_c^k(\mathbf{R}^n)$
 - properties of, 300
- addition of matrices, 65
- affine mapping, 64
- annulus, 267
- Arithmetic–Geometric Mean Inequality, 203
- ball, 42
 - n -dimensional of radius r , 277
 - closed, 171
- basis of \mathbf{R}^n , 22
 - negative, 108
 - positive, 108
 - standard, 21
- Bernoulli's Inequality, 3
- Bolzano–Weierstrass Property
 - in \mathbf{R} , 45
 - in \mathbf{R}^n , 45
- boundary
 - of a closed ball, 171
- boundary mapping, 374
- boundary point of a set, 281
- bounded
 - function, 216, 220
 - mapping, 220
 - set, 44
- box
 - compact, 219
 - face of, 99
 - unit, 99
- cartesian product, 219
- Cauchy–Riemann equations, 390
- Cauchy–Schwarz Inequality, 26
- centroid, 260
- chain, 372
- Chain Rule, 128
 - in coordinates, 141
- change of scale principle, 277
- change of variable mapping
 - for cylindrical coordinates, 270
 - for polar coordinates, 266
 - for spherical coordinates, 272
- Change of Variable Theorem
 - for differential forms, 370
 - for multiple integrals, 269
 - for one-variable integrals, 240
- characteristic function of a set, 244
- characteristic polynomial, 158
- closed set, 43
- closure of a set, 281
- common refinement of two partitions, 223
- compact set, 46
- complement of a set, 170
- completeness of the real number system,
 - 2
 - as a binary search criterion, 2
 - as a monotonic sequence criterion, 2
 - as a set-bound criterion, 2
- component of a vector
 - normal, 32
 - parallel, 32
- componentwise nature of linearity, 57
- connected set, 47

- continuity
 - at a point, 36
 - componentwise nature of, 38
 - ε - δ at a point, 230
 - ε - δ on a set, 230
 - linearity of, 37
 - on a set, 36
 - pointwise, 233
 - preservation under composition, 37
 - sequential at a point, 230
 - sequential on a set, 230
 - uniform on a set, 233
- convergence
 - componentwise nature of, 36
 - linearity of, 36
- Critical Point Theorem, 151
- cross product
 - definition of, 112
 - in \mathbf{R}^n , 335
 - properties of, 113
- cube
 - singular, 372
 - standard, 372
- curl of a vector-valued mapping, 359
 - in polar coordinates, 394
- cylindrical coordinates, 270
- derivative
 - linearity of the, 128
 - of a differential form, 355
 - of a mapping, 122
 - uniqueness of, 125
- determinant, 79, 89
 - of a basis of \mathbf{R}^n , 108
 - properties of, 91
- differentiability
 - componentwise nature of, 126
- differential form, 337
 - closed, 360
 - derivative of, 355
 - exact, 360
 - pullback of, 363
 - standard presentation of, 340
- differential forms
 - integration of, 339
 - meaning of integration of, 348
 - Product Rule for, 356
 - syntax of, 338
 - wedge product of, 353
- differentiation under the integral sign, 262
- directional derivative, 159
- discrete set, 49
- distance in \mathbf{R}^n , 29
 - properties of, 30
- divergence of a vector-valued mapping, 359
 - in polar coordinates, 394
- Divergence Theorem, 387
 - geometric interpretation of, 387
- Euclidean space, 17
- Euclid's Least Area Problem, 199
- Euler's identity for homogeneous functions, 360
- Extreme Value Theorem, 3, 47
- field axioms, 1
- flow integral
 - 2-dimensional
 - coordinate, 332
 - intrinsic, 332
 - 3-dimensional
 - coordinate, 333
 - general, 334
- flux integral
 - 2-dimensional
 - coordinate, 333
 - intrinsic, 333
 - 3-dimensional
 - coordinate, 334
 - intrinsic, 333
 - general
 - coordinate, 335
 - intrinsic, 334
- Fubini's Theorem, 253
- Fundamental Theorem of Integral Calculus, 4
 - Generalized, 379
 - one-variable
 - first version, 239
 - second version, 240
- gamma function, 278
- Gauss's Theorem, 387
- gradient, 160
- gradient of a scalar-valued function, 359
- Gram–Schmidt process, 33

- graph of a function, 246
- Green's identities, 389
- Green's Theorem, 384
 - geometric interpretations of, 385
- harmonic equation, 360
- Hölder's Inequality, 205
- homogeneous function, 144, 360
- identity matrix, 67
- Implicit Function Theorem, 190
- inner product
 - of two vectors, 24
 - properties of, 24
- Integrability Criterion, 226
- integral
 - iterated, 252
 - lower, of a function over a box, 225
 - of a differential form over a chain, 372
 - of a differential form over a surface, 339
 - of a function over a box, 225
 - of a function over a surface, 329
 - upper, of a function over a box, 225
- integral curve, 165
- interior point of a set, 122
- Intermediate Value Theorem, 3
- interval
 - compact, 215
- Inverse Function Theorem, 174
- inverse image characterization of
 - continuity, 172
- inverse image of a mapping, 171
- isolated point of a set, 43
- Jacobian matrix, 122
- Known-Integrable Function, 250
- Lagrange Multiplier Condition, 205
- Laplace's equation, 148, 360
- Laplacian operator, 388
- length
 - of an interval, 216
- level set
 - of a function, 161
 - of a mapping, 181
- limit of a function, 49
- limit point of a set, 42
- line in \mathbf{R}^3 , 115
- linear combination, 22
- linear equation
 - homogeneous, 78
 - inhomogeneous, 78
- Linear Invertibility Theorem, 84
- linear mapping, 54
- lower sum, 217, 220
- mapping
 - affine, 64
 - component functions of, 37
 - from \mathbf{R}^n to \mathbf{R}^m , 33
 - from A to \mathbf{R}^m , 34
 - linear, 54
 - null, 50
- matrix, 59
 - echelon, 75
 - elementary, 71
 - Hessian, 151
 - indefinite, 154
 - lower triangular, 78
 - negative definite, 154
 - new column of, 75
 - old column of, 75
 - orthogonal, 85
 - positive definite, 154
 - recombine, 71
 - scale, 71
 - skew-symmetric, 85
 - transpose of, 69
 - transposition, 71
 - triangular, 93
- matrix inversion algorithm, 76
- matrix multiplication
 - properties of, 67
- Max/min Test
 - general, 159
 - two-variable, 156
- Mean Value Theorem, 4
- modulus
 - of a vector, 25
 - properties of, 25
- multilinear function, 134
- multiplication
 - matrix-by-matrix, 66
 - matrix-by-vector, 59
 - scalar-by-matrix, 65

- nilpotence of d , 358
- norm of a linear mapping, 64
- open set, 170
- order axioms, 2
- ordered k -tuple from $\{1, \dots, n\}$, 337
 - increasing, 340
- orthogonal vectors, 30
- Pappas's Theorem, 277
- parallelepiped, 101
 - face of, 101
- parameter domain of a surface, 324
- partial derivative, 135
 - mixed, 146
 - second order, 146
- partition
 - of a box, 219
 - of an interval, 216
- path-connected set, 49
- permutation, 88
 - inversion in, 88
 - sign of, 88
- Persistence of Inequality, 40
- plane in \mathbf{R}^3 , 115
- polar coordinates, 265
- Product Rule
 - for differential forms, 356
- product rule
 - multivariable, 131
- pullback
 - contravariance of, 367
 - properties of, 365
- pullback of a differential form, 363
- Pullback Theorem, 370
- Pullback–Determinant Theorem, 365
- quotient rule
 - multivariable, 131
- real number system, 1
- refinement of a partition, 221
- reparameterization of a surface, 330
 - orientation-preserving, 350
 - orientation-reversing, 350
- saddle point, 152
- scalar multiplication, 19
- sequence, 35
 - convergence of, 35
 - limit of, 35
 - null, 35
- sequential characterization
 - of bounded sets, 45
 - of closed sets, 43
 - of compact sets, 46
 - of limit points, 43
- simplex, 378
 - n -dimensional of side r , 264
 - unit in \mathbf{R}^3 , 260
- Size Bounds for a vector, 29
- smooth mapping, 324
- Snell's Law, 201
- spherical coordinates, 271
- Stokes's Theorem, 386
 - geometric interpretation of, 387
- subbox of a partition, 219
- subinterval of a partition, 216
- subsequence, 44
- substitution
 - forward, in one-variable integrals, 241
 - reverse, in one-variable integrals, 241
- surface, 324
- Taylor polynomial, 6
- Taylor remainder, 7
 - integral form, 16
 - Lagrange form, 10
- Taylor's Theorem, 10
- The Continuous Image of a Compact Set is Compact, 46
- topological property, 47
- Triangle Inequality, 28
 - full, 28
- upper sum, 217, 220
- Vandermonde matrix, 96
- vector addition, 18
- vector space axioms, 20
- vector space over a field, 21
- volume
 - k -dimensional in \mathbf{R}^n
 - of a parallelepiped, 327
 - of a surface, 329
 - of a box, 219
 - of a set, 244
- volume form, 350
- Volume Zero Criterion, 245

wedge product, 353
properties of, 354

skew symmetry of, 354