ANALYSIS ON MANIFOLDS

James R. Munkres

Analysis on Manifolds

James R. Munkres

Massachusetts Institute of Technology
Cambridge, Massachusetts

ADVANCED BOOK PROGRAM

ADDISON-WESLEY PUBLISHING COMPANY
The Advanced Book Program
Redwood City, California • Menlo Park, California • Reading, Massachusetts
New York • Don Mills, Ontario • Wokingham, United Kingdom • Amsterdam
Bonn • Sydney • Singapore • Tokyo • Madrid • San Juan

Publisher: *Allan M. Wylde*
Production Manager: *Jan V. Benes*
Marketing Manager: *Laura Likely*
Electronic Composition: *Peter Vacek*
Cover Design: *Iva Frank*

Library of Congress Cataloging-in-Publication Data

Munkres, James R., 1930–
Analysis on manifolds/James R. Munkres.
p. cm.
Includes bibliographical references.
1. Mathematical analysis. 2. Manifolds (Mathematics)
QA300.M75 1990
516.3'6'20–dc20 91-39786
ISBN 0-201-51035-9 CIP

This book was prepared using the TEX typesetting language.

ABCDEFGHIJ-MA-943210

Analysis on Manifolds

Preface

This book is intended as a text for a course in analysis, at the senior or first-year graduate level.

A year-long course in real analysis is an essential part of the preparation of any potential mathematician. For the first half of such a course, there is substantial agreement as to what the syllabus should be. Standard topics include: sequence and series, the topology of metric spaces, and the derivative and the Riemannian integral for functions of a single variable. There are a number of excellent texts for such a course, including books by Apostol [A], Rudin [Ru], Goldberg [Go], and Royden [Ro], among others.

There is no such universal agreement as to what the syllabus of the second half of such a course should be. Part of the problem is that there are simply too many topics that belong in such a course for one to be able to treat them all within the confines of a single semester, at more than a superficial level.

At M.I.T., we have dealt with the problem by offering two independent second-term courses in analysis. One of these deals with the derivative and the Riemannian integral for functions of several variables, followed by a treatment of differential forms and a proof of Stokes' theorem for manifolds in euclidean space. The present book has resulted from my years of teaching this course. The other deals with the Lebesque integral in euclidean space and its applications to Fourier analysis.

Prequisites

As indicated, we assume the reader has completed a one-term course in analysis that included a study of metric spaces and of functions of a single variable. We also assume the reader has some background in linear algebra, including vector spaces and linear transformations, matrix algebra, and determinants.

The first chapter of the book is devoted to reviewing the basic results from linear algebra and analysis that we shall need. Results that are truly basic are stated without proof, but proofs are provided for those that are sometimes omitted in a first course. The student may determine from a perusal of this chapter whether his or her background is sufficient for the rest of the book.

How much time the instructor will wish to spend on this chapter will depend on the experience and preparation of the students. I usually assign Sections 1 and 3 as reading material, and discuss the remainder in class.

How the book is organized

The main part of the book falls into two parts. The first, consisting of Chapter 2 through 4, covers material that is fairly standard: derivatives, the inverse function theorem, the Riemann integral, and the change of variables theorem for multiple integrals. The second part of the book is a bit more sophisticated. It introduces manifolds and differential forms in $\mathbf{R}^n$, providing the framework for proofs of the n-dimensional version of Stokes' theorem and of the Poincaré lemma.

A final chapter is devoted to a discussion of abstract manifolds; it is intended as a transition to more advanced texts on the subject.

The dependence among the chapters of the book is expressed in the following diagram:

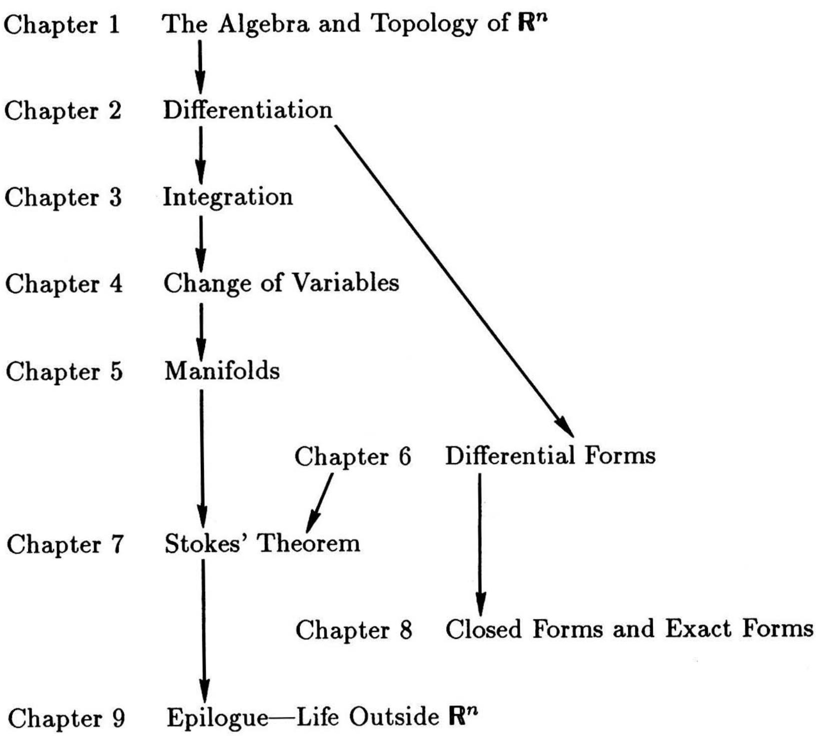

Certain sections of the books are marked with an asterisk; these sections may be omitted without loss of continuity. Similarly, certain theorems that may be omitted are marked with asterisks. When I use the book in our undergraduate analysis sequence, I usually omit Chapter 8, and assign Chapter 9 as reading. With graduate students, it should be possible to cover the entire book.

At the end of each section is a set of exercises. Some are computational in nature; students find it illuminating to know that one can compute the volume of a five-dimensional ball, even if the practical applications are limited! Other exercises are theoretical in nature, requiring that the student analyze carefully the theorems and proofs of the preceding section. The more difficult exercises are marked with asterisks, but none is unreasonably hard.

Acknowledgements

Two pioneering works in this subject demonstrated that such topics as manifolds and differential forms could be discussed with undergraduates. One is the set of notes used at Princeton c. 1960, written by Nickerson, Spencer, and Steenrod [N-S-S]. The second is the book by Spivak [S]. Our indebtedness to these sources is obvious. A more recent book on these topics is the one by Guillemin and Pollack [G-P]. A number of texts treat this material at a more advanced level. They include books by Boothby [B], Abraham, Mardsen, and Raitu [A-M-R], Berger and Gostiaux [B-G], and Fleming [F]. Any of them would be suitable reading for the student who wishes to pursue these topics further.

I am indebted to Sigurdur Helgason and Andrew Browder for helpful comments. To Ms. Viola Wiley go my thanks for typing the original set of lecture notes on which the book is based. Finally, thanks is due to my students at M.I.T., who endured my struggles with this material, as I tried to learn how to make it understandable (and palatable) to them!

J.R.M.

Contents

1

The Algebra and Topology of $\mathbb{R}^n$

§1. REVIEW OF LINEAR ALGEBRA

Vector spaces

Suppose one is given a set V of objects, called **vectors**. And suppose there is given an operation called **vector addition**, such that the sum of the vectors $\mathbf{x}$ and $\mathbf{y}$ is a vector denoted $\mathbf{x}+\mathbf{y}$. Finally, suppose there is given an operation called **scalar multiplication**, such that the product of the scalar (i.e., real number) c and the vector $\mathbf{x}$ is a vector denoted $c\mathbf{x}$.

The set V, together with these two operations, is called a **vector space** (or **linear space**) if the following properties hold for all vectors $\mathbf{x}$, $\mathbf{y}$, $\mathbf{z}$ and all scalars c, d:

(1) $\mathbf{x}+\mathbf{y}=\mathbf{y}+\mathbf{x}$.

(2) $\mathbf{x}+(\mathbf{y}+\mathbf{z})=(\mathbf{x}+\mathbf{y})+\mathbf{z}$.

(3) There is a unique vector $\mathbf{0}$ such that $\mathbf{x}+\mathbf{0}=\mathbf{x}$ for all $\mathbf{x}$.

(4) $\mathbf{x}+(-1)\mathbf{x}=\mathbf{0}$.

(5) $1\mathbf{x}=\mathbf{x}$.

(6) $c(d\mathbf{x})=(cd)\mathbf{x}$.

(7) $(c+d)\mathbf{x}=c\mathbf{x}+d\mathbf{x}$.

(8) $c(\mathbf{x}+\mathbf{y})=c\mathbf{x}+c\mathbf{y}$.

One example of a vector space is the set $\mathbb{R}^n$ of all n-tuples of real numbers, with component-wise addition and multiplication by scalars. That is, if $\mathbf{x} = (x_1, \ldots, x_n)$ and $\mathbf{y} = (y_1, \ldots, y_n)$, then

$$\mathbf{x} + \mathbf{y} = (x_1 + y_1, \ldots, x_n + y_n),$$
$$c\mathbf{x} = (cx_1, \ldots, cx_n).$$

The vector space properties are easy to check.

If V is a vector space, then a subset W of V is called a **linear subspace** (or simply, a **subspace**) of V if for every pair $\mathbf{x}$,$\mathbf{y}$ of elements of W and every scalar c, the vectors $\mathbf{x} + \mathbf{y}$ and $c\mathbf{x}$ belong to W. In this case, W itself satisfies properties (1)–(8) if we use the operations that W inherits from V, so that W is a vector space in its own right.

In the first part of this book, $\mathbb{R}^n$ and its subspaces are the only vector spaces with which we shall be concerned. In later chapters we shall deal with more general vector spaces.

Let V be a vector space. A set $\mathbf{a}_1, \ldots, \mathbf{a}_m$ of vectors in V is said to **span** V if to each $\mathbf{x}$ in V, there corresponds *at least* one m-tuple of scalars $c_1, \ldots, c_m$ such that

$$\mathbf{x} = c_1\mathbf{a}_1 + \cdots + c_m\mathbf{a}_m.$$

In this case, we say that $\mathbf{x}$ can be written as a **linear combination** of the vectors $\mathbf{a}_1, \ldots, \mathbf{a}_m$.

The set $\mathbf{a}_1, \ldots, \mathbf{a}_m$ of vectors is said to be **independent** if to each $\mathbf{x}$ in V there corresponds *at most* one m-tuple of scalars $c_1, \ldots, c_m$ such that

$$\mathbf{x} = c_1\mathbf{a}_1 + \cdots + c_m\mathbf{a}_m.$$

Equivalently, $\{\mathbf{a}_1, \ldots, \mathbf{a}_m\}$ is independent if to the zero vector $\mathbf{0}$ there corresponds only one m-tuple of scalars $d_1, \ldots, d_m$ such that

$$\mathbf{0} = d_1\mathbf{a}_1 + \cdots + d_m\mathbf{a}_m,$$

namely the scalars $d_1 = d_2 = \cdots = d_m = 0$.

If the set of vectors $\mathbf{a}_1, \ldots, \mathbf{a}_m$ both spans V and is independent, it is said to be a **basis** for V.

One has the following result:

Theorem 1.1. *Suppose V has a basis consisting of m vectors. Then any set of vectors that spans V has at least m vectors, and any set of vectors of V that is independent has at most m vectors. In particular, any basis for V has exactly m vectors.* □

If V has a basis consisting of m vectors, we say that m is the **dimension** of V. We make the convention that the vector space consisting of the zero vector alone has dimension zero.

It is easy to see that $\mathbb{R}^n$ has dimension n. (Surprise!) The following set of vectors is called the **standard basis** for $\mathbb{R}^n$:

$$\begin{aligned}\mathbf{e}_1 &= (1,0,0,\dots,0),\\ \mathbf{e}_2 &= (0,1,0,\dots,0),\\ &\dots\\ \mathbf{e}_n &= (0,0,0,\dots,1).\end{aligned}$$

The vector space $\mathbb{R}^n$ has many other bases, but any basis for $\mathbb{R}^n$ must consist of precisely n vectors.

One can extend the definitions of *spanning, independence,* and *basis* to allow for infinite sets of vectors; then it is possible for a vector space to have an infinite basis. (See the exercises.) However, we shall not be concerned with this situation.

Because $\mathbb{R}^n$ has a finite basis, so does every subspace of $\mathbb{R}^n$. This fact is a consequence of the following theorem:

Theorem 1.2. *Let V be a vector space of dimension m. If W is a linear subspace of V (different from V), then W has dimension less than m. Furthermore, any basis $\mathbf{a}_1,\dots,\mathbf{a}_k$ for W may be extended to a basis $\mathbf{a}_1,\dots,\mathbf{a}_k$, $\mathbf{a}_{k+1},\dots,\mathbf{a}_m$ for V.* □

Inner products

If V is a vector space, an **inner product** on V is a function assigning, to each pair $\mathbf{x}$, $\mathbf{y}$ of vectors of V, a real number denoted $\langle\mathbf{x},\mathbf{y}\rangle$, such that the following properties hold for all $\mathbf{x}$, $\mathbf{y}$, $\mathbf{z}$ in V and all scalars c:

(1) $\langle\mathbf{x},\mathbf{y}\rangle = \langle\mathbf{y},\mathbf{x}\rangle$.

(2) $\langle\mathbf{x}+\mathbf{y},\mathbf{z}\rangle = \langle\mathbf{x},\mathbf{z}\rangle + \langle\mathbf{y},\mathbf{z}\rangle$.

(3) $\langle c\mathbf{x},\mathbf{y}\rangle = c\langle\mathbf{x},\mathbf{y}\rangle = \langle\mathbf{x},c\mathbf{y}\rangle$.

(4) $\langle\mathbf{x},\mathbf{x}\rangle > 0$ if $\mathbf{x} \neq \mathbf{0}$.

A vector space V together with an inner product on V is called an **inner product space.**

A given vector space may have many different inner products. One particularly useful inner product on $\mathbb{R}^n$ is defined as follows: If $\mathbf{x} = (x_1,\dots,x_n)$ and $\mathbf{y} = (y_1,\dots,y_n)$, we define

$$\langle\mathbf{x},\mathbf{y}\rangle = x_1y_1 + \cdots + x_ny_n.$$

The properties of an inner product are easy to verify. This is the inner product we shall commonly use in $\mathbb{R}^n$. It is sometimes called the **dot product**; we denote it by $\langle\mathbf{x},\mathbf{y}\rangle$ rather than $\mathbf{x}\cdot\mathbf{y}$ to avoid confusion with the matrix product, which we shall define shortly.

If V is an inner product space, one defines the **length** (or **norm**) of a vector of V by the equation

$$\|\mathbf{x}\| = \langle \mathbf{x}, \mathbf{y} \rangle^{1/2}.$$

The norm function has the following properties:

(1) $\|\mathbf{x}\| > 0$ if $\mathbf{x} \neq \mathbf{0}$.

(2) $\|c\mathbf{x}\| = |c|\,\|\mathbf{x}\|$.

(3) $\|\mathbf{x}+\mathbf{y}\| \leq \|\mathbf{x}\| + \|\mathbf{y}\|$.

The third of these properties is the only one whose proof requires some work; it is called the **triangle inequality**. (See the exercises.) An equivalent form of this inequality, which we shall frequently find useful, is the inequality

(3′) $\|\mathbf{x}-\mathbf{y}\| \geq \|\mathbf{x}\| - \|\mathbf{y}\|$.

Any function from V to the reals $\mathbb{R}$ that satisfies properties (1)–(3) just listed is called a **norm** on V. The length function derived from an inner product is one example of a norm, but there are other norms that are not derived from inner products. On $\mathbb{R}^n$, for example, one has not only the familiar norm derived from the dot product, which is called the **euclidean norm,** but one has also the **sup norm,** which is defined by the equation

$$|\mathbf{x}| = \max\{|x_1|, \ldots, |x_n|\}.$$

The sup norm is often more convenient to use than the euclidean norm. We note that these two norms on $\mathbb{R}^n$ satisfy the inequalities

$$|\mathbf{x}| \leq \|\mathbf{x}\| \leq \sqrt{n}|\mathbf{x}|.$$

Matrices

A **matrix** A is a rectangular array of numbers. The general number appearing in the array is called an **entry** of A. If the array has n rows and m columns, we say that A has **size** n by m, or that A is "an n by m matrix." We usually denote the entry of A appearing in the i^{th} row and j^{th} column by a_{ij}; we call i the **row index** and j the **column index** of this entry.

If A and B are matrices of size n by m, with general entries a_{ij} and b_{ij}, respectively, we define $A + B$ to be the n by m matrix whose general entry is $a_{ij} + b_{ij}$, and we define cA to be the n by m matrix whose general entry is ca_{ij}. With these operations, the set of all n by m matrices is a vector space; the eight vector space properties are easy to verify. This fact is hardly surprising, for an n by m matrix is very much like an nm-tuple; the only difference is that the numbers are written in a rectangular array instead of a linear array.

The set of matrices has, however, an additional operation, called **matrix multiplication.** If A is a matrix of size n by m, and if B is a matrix of size m by p, then the **product** $A \cdot B$ is defined to be the matrix C of size n by p whose general entry c_{ij} is given by the equation

$$c_{ij} = \sum_{k=1}^{m} a_{ik} b_{kj}.$$

This product operation satisfies the following properties, which are straightforward to verify:

(1) $A \cdot (B \cdot C) = (A \cdot B) \cdot C$.

(2) $A \cdot (B + C) = A \cdot B + A \cdot C$.

(3) $(A + B) \cdot C = A \cdot C + B \cdot C$.

(4) $(cA) \cdot B = c(A \cdot B) = A \cdot (cB)$.

(5) For each k, there is a k by k matrix I_k such that if A is any n by m matrix,

$$I_n \cdot A = A \quad \text{and} \quad A \cdot I_m = A.$$

In each of these statements, we assume that the matrices involved are of appropriate sizes, so that the indicated operations may be performed.

The matrix I_k is the matrix of size k by k whose general entry δ_{ij} is defined as follows: $\delta_{ij} = 0$ if $i \neq j$, and $\delta_{ij} = 1$ if $i = j$. The matrix I_k is called the **identity matrix** of size k by k; it has the form

$$I_k = \begin{bmatrix} 1 & 0 & \dots & 0 \\ 0 & 1 & \dots & 0 \\ & \dots & \dots & \\ 0 & 0 & \dots & 1 \end{bmatrix},$$

with entries of 1 on the "main diagonal" and entries of 0 elsewhere.

We extend to matrices the sup norm defined for n-tuples. That is, if A is a matrix of size n by m with general entry a_{ij}, we define

$$|A| = \max\{|a_{ij}|;\ i = 1, \dots, n \text{ and } j = 1, \dots, m\}.$$

The three properties of a norm are immediate, as is the following useful result:

Theorem 1.3. *If A has size n by m, and B has size m by p, then*

$$|A \cdot B| \leq m|A|\,|B|. \quad \square$$

Linear transformations

If V and W are vector spaces, a function $T : V \to W$ is called a **linear transformation** if it satisfies the following properties, for all $\mathbf{x}$, $\mathbf{y}$ in V and all scalars c:

(1) $T(\mathbf{x} + \mathbf{y}) = T(\mathbf{x}) + T(\mathbf{y})$.

(2) $T(c\mathbf{x}) = cT(\mathbf{x})$.

If, in addition, T carries V onto W in a one-to-one fashion, then T is called a **linear isomorphism.**

One checks readily that if $T : V \to W$ is a linear transformation, and if $S : W \to X$ is a linear transformation, then the composite $S \circ T : V \to X$ is a linear transformation. Furthermore, if $T : V \to W$ is a linear isomorphism, then $T^{-1} : W \to V$ is also a linear isomorphism.

A linear transformation is uniquely determined by its values on basis elements, and these values may be specified arbitrarily. That is the substance of the following theorem:

Theorem 1.4. *Let V be a vector space with basis $\mathbf{a}_1, \dots, \mathbf{a}_m$. Let W be a vector space. Given any m vectors $\mathbf{b}_1, \dots, \mathbf{b}_m$ in W, there is exactly one linear transformation $T : V \to W$ such that, for all i, $T(\mathbf{a}_i) = \mathbf{b}_i$.* □

In the special case where V and W are "tuple spaces" such as $\mathbb{R}^m$ and $\mathbb{R}^n$, matrix notation gives us a convenient way of specifying a linear transformation, as we now show.

First we discuss row matrices and column matrices. A matrix of size 1 by n is called a **row matrix;** the set of all such matrices bears an obvious resemblance to $\mathbb{R}^n$. Indeed, under the one-to-one correspondence

$$(x_1, \dots, x_n) \longrightarrow [x_1 \cdots x_n]$$

the vector space operations also correspond. Thus this correspondence is a linear isomorphism. Similarly, a matrix of size n by 1 is called a **column matrix;** the set of all such matrices also bears an obvious resemblance to $\mathbb{R}^n$. Indeed, the correspondence

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

is a linear isomorphism.

The second of these isomorphisms is particularly useful when studying linear transformations. Suppose for the moment that we represent elements of $\mathbb{R}^m$ and $\mathbb{R}^n$ by column matrices rather than by tuples. If A is a fixed n by m matrix, let us define a function $T : \mathbb{R}^m \longrightarrow \mathbb{R}^n$ by the equation

$$T(\mathbf{x}) = A \cdot \mathbf{x}.$$

The properties of matrix product imply immediately that T is a linear transformation.

In fact, every linear transformation of $\mathbb{R}^m$ to $\mathbb{R}^n$ has this form. The proof is easy. Given T, let $\mathbf{b}_1, \ldots, \mathbf{b}_m$ be the vectors of $\mathbb{R}^n$such that $T(\mathbf{e}_j) = \mathbf{b}_j$. Then let A be the n by m matrix $A = [\mathbf{b}_1 \cdots \mathbf{b}_m]$ with successive columns $\mathbf{b}_1, \ldots, \mathbf{b}_m$. Since the identity matrix has columns $\mathbf{e}_1, \ldots, \mathbf{e}_m$, the equation $A \cdot I_m = A$ implies that $A \cdot \mathbf{e}_j = \mathbf{b}_j$ for all j. Then $A \cdot \mathbf{e}_j = T(\mathbf{e}_j)$ for all j; it follows from the preceding theorem that $A \cdot \mathbf{x} = T(\mathbf{x})$ for all $\mathbf{x}$.

The convenience of this notation leads us to make the following convention:

Convention. *Throughout, we shall represent the elements of* $\mathbb{R}^n$ *by column matrices, unless we specifically state otherwise.*

Rank of a matrix

Given a matrix A of size n by m, there are several important linear spaces associated with A. One is the space spanned by the columns of A, looked at as column matrices (equivalently, as elements of $\mathbb{R}^n$). This space is called the **column space** of A, and its dimension is called the **column rank** of A. Because the column space of A is spanned by m vectors, its dimension can be no larger than m; because it is a subspace of $\mathbb{R}^n$, its dimension can be no larger than n.

Similarly, the space spanned by the rows of A, looked at as row matrices (or as elements of $\mathbb{R}^m$) is called the **row space** of A, and its dimension is called the **row rank** of A.

The following theorem is of fundamental importance:

Theorem 1.5. *For any matrix A, the row rank of A equals the column rank of A.* □

Once one has this theorem, one can speak merely of the **rank** of a matrix A, by which one means the number that equals both the row rank of A and the column rank of A.

The rank of a matrix A is an important number associated with A. One cannot in general determine what this number is by inspection. However, there is a relatively simple procedure called **Gauss-Jordan reduction** that can be used for finding the rank of a matrix. (It is used for other purposes as well.) We assume you have seen it before, so we merely review its major features here.