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 will 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 has a *dimension*, which is, roughly speaking, the number of independent numbers (or "parameters") needed to specify a point. The prototype of

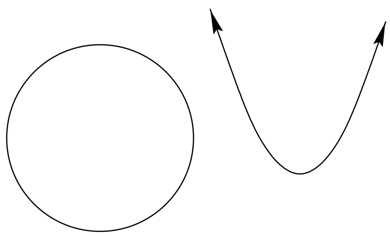

FIGURE 1.1. Plane curves.

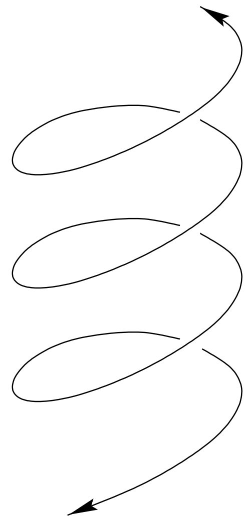

FIGURE 1.2. Space curve.

an n-dimensional manifold is n-dimensional Euclidean space $\mathbb{R}^n$, in which each point *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 commonly called *curves* (although they need not be "curved" in the ordinary sense of the word). The simplest example is the real line; other examples are provided by familiar plane curves such as circles, parabolas, or the graph of any continuous function of the form $y = f(x)$ (Figure 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 (Figure 1.2).

In each of these examples, a point on the curve can be unambiguously specified by a single real number. For example, a point on the real line *is* a real number. We might specify 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 curve.

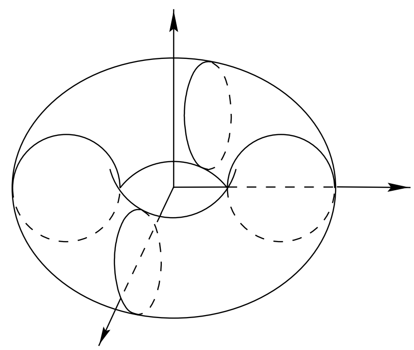

FIGURE 1.3. Doughnut surface.

Manifolds of dimension 2 are *surfaces*. The two most common examples are planes and spheres. (When mathematicians speak of a sphere, we invariably mean a spherical *surface*, which is 2-dimensional, not a solid ball, which is 3-dimensional.) Other familiar surfaces include cylinders, ellipsoids, paraboloids, and the doughnut-shaped surface in $\mathbb{R}^3$ obtained by revolving a circle around the z-axis (Figure 1.3). (This doughnut-shaped surface is often called a *torus*, but we will reserve that name for a slightly different but closely related object, to be introduced in the next chapter.)

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

The only higher-dimensional manifold that we can visualize is Euclidean 3-space. 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

$$(x_1)^2 + (x_2)^2 + (x_3)^2 + (x_4)^2 = 1 \tag{1.1}$$

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 will say that two subsets of Euclidean spaces $U \subset \mathbb{R}^k$, $V \subset \mathbb{R}^n$ are *topologically equivalent* or *homeomorphic* (Greek for "same form") if there exists a one-to-one correspondence $\varphi\colon U \to V$ such that both φ and its inverse are continuous maps. (Such a correspondence is called a *homeomorphism*.) 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 preliminary definition of manifolds. An *n-dimensional manifold* (*n-manifold* for short) is 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.

Why Study Manifolds?

What follows is an incomplete survey of some of the fields of mathematics in which manifolds play an important role.

Topology

Roughly speaking, topology is the branch of mathematics that is concerned with properties of sets that are unchanged by "continuous deformations." More accurately, a topological property is one that is preserved by homeomorphisms.

The subject in its modern form was invented a century ago 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 companion papers in 1899–1905, 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 topology was invented 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 (we will do so in Chapter 2). Similarly, with a little more work, you

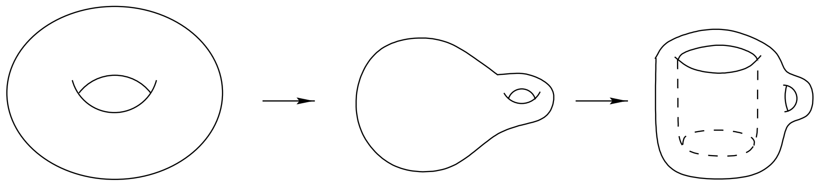

FIGURE 1.4. Deforming a doughnut into a coffee cup.

should be able to see how a doughnut surface 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 (Figure 1.4). Once you decide on an explicit set 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 the doughnut surface. On the other hand, a little reflection will probably convince you that there is *no* homeomorphism from the sphere to the doughnut surface: 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 the sphere and the doughnut—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. (There are precise mathematical definitions of the terms "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 (Figure 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, LMM94, SFL98].

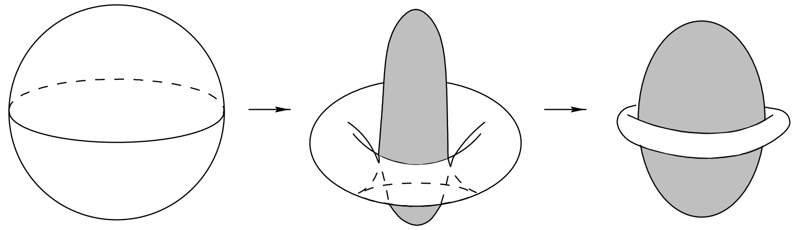

FIGURE 1.5. Turning a sphere inside out (with a crease).

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 preserved by homeomorphisms. Intuitively, the property that distinguishes the sphere from the doughnut surface 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 count 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 groups. Then the "size" of the fundamental group is a measure of the number of holes possessed by the manifold. The study of the fundamental group will occupy 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 problem of classifying manifolds. Ideally, 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, together with a list of computable topological invariants that could be used to decide where on the list any given manifold belongs. Precisely such a theorem is known for surfaces: It says that every compact surface is homeomorphic to a sphere, or to a doughnut surface with a finite number of holes, or to a connected sum of projective planes. (We will define these terms and prove the theorem in Chapter 6.)

For higher-dimensional manifolds, the situation is much more complicated. For example, Poincaré conjectured around 1900 that any compact 3-

manifold whose fundamental group is the trivial (one-element) group must be homeomorphic to the 3-sphere. For a long time, topologists thought of this as the simplest first step in a potential classification of 3-manifolds. But although analogous conjectures have been made for higher-dimensional manifolds and were proved in the intervening years (for 5-manifolds and higher by Stephen Smale in 1961 [Sma61], and for 4-manifolds by Michael Freedman in 1982 [Fre82]), the original Poincaré conjecture remains as of this writing a preeminent unsolved problem in topology. The best hope for a classification of 3-manifolds is the *geometrization conjecture* made in the 1970s by William Thurston (see [Sco83, Thu97] for an explanation), which says, roughly, that every compact 3-manifold can be cut into finitely many pieces each of which admits one of eight (mostly non-Euclidean) geometric structures. Since the manifolds with geometric structures are much better understood, a proof of this conjecture would go a long way toward providing a complete classification of 3-manifolds; in particular, it would imply that the Poincaré conjecture is true.

In dimensions 4 and higher, on the other hand, 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.

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," while 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 it 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 transformations. 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.

Complex Analysis

Complex analysis is the study of holomorphic (i.e., complex analytic) functions. Some such 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 not possible to define a global continuous square root function on the complex plane. To see why, write z in polar coordinates as $z = re^{i\theta}$. 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 define 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

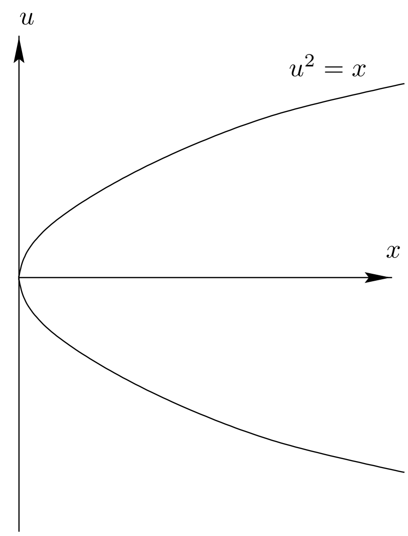

FIGURE 1.6. Graph of the two branches of the real square root.

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 (Figure 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 $(z, w) \in \mathbb{C} \times \mathbb{C}$ such that $w^2 = z$. Over each small disk $U \subset \mathbb{C}$ 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 (Figure 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, since 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.

One of the fundamental tasks of complex analysis is to understand the topological type (number of "holes" or "handles") of the Riemann surface

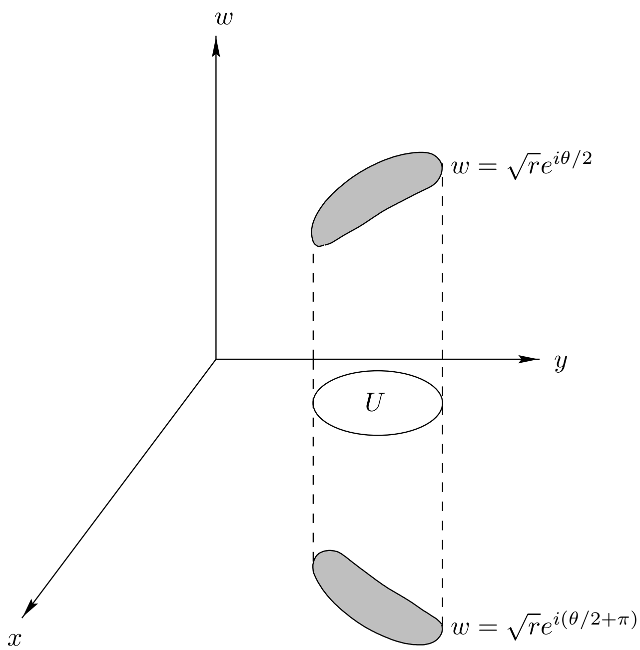

FIGURE 1.7. Two branches of the complex square root.

of a given function, and how it relates to the analytic properties of the function.

Algebra

One of the most important objects studied in abstract algebra is the *general linear group* $\mathrm{GL}(n,\mathbb{R})$, which is the group of $n \times n$ invertible real matrices. 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, $\mathrm{GL}(n,\mathbb{R})$ is an open subset of $\mathbb{R}^{n^2}$, and is therefore an n^2-dimensional manifold. Similarly, the *complex general linear group* $\mathrm{GL}(n,\mathbb{C})$ 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 a central role 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

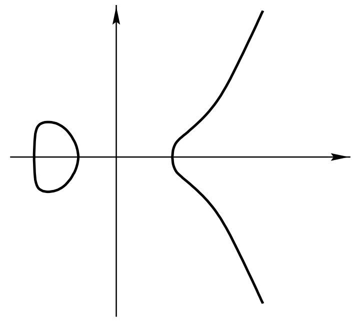

FIGURE 1.8. A plane curve with disconnected pieces.

encountered examples are the *special linear group* $\mathrm{SL}(n,\mathbb{R}) \subset \mathrm{GL}(n,\mathbb{R})$, consisting of matrices with determinant 1; the *orthogonal group* $\mathrm{O}(n) \subset \mathrm{GL}(n,\mathbb{R})$, consisting of matrices whose columns are orthonormal; the *special orthogonal group* $\mathrm{SO}(n) = \mathrm{O}(n) \cap \mathrm{SL}(n,\mathbb{R})$; and their complex analogues, the *complex special linear group* $\mathrm{SL}(n,\mathbb{C}) \subset \mathrm{GL}(n,\mathbb{C})$, the *unitary group* $\mathrm{U}(n) \subset \mathrm{GL}(n,\mathbb{C})$, and the *special unitary group* $\mathrm{SU}(n) = \mathrm{U}(n) \cap \mathrm{SL}(n,\mathbb{C})$.

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 $\mathrm{SO}(2)$ is topologically equivalent to a circle, $\mathrm{SU}(2)$ 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.

Algebraic Geometry

Algebraic geometers study the geometric 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 (Figure 1.8)? Does a plane curve have any self crossings (Figure 1.9) or "cusps" (points where the tangent vector does not vary continuously—Figure 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), since polynomial equations always have the expected number of solutions (counted with multiplicity) in that case. The most deeply studied case is the complex field; in this context the solution set to a system of complex

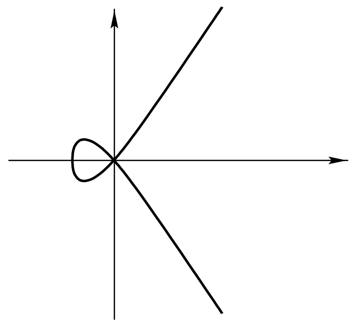

FIGURE 1.9. A self crossing.

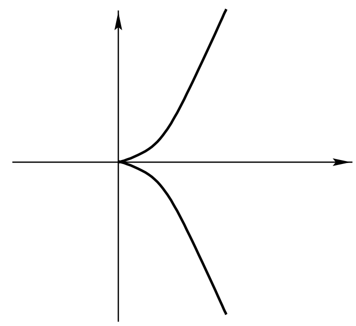

FIGURE 1.10. A cusp.

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, cusps, 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 as meeting; the resulting space 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 geometric 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.

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.

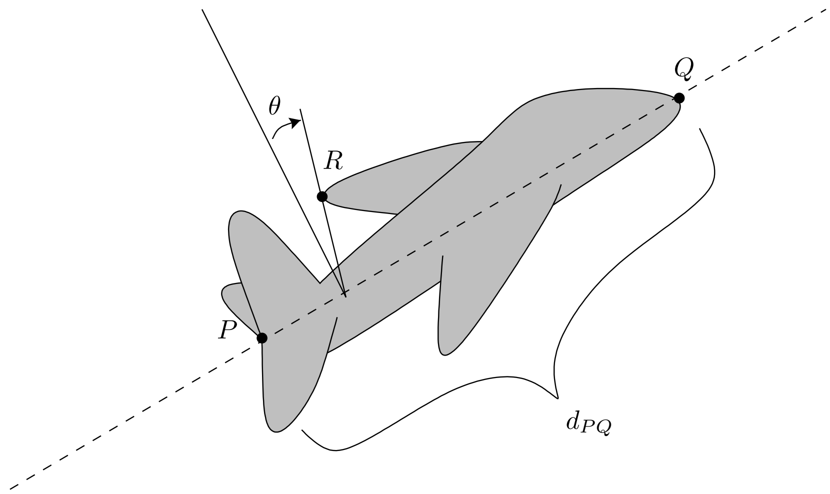

FIGURE 1.11. A rigid body in space.

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 (Figure 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 determine the position of the body, we can arbitrarily specify the coordinates of P in space (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 some reference plane (one more parameter). Thus the set of possible positions of the body is a certain 6-dimensional manifold $M \subset \mathbb{R}^9$.

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, also 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 history) 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 what kinds of 4-manifolds exist and what kinds of Lorentz metrics they can carry. 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 point singularity in a finite 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.

Quantum Field Theory

The theory of elementary particle interactions, called quantum field theory, has become increasingly geometric in recent decades. In particular, the latest attempts to unify quantum theory and gravitation have led to ever more interesting and exotic geometric structures. The approach to quantum gravity that is currently considered most promising by many physicists is *string theory*, in which manifolds appear in several different starring roles.

First, in order to obtain a consistent 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 cylinder would appear to be one-dimensional when viewed from far enough away. 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 quantum field theory that it is possible to rule out most cross sections a priori. It currently appears that a consistent theory can be constructed only if the cross section is a certain kind of 6-dimensional manifold known as a *Calabi–Yau manifold*. These developments in physics have stimulated profound developments in the mathematical understanding of 6-manifolds in general and Calabi-Yau manifolds in particular.

Another role that manifolds play in string theory is in describing the history of an elementary particle. One of the central tenets of string theory is that particles should be represented not as points, but as tiny 1-dimensional objects ("strings") moving through spacetime. As a particle moves, 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 Calabi-Yau cross section, and the macroscopic four-dimensional spacetime that we see.

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.