

John M. Lee

Introduction to Topological Manifolds

Second Edition

2011

 Springer

Chapter 1

Introduction

A course on manifolds differs from most other introductory graduate mathematics courses in that the subject matter is often completely unfamiliar. Most beginning graduate students have had undergraduate courses in algebra and analysis, so that graduate courses in those areas are continuations of subjects they have already begun to study. But it is possible to get through an entire undergraduate mathematics education, at least in the United States, without ever hearing the word “manifold.”

One reason for this anomaly is that even the definition of manifolds involves rather a large number of technical details. For example, in this book the formal definition does not come until the end of Chapter 2. Since it is disconcerting to embark on such an adventure without even knowing what it is about, we devote this introductory chapter to a nonrigorous definition of manifolds, an informal exploration of some examples, and a consideration of where and why they arise in various branches of mathematics.

What Are Manifolds?

Let us begin by describing informally how one should think about manifolds. The underlying idea is that manifolds are like curves and surfaces, except, perhaps, that they might be of higher dimension. Every manifold comes with a specific nonnegative integer called its *dimension*, which is, roughly speaking, the number of independent numbers (or “parameters”) needed to specify a point. The prototype of an n -dimensional manifold is n -dimensional Euclidean space \mathbb{R}^n , in which each point literally *is* an n -tuple of real numbers.

An n -dimensional manifold is an object modeled *locally* on \mathbb{R}^n ; this means that it takes exactly n numbers to specify a point, at least if we do not stray too far from a given starting point. A physicist would say that an n -dimensional manifold is an object with n *degrees of freedom*.

Manifolds of dimension 1 are just lines and curves. The simplest example is the real line; other examples are provided by familiar plane curves such as circles,




Fig. 1.1: Plane curves.




Fig. 1.2: Space curve.

parabolas, or the graph of any continuous function of the form $y = f(x)$ (Fig. 1.1). Still other familiar 1-dimensional manifolds are space curves, which are often described parametrically by equations such as $(x, y, z) = (f(t), g(t), h(t))$ for some continuous functions f, g, h (Fig. 1.2).

In each of these examples, a point can be unambiguously specified by a single real number. For example, a point on the real line *is* a real number. We might identify a point on the circle by its angle, a point on a graph by its x -coordinate, and a point on a parametrized curve by its parameter t . Note that although a parameter value determines a point, different parameter values may correspond to the same point, as in the case of angles on the circle. But in every case, as long as we stay close to some initial point, there is a one-to-one correspondence between nearby real numbers and nearby points on the line or curve.

Manifolds of dimension 2 are *surfaces*. The most common examples are planes and spheres. (When mathematicians speak of a sphere, we invariably mean a spherical *surface*, not a solid ball. The familiar unit sphere in \mathbb{R}^3 is 2-dimensional, whereas the solid ball is 3-dimensional.) Other familiar surfaces include cylinders, ellipsoids, paraboloids, hyperboloids, and the torus, which can be visualized as a doughnut-shaped surface in \mathbb{R}^3 obtained by revolving a circle around the z -axis (Fig. 1.3).

In these cases two coordinates are needed to determine a point. For example, on a plane we typically use Cartesian or polar coordinates; on a sphere we might use latitude and longitude; and on a torus we might use two angles. As in the 1-dimensional case, the correspondence between points and pairs of numbers is in general only local.




Fig. 1.3: Doughnut surface.

The only higher-dimensional manifold that we can easily visualize is Euclidean 3-space (or parts of it). But it is not hard to construct subsets of higher-dimensional Euclidean spaces that might reasonably be called manifolds. First, any open subset of \mathbb{R}^n is an n -manifold for obvious reasons. More interesting examples are obtained by using one or more equations to “cut out” lower-dimensional subsets. For example, the set of points (x_1, x_2, x_3, x_4) in \mathbb{R}^4 satisfying the equation

$$(x_1)^2 + (x_2)^2 + (x_3)^2 + (x_4)^2 = 1 \quad (1.1)$$

is called the (unit) 3-sphere. It is a 3-dimensional manifold because in a neighborhood of any given point it takes exactly three coordinates to specify a nearby point: starting at, say, the “north pole” $(0, 0, 0, 1)$, we can solve equation (1.1) for x_4 , and then each nearby point is uniquely determined by choosing appropriate (small) (x_1, x_2, x_3) coordinates and setting $x_4 = (1 - (x_1)^2 - (x_2)^2 - (x_3)^2)^{1/2}$. Near other points, we may need to solve for different variables, but in each case three coordinates suffice.

The key feature of these examples is that an n -dimensional manifold “looks like” \mathbb{R}^n locally. To make sense of the intuitive notion of “looks like,” we say that two subsets of Euclidean spaces $U \subseteq \mathbb{R}^k$, $V \subseteq \mathbb{R}^n$ are *topologically equivalent* or *homeomorphic* (from the Greek for “similar form”) if there exists a one-to-one correspondence $\varphi: U \rightarrow V$ such that both φ and its inverse are continuous maps. (Such a correspondence is called a *homeomorphism*.) Let us say that a subset M of some Euclidean space \mathbb{R}^k is *locally Euclidean of dimension n* if every point of M has a neighborhood in M that is topologically equivalent to a ball in \mathbb{R}^n .

Now we can give a provisional definition of manifolds. We can think of an n -dimensional manifold (n -manifold for short) as a subset of some Euclidean space \mathbb{R}^k that is locally Euclidean of dimension n . Later, after we have developed more machinery, we will give a considerably more general definition; but this one will get us started.




Fig. 1.4: Deforming a doughnut into a coffee cup.

Why Study Manifolds?

What follows is an incomplete survey of some of the fields of mathematics in which manifolds play an important role. This is not an overview of what we will be discussing in this book; to treat all of these topics adequately would take at least a dozen books of this size. Rather, think of this section as a glimpse at the vista that awaits you once you've learned to handle the basic tools of the trade.

Topology

Roughly speaking, topology is the branch of mathematics that is concerned with properties of sets that are unchanged by “continuous deformations.” Somewhat more accurately, a *topological property* is one that is preserved by homeomorphisms.

The subject in its modern form was invented near the end of the nineteenth century by the French mathematician Henri Poincaré, as an outgrowth of his attempts to classify geometric objects that appear in analysis. In a seminal 1895 paper titled *Analysis Situs* (the old name for topology, Latin for “analysis of position”) and a series of five companion papers published over the next decade, Poincaré laid out the main problems of topology and introduced an astonishing array of new ideas for solving them. As you read this book, you will see that his name is written all over the subject. In the intervening century, topology has taken on the role of providing the foundations for just about every branch of mathematics that has any use for a concept of “space.” (An excellent historical account of the first six decades of the subject can be found in [Die89].)

Here is a simple but telling example of the kind of problem that topological tools are needed to solve. Consider two surfaces in space: a sphere and a cube. It should not be hard to convince yourself that the cube can be continuously deformed into the sphere without tearing or collapsing it. It is not much harder to come up with an explicit formula for a homeomorphism between them (as we will do in Chapter 2). Similarly, with a little more work, you should be able to see how the surface of a doughnut can be continuously deformed into the surface of a one-handled coffee cup, by stretching out one half of the doughnut to become the cup, and shrinking the other half to become the handle (Fig. 1.4). Once you decide on an explicit set




Fig. 1.5: Turning a sphere inside out (with a crease).

of equations to define a “coffee-cup surface” in \mathbb{R}^3 , you could in principle come up with a set of formulas to describe a homeomorphism between it and a torus. On the other hand, a little reflection will probably convince you that there is no homeomorphism from a sphere to a torus: any such map would have to tear open a “hole” in the sphere, and thus could not be continuous.

It is usually relatively straightforward (though not always easy!) to prove that two manifolds are topologically equivalent once you have convinced yourself intuitively that they are: just write down an explicit homeomorphism between them. What is much harder is to prove that two manifolds are not homeomorphic—even when it seems “obvious” that they are not as in the case of a sphere and a torus—because you would need to show that no one, no matter how clever, could find such a map.

History abounds with examples of operations that mathematicians long believed to be impossible, only to be proved wrong. Here is an example from topology. Imagine a spherical surface colored white on the outside and gray on the inside, and imagine that it can move freely in space, including passing freely through itself. Under these conditions you could turn the sphere inside out by continuously deforming it, so that the gray side ends up facing out, but it seems obvious that in so doing you would have to introduce a crease somewhere. (It is possible to give precise mathematical definitions of what we mean by “continuously deforming” and “creases,” but you do not need to know them to get the general idea.) The simplest way to proceed would be to push the northern hemisphere down and the southern hemisphere up, allowing them to pass through each other, until the two hemispheres had switched places (Fig. 1.5); but this would introduce a crease along the equator. The topologist Stephen Smale stunned the mathematical community in 1958 [Sma58] when he proved it was possible to turn the sphere inside out without introducing any creases. Several ways to do this are beautifully illustrated in video recordings [Max77, LMM95, SFL98].

The usual way to prove that two manifolds are not topologically equivalent is by finding *topological invariants*: properties (which could be numbers or other mathematical objects such as groups, matrices, polynomials, or vector spaces) that are preserved by homeomorphisms. If two manifolds have different invariants, they cannot be homeomorphic.

It is evident from the examples above that geometric properties such as circumference and area are not topological invariants, because they are not generally pre-

served by homeomorphisms. Intuitively, the property that distinguishes a sphere from a torus is the fact that the latter has a “hole,” while the former does not. But it turns out that giving a precise definition of what is meant by a hole takes rather a lot of work.

One invariant that is commonly used to detect holes in a manifold is called the *fundamental group* of the manifold, which is a group (in the algebraic sense) attached to each manifold in such a way that homeomorphic manifolds have isomorphic fundamental groups. Different elements of the fundamental group represent inequivalent ways that a “loop,” or continuous closed path, can be drawn in the manifold, with two loops considered equivalent if one can be continuously deformed into the other while remaining in the manifold. The number of such inequivalent loops—in some sense, the “size” of the fundamental group—is one measure of the number of holes possessed by the manifold. A manifold in which every loop can be continuously shrunk to a single point has the trivial (one-element) group as its fundamental group; such a manifold is said to be *simply connected*. For example, a sphere is simply connected, but a torus is not. We will prove this rigorously in Chapter 8; but you can probably convince yourself intuitively that this is the case if you imagine stretching a rubber band around part of each surface and seeing if it can shrink itself to a point. On the sphere, no matter where you place the rubber band initially, it can always shrink down to a single point while remaining on the surface. But on the surface of a doughnut, there are at least two places to place the rubber band so that it cannot be shrunk to a point without leaving the surface (one goes around the hole in the middle of the doughnut, and the other goes around the part that would be solid if it were a real doughnut).

The study of the fundamental group occupies a major portion of this book. It is the starting point for *algebraic topology*, which is the subject that studies topological properties of manifolds (or other geometric objects) by attaching algebraic structures such as groups and rings to them in a topologically invariant way.

One of the most important problems of topology is the search for a classification of manifolds up to topological equivalence. Ideally, for each dimension n , one would like to produce a list of n -dimensional manifolds, and a theorem that says every n -dimensional manifold is homeomorphic to exactly one on the list. The theorem would be even better if it came with a list of computable topological invariants that could be used to decide where on the list any given manifold belongs. To make the problem more tractable, it is common to restrict attention to *compact* manifolds, which can be thought of as those that are homeomorphic to closed and bounded subsets of some Euclidean space.

Precisely such a classification theorem is known for 2-manifolds. The first part of the theorem says that every compact 2-manifold is homeomorphic to one of the following: a sphere, or a doughnut surface with $n \geq 1$ holes, or a connected sum of $n \geq 1$ projective planes. The second part says that no two manifolds on this list are homeomorphic to each other. We will define these terms and prove the first part of the theorem in Chapter 6, and in Chapter 10 we will use the technology provided by the fundamental group to prove the second part.

For higher-dimensional manifolds, the situation is much more complicated. The most delicate classification problem is that for compact 3-manifolds. It was already known to Poincaré that the 3-sphere is simply connected (we will prove this in Chapter 7), a property that distinguished it from all other examples of compact 3-manifolds known in his time. In the last of his five companion papers to *Analysis Situs*, Poincaré asked if it were possible to find a compact 3-manifold that is simply connected and yet *not* homeomorphic to the 3-sphere. Nobody ever found one, and the conjecture that every simply connected compact 3-manifold is homeomorphic to the 3-sphere became known as the *Poincaré conjecture*. For a long time, topologists thought of this as the simplest first step in a potential classification of 3-manifolds, but it resisted proof for a century, even as analogous conjectures were made and proved in higher dimensions (for 5-manifolds and higher by Stephen Smale in 1961 [Sma61], and for 4-manifolds by Michael Freedman in 1982 [Fre82]).

The intractability of the original 3-dimensional Poincaré conjecture led to its being acknowledged as the most important topological problem of the twentieth century, and many strategies were introduced for proving it. Surprisingly, the strategy that eventually succeeded involved techniques from differential geometry and partial differential equations, not just from topology. These techniques require far more groundwork than we are able to cover in this book, so we are not able to treat them here. But because of the significance of the Poincaré conjecture in the general theory of topological manifolds, it is worth saying a little more about its solution.

A major leap forward in our understanding of 3-manifolds occurred in the 1970s, when William Thurston formulated a much more powerful conjecture, now known as the *Thurston geometrization conjecture*. Thurston conjectured that every compact 3-manifold has a “geometric decomposition,” meaning that it can be cut along certain surfaces into finitely many pieces, each of which admits one of eight highly uniform (but mostly non-Euclidean) geometric structures. Since the manifolds with geometric structures are much better understood, the geometrization conjecture gives a nearly complete classification of 3-manifolds (but not yet complete, because there are still open questions about how many manifolds with certain non-Euclidean geometric structures exist). In particular, since the only compact, simply connected 3-manifold with a geometric decomposition is the 3-sphere, the geometrization conjecture implies the Poincaré conjecture.

The most important advance came in the 1980s, when Richard Hamilton introduced a tool called the *Ricci flow* for proving the existence of geometric decompositions. This is a partial differential equation that starts with an arbitrary geometric structure on a manifold and forces it to evolve in a way that tends to make its geometry increasingly uniform as time progresses, except in certain places where the curvature grows without bound. Hamilton proposed to use the places where the curvature becomes very large during the flow as a guide to where to cut the manifold, and then try to prove that the flow approaches one of the eight uniform geometries on each of the remaining pieces after the cuts are made. Hamilton made significant progress in implementing his program, but the technical details were formidable, requiring deep insights from topology, geometry, and partial differential equations.

In 2003, Russian mathematician Grigori Perelman figured out how to overcome the remaining technical obstacles in Hamilton's program, and completed the proof of the geometrization conjecture and thus the Poincaré conjecture. Thus the greatest challenge of twentieth century topology has been solved, paving the way for a much deeper understanding of 3-manifolds. Perelman's proof of the Poincaré conjecture is described in detail in the book [MT07].

In dimensions 4 and higher, there is no hope for a complete classification: it was proved in 1958 by A. A. Markov that there is no algorithm for classifying manifolds of dimension greater than 3 (see [Sti93]). Nonetheless, there is much that can be said using sophisticated combinations of techniques from algebraic topology, differential geometry, partial differential equations, and algebraic geometry, and spectacular progress was made in the last half of the twentieth century in understanding the variety of manifolds that exist. The topology of 4-manifolds, in particular, is currently a highly active field of research.

Vector Analysis

One place where you have already seen some examples of manifolds is in elementary vector analysis: the study of vector fields, line integrals, surface integrals, and vector operators such as the divergence, gradient, and curl. A line integral is, in essence, an integral over a 1-manifold, and a surface integral is an integral over a 2-manifold. The tools and theorems of vector analysis lie at the heart of the classical Maxwell theory of electromagnetism, for example.

Even in elementary treatments of vector analysis, topological properties play a role. You probably learned that if a vector field is the gradient of a function on some open domain in \mathbb{R}^3 , then its curl is identically zero. For certain domains, such as rectangular solids, the converse is true: every vector field whose curl is identically zero is the gradient of a function. But there are some domains for which this is not the case. For example, if $r = \sqrt{x^2 + y^2}$ denotes the distance from the z -axis, the vector field whose component functions are $(-y/r^2, x/r^2, 0)$ is defined everywhere in the domain D consisting of \mathbb{R}^3 with the z -axis removed, and has zero curl. It would be the gradient of the polar angle function $\theta = \tan^{-1}(y/x)$, except that there is no way to define the angle function continuously on all of D .

The question of whether every curl-free vector field is a gradient can be rephrased in such a way that it makes sense on a manifold of any dimension, provided the manifold is sufficiently "smooth" that one can take derivatives. The answer to the question, surprisingly, turns out to be a purely topological one. If the manifold is simply connected, the answer is yes, but in general simple connectivity is not necessary. The precise criterion that works for manifolds in all dimensions involves the concept of *homology* (or rather, its closely related cousin *cohomology*), which is an alternative way of measuring "holes" in a manifold. We give a brief introduction to homology and cohomology in Chapter 13 of this book; a more thorough treatment of the relationship between gradients and topology can be found in [Lee02].

Geometry

The principal objects of study in Euclidean plane geometry, as you encountered it in secondary school, are figures constructed from portions of lines, circles, and other curves—in other words, 1-manifolds. Similarly, solid geometry is concerned with figures made from portions of planes, spheres, and other 2-manifolds. The properties that are of interest are those that are invariant under rigid motions. These include simple properties such as lengths, angles, areas, and volumes, as well as more sophisticated properties derived from them such as curvature. The curvature of a curve or surface is a quantitative measure of how it bends and in what directions; for example, a positively curved surface is “bowl-shaped,” whereas a negatively curved one is “saddle-shaped.”

Geometric theorems involving curves and surfaces range from the trivial to the very deep. A typical theorem you have undoubtedly seen before is the angle-sum theorem: the sum of the interior angles of any Euclidean triangle is π radians. This seemingly trivial result has profound generalizations to the study of curved surfaces, where angles may add up to more or less than π depending on the curvature of the surface. The high point of surface theory is the Gauss–Bonnet theorem: for a closed, bounded surface in \mathbb{R}^3 , this theorem expresses the relationship between the total curvature (i.e., the integral of curvature with respect to area) and the number of holes the surface has. If the surface is topologically equivalent to an n -holed doughnut surface, the theorem says that the total curvature is exactly equal to $4\pi - 4\pi n$. In the case $n = 1$ this implies that no matter how a one-holed doughnut surface is bent or stretched, the regions of positive and negative curvature will always precisely cancel each other out so that the total curvature is zero.

The introduction of manifolds has allowed the study of geometry to be carried into higher dimensions. The appropriate setting for studying geometric properties in arbitrary dimensions is that of *Riemannian manifolds*, which are manifolds on which there is a rule for measuring distances and angles, subject to certain natural restrictions to ensure that these quantities behave analogously to their Euclidean counterparts. The properties of interest are those that are invariant under *isometries*, or distance-preserving homeomorphisms. For example, one can study the relationship between the curvature of an n -dimensional Riemannian manifold (a local property) and its global topological type. A typical theorem is that a complete Riemannian n -manifold whose curvature is everywhere larger than some fixed positive number must be compact and have a finite fundamental group (not too many holes). The search for such relationships is one of the principal activities in Riemannian geometry, a thriving field of contemporary research. See Chapter 1 of [Lee97] for an informal introduction to the subject.

Algebra

One of the most important objects studied in abstract algebra is the *general linear group* $GL(n, \mathbb{R})$, which is the group of $n \times n$ invertible real matrices, with matrix multiplication as the group operation. As a set, it can be identified with a subset of n^2 -dimensional Euclidean space, simply by stringing all the matrix entries out in a row. Since a matrix is invertible if and only if its determinant is nonzero, $GL(n, \mathbb{R})$ is an open subset of \mathbb{R}^{n^2} , and is therefore an n^2 -dimensional manifold. Similarly, the *complex general linear group* $GL(n, \mathbb{C})$ is the group of $n \times n$ invertible complex matrices; it is a $2n^2$ -manifold, because we can identify \mathbb{C}^{n^2} with \mathbb{R}^{2n^2} .

A *Lie group* is a group (in the algebraic sense) that is also a manifold, together with some technical conditions to ensure that the group structure and the manifold structure are compatible with each other. They play central roles in differential geometry, representation theory, and mathematical physics, among many other fields. The most important Lie groups are subgroups of the real and complex general linear groups. Some commonly encountered examples are the *special linear group* $SL(n, \mathbb{R}) \subseteq GL(n, \mathbb{R})$, consisting of matrices with determinant 1; the *orthogonal group* $O(n) \subseteq GL(n, \mathbb{R})$, consisting of matrices whose columns are orthonormal; the *special orthogonal group* $SO(n) = O(n) \cap SL(n, \mathbb{R})$; and their complex analogues, the *complex special linear group* $SL(n, \mathbb{C}) \subseteq GL(n, \mathbb{C})$, the *unitary group* $U(n) \subseteq GL(n, \mathbb{C})$, and the *special unitary group* $SU(n) = U(n) \cap SL(n, \mathbb{C})$.

It is important to understand the topological structure of a Lie group and how its topological structure relates to its algebraic structure. For example, it can be shown that $SO(2)$ is topologically equivalent to a circle, $SU(2)$ is topologically equivalent to the 3-sphere, and any connected abelian Lie group is topologically equivalent to a Cartesian product of circles and lines. Lie groups provide a rich source of examples of manifolds in all dimensions.

Complex Analysis

Complex analysis is the study of holomorphic (i.e., complex analytic) functions. If f is any complex-valued function of a complex variable, its *graph* is a subset of $\mathbb{C}^2 = \mathbb{C} \times \mathbb{C}$, namely $\{(z, w) : w = f(z)\}$. More generally, the graph of a holomorphic function of n complex variables is a subset of $\mathbb{C}^n \times \mathbb{C} = \mathbb{C}^{n+1}$. Because the set \mathbb{C} of complex numbers is naturally identified with \mathbb{R}^2 , and therefore the n -dimensional complex Euclidean space \mathbb{C}^n can be identified with \mathbb{R}^{2n} , we can consider graphs of holomorphic functions as manifolds, just as we do for real-valued functions.

Some holomorphic functions are naturally “multiple-valued.” A typical example is the complex square root. Except for zero, every complex number has two distinct square roots. But unlike the case of positive real numbers, where we can always unambiguously choose the positive square root to denote by the symbol \sqrt{x} , it is




Fig. 1.6: Graph of the two branches of the real square root.

not possible to define a global continuous square root function on the complex plane. To see why, write z in polar coordinates as $z = r e^{i\theta} = r(\cos \theta + i \sin \theta)$. Then the two square roots of z can be written $\sqrt{r} e^{i\theta/2}$ and $\sqrt{r} e^{i(\theta/2+\pi)}$. As θ increases from 0 to 2π , the first square root goes from the positive real axis through the upper half-plane to the negative real axis, while the second goes from the negative real axis through the lower half-plane to the positive real axis. Thus whichever continuous square root function we start with on the positive real axis, we are forced to choose the other after having made one circuit around the origin.

Even though a “two-valued function” is properly considered as a relation and not really a function at all, we can make sense of the *graph* of such a relation in an unambiguous way. To warm up with a simpler example, consider the two-valued square root “function” on the nonnegative real axis. Its graph is defined to be the set of pairs $(x, u) \in \mathbb{R} \times \mathbb{R}$ such that $u = \pm\sqrt{x}$, or equivalently $u^2 = x$. This is a parabola opening in the positive x direction (Fig. 1.6), which we can think of as the two “branches” of the square root.

Similarly, the graph of the two-valued complex square root “function” is the set of pairs $(z, w) \in \mathbb{C}^2$ such that $w^2 = z$. Over each small disk $U \subseteq \mathbb{C}$ that does not contain 0, this graph has two branches or “sheets,” corresponding to the two possible continuous choices of square root function on U (Fig. 1.7). If you start on one sheet above the positive real axis and pass once around the origin in the counterclockwise direction, you end up on the other sheet. Going around once more brings you back to the first sheet.

It turns out that this graph in \mathbb{C}^2 is a 2-dimensional manifold, of a special type called a *Riemann surface*: this is essentially a 2-manifold on which there is some way to define holomorphic functions. Riemann surfaces are of great importance in complex analysis, because any holomorphic function gives rise to a Riemann surface by a procedure analogous to the one we sketched above. The surface we constructed turns out to be topologically equivalent to a plane, but more complicated functions can give rise to more complicated surfaces. For example, the two-valued “function” $f(z) = \pm\sqrt{z^3 - z}$ yields a Riemann surface that is homeomorphic to a plane with one “handle” attached.




Fig. 1.7: Two branches of the complex square root.

One of the fundamental tasks of complex analysis is to understand the topological type (number of “holes” or “handles”) of the Riemann surface of a given function, and how it relates to the analytic properties of the function.

Algebraic Geometry

Algebraic geometers study the geometric and topological properties of solution sets to systems of polynomial equations. Many of the basic questions of algebraic geometry can be posed very naturally in the elementary context of plane curves defined by polynomial equations. For example: How many intersection points can one expect between two plane curves defined by polynomials of degrees k and l ? (Not more than kl , but sometimes fewer.) How many disconnected “pieces” does the solution set to a particular polynomial equation have (Fig. 1.8)? Does a plane curve have any self-crossings (Fig. 1.9) or “cusps” (points where the tangent vector does not vary continuously—Fig. 1.10)?

But the real power of algebraic geometry becomes evident only when one focuses on polynomials with coefficients in an algebraically closed field (one in which every polynomial decomposes into a product of linear factors), because polynomial equations always have the expected number of solutions (counted with multiplicity) in that case. The most extensively studied case is the complex field; in this context the solution set to a system of complex polynomials in n variables is a certain geometric object in \mathbb{C}^n called an *algebraic variety*, which (except for a small subset where there might be self-crossings or more complicated kinds of behavior) is a manifold. The subject becomes even more interesting if one enlarges \mathbb{C}^n by adding “ideal points at infinity” where parallel lines or asymptotic curves can be thought of




Fig. 1.8: A plane curve with disconnected pieces.




Fig. 1.9: A self-crossing.




Fig. 1.10: A cusp.

as meeting; the resulting set is called *complex projective space*, and is an extremely important manifold in its own right.

The properties of interest are those that are invariant under projective transformations (the natural changes of coordinates on projective space). One can ask such questions as these: Is a given variety a manifold, or does it have singular points (points where it fails to be a manifold)? If it is a manifold, what is its topological type? If it is not a manifold, what is the topological structure of its singular set, and how does that set change when one varies the coefficients of the polynomials slightly? If two varieties are homeomorphic, are they equivalent under a projective transformation? How many times and in what way do two or more varieties intersect?

Algebraic geometry has contributed a prodigious supply of examples of manifolds. In particular, much of the recent progress in understanding 4-dimensional manifolds has been driven by the wealth of examples that arise as algebraic varieties.

Computer Graphics

The job of a computer graphics program is to generate realistic images of 3-dimensional objects, for such applications as movies, simulators, industrial design, and computer games. The surfaces of the objects being modeled are usually represented as 2-dimensional manifolds.

A surface for which a simple equation is known—a sphere, for example—is easy to model on a computer. But there is no single equation that describes the surface of an airplane or a dinosaur. Thus computer graphics designers typically create models of surfaces by specifying multiple coordinate patches, each of which represents a small region homeomorphic to a subset of \mathbb{R}^2 . Such regions can be described by simple polynomial functions, called *splines*, and the program can ensure that the various splines fit together to create an appropriate global surface. Analyzing the tangent plane at each point of a surface is important for understanding how light reflects and scatters from the surface; and analyzing the curvature is important to ensure that adjacent splines fit together smoothly without visible “seams.” If it is necessary to create a model of an already existing surface rather than one being designed from scratch, then it is necessary for the program to find an efficient way to subdivide the surface into small pieces, usually triangles, which can then be represented by splines.

Computer graphics programmers, designers, and researchers make use of many of the tools of manifold theory: coordinate charts, parametrizations, triangulations, and curvature, to name just a few.

Classical Mechanics

Classical mechanics is the study of systems that obey Newton’s laws of motion. The positions of all the objects in the system at any given time can be described by a set of numbers, or coordinates; typically, these are not independent of each other but instead must satisfy some relations. The relations can usually be interpreted as defining a manifold in some Euclidean space.

For example, consider a rigid body moving through space under the influence of gravity. If we choose three noncollinear points P , Q , and R on the body (Fig. 1.11), the position of the body is completely specified once we know the coordinates of these three points, which correspond to a point in \mathbb{R}^9 . However, the positions of the three points cannot all be specified arbitrarily: because the body is rigid, they are subject to the constraint that the distances between pairs of points are fixed. Thus, to position the body in space, we can arbitrarily specify the coordinates of P (three parameters), and then we can specify the position of Q by giving, say, its latitude and longitude on the sphere of radius d_{PQ} , the fixed distance between P and Q (two more parameters). Finally, having determined the position of the two points P and Q , the only remaining freedom is to rotate R around the line PQ ; so we can specify the position of R by giving the angle θ that the plane PQR makes with




Fig. 1.11: A rigid body in space.

some reference plane (one more parameter). Thus the set of possible positions of the body is a certain 6-dimensional manifold $M \subseteq \mathbb{R}^9$.

Newton's second law of motion expresses the acceleration of the object—that is, the second derivatives of the coordinates of P , Q , R —in terms of the force of gravity, which is a certain function of the object's position. This can be interpreted as a system of second-order ordinary differential equations for the position coordinates, whose solutions are all the possible paths the rigid body can take on the manifold M .

The study of classical mechanics can thus be interpreted as the study of ordinary differential equations on manifolds, a subject known as *smooth dynamical systems*. A wealth of interesting questions arise in this subject: How do solutions behave over the long term? Are there any equilibrium points or periodic trajectories? If so, are they *stable*; that is, do nearby trajectories stay nearby? A good understanding of manifolds is necessary to fully answer these questions.

General Relativity

Manifolds play a decisive role in Einstein's general theory of relativity, which describes the interactions among matter, energy, and gravitational forces. The central assertion of the theory is that *spacetime* (the collection of all points in space at all times in the history of the universe) can be modeled by a 4-dimensional manifold that carries a certain kind of geometric structure called a *Lorentz metric*; and this metric satisfies a system of partial differential equations called the *Einstein field*

equations. Gravitational effects are then interpreted as manifestations of the curvature of the Lorentz metric.

In order to describe the global structure of the universe, its history, and its possible futures, it is important to understand first of all which 4-manifolds can carry Lorentz metrics, and for each such manifold how the topology of the manifold influences the properties of the metric. There are especially interesting relationships between the local geometry of spacetime (as reflected in the local distribution of matter and energy) and the global topological structure of the universe; these relationships are similar to those described above for Riemannian manifolds, but are more complicated because of the introduction of forces and motion into the picture. In particular, if we assume that on a cosmic scale the universe looks approximately the same at all points and in all directions (such a spacetime is said to be *homogeneous* and *isotropic*), then it turns out there is a critical value for the average density of matter and energy in the universe: above this density, the universe closes up on itself spatially and will collapse to a one-point singularity in a finite amount of time (the “big crunch”); below it, the universe extends infinitely far in all directions and will expand forever. Interestingly, physicists’ best current estimates place the average density rather near the critical value, and they have so far been unable to determine whether it is above or below it, so they do not know whether the universe will go on existing forever or not.

String Theory

One of the most fundamental and perplexing challenges for modern physics is to resolve the incompatibilities between quantum theory and general relativity. An approach that some physicists consider very promising is called *string theory*, in which manifolds appear in several different starring roles.

One of the central tenets of string theory is that elementary particles should be modeled as vibrating submicroscopic 1-dimensional objects, called “strings,” instead of points. This approach promises to resolve many of the contradictions that plagued previous attempts to unify gravity with the other forces of nature. But in order to obtain a consistent string theory, it seems to be necessary to assume that spacetime has more than four dimensions. We experience only four of them directly, because the dimensions beyond four are so tightly “curled up” that they are not visible on a macroscopic scale, much as a long but microscopically narrow 2-dimensional cylinder would appear to be 1-dimensional when viewed on a large enough scale. The topological properties of the manifold that appears as the “cross-section” of the curled-up dimensions have such a profound effect on the observable dynamics of the resulting theory that it is possible to rule out most cross-sections a priori.

Several different kinds of string theory have been constructed, but all of them give consistent results only if the cross-section is a certain kind of 6-dimensional manifold known as a *Calabi–Yau manifold*. More recently, evidence has been un-

covered that all of these string theories are different limiting cases of a single underlying theory, dubbed *M-theory*, in which the cross-section is a 7-manifold. These developments in physics have stimulated profound advancements in the mathematical understanding of manifolds of dimensions 6 and 7, and Calabi–Yau manifolds in particular.

Another role that manifolds play in string theory is in describing the history of an elementary particle. As a string moves through spacetime, it traces out a 2-dimensional manifold called its *world sheet*. Physical phenomena arise from the interactions among these different topological and geometric structures: the world sheet, the 6- or 7-dimensional cross-section, and the macroscopic 4-dimensional spacetime that we see.

It is still too early to predict whether string theory will turn out to be a useful description of the physical world. But it has already established a lasting place for itself in mathematics.

Manifolds are used in many more areas of mathematics than the ones listed here, but this brief survey should be enough to show you that manifolds have a rich assortment of applications. It is time to get to work.