
Introduction

This book is about differential equations—a **very** big subject! It is so extensive, in fact, that we could not hope to cover it completely even in a book many times this size. So we will have to be selective. In the first place, we will restrict our attention almost entirely to *equations of evolution*. That is to say, we will be considering quantities q that depend on a “time” variable t , and we will be considering mainly *initial value problems*. This is the problem of predicting the value of such a quantity q at a time t_1 from its value at some (usually earlier) “initial” time t_0 , assuming that we know the “law of evolution” of q . The latter will always be a “differential equation” that tells us how to compute the rate at which q is changing from a knowledge of its current value. While we will concentrate mainly on the easier case of an ordinary differential equation (ODE), where the quantity q depends **only** on the time, we will on occasion consider the partial differential equation (PDE) case, where q depends also on other “spatial variables” x as well as the time t and where the partial derivatives of q with respect to these spatial variables can enter into the law determining its rate of change with respect to time.

Our principal goal will be to help you develop a good intuition for equations of evolution and how they can be used to model a large variety of time-dependent processes—in particular those that arise in the study of classical mechanics. To this end we will stress various metaphors that we hope will encourage you to get started thinking creatively about differential equations and their solutions.

But wait! Just who is this “you” we are addressing? Every textbook author has in mind at least a rough image of some prototypical

student for whom he is writing, and since the assumed background and abilities of this model student are sure to have an important influence on how the book gets written, it is only fair that we give you some idea of our own preconceptions about you.

We are assuming that, at a minimum, the usual reader of this book will have completed the equivalent of two years of undergraduate mathematics in a U.S. college or university and, in particular, will have had a solid introduction to linear algebra and to multi-variable (aka “advanced”) calculus. But in all honesty, we have in mind some other hoped-for qualities in our reader, principally that he or she is accustomed to and enjoys seeing mathematics presented conceptually and not as a collection of cookbook methods for solving standard exercises. And finally we hope our readers enjoy working out mathematical details on their own. We will give frequent exercises (usually with liberal hints) that ask the student to fill in some details of a proof or derive a corollary.

A related question is how we expect this book to be used. We would of course be delighted to hear that it has been adopted as the assigned text for many junior and senior level courses in differential equations (and perhaps not surprisingly we would be happy using it ourselves in teaching such a course). But we realize that the book we have written diverges in many ways from the current “standard model” of an ODE text, so it is our real hope and expectation that many students, particularly those of the sort described above, will find it a challenging but helpful source from which to learn about ODEs, either on their own or as a supplement to a more standard assigned text while taking an ODE course.

We should mention here—and explain—a somewhat unusual feature of our exposition. The book consists of two parts that we will refer to as “text” and “appendices”. The text is made up of five chapters that together contain about two-thirds of the material, while the appendices consist of ten shorter mini-chapters. Our aim was to make the text relatively easy reading by relegating the more difficult and technical material to the appendices. A reader should be able to get a quick overview of the subject matter of one or more chapters by just reading the text and ignoring the references to material in the

appendices. Later, when ready to go deeper or to check an omitted proof, a reading of the relevant appendices should satisfy the reader's hunger for more detail.

Finally we would like to discuss “visual aids”—that is, the various kinds of diagrams and pictures that make it easier for a student to internalize a complicated mathematical concept upon meeting it for the first time. Both of the authors have been very actively involved with the development of software tools for creating such mathematical visualizations and with investigating techniques for using them to enhance the teaching and learning of mathematics, and paradoxically that has made it difficult for us to choose appropriate figures for our text. Indeed, recent advances in technology, in particular the explosive development of the Internet and in particular of the World Wide Web, have not only made it easy to provide visual material online, but moreover the expressiveness possible using the interactive and animated multimedia tools available in the virtual world of the Internet far surpasses that of the classic static diagrams that have traditionally been used in printed texts. As a result we at first considered omitting diagrams entirely from this text, but in the end we decided on a dual approach. We have used traditional diagrams in the text where we felt that they would be useful, and in addition we have placed a much richer assortment of visual material online to accompany the text. Our publisher, the American Mathematical Society, has agreed to set aside a permanent area on its own website to be devoted to this book, and throughout the text you will find references to this area that we will refer to as the “Web Companion”.¹ Here, organized by chapter and section, you will find visualizations that go far beyond anything we could hope to put in the pages of a book—static diagrams, certainly, but in addition Flash animations, Java applets, QuickTime movies, Mathematica, Matlab, Maple Notebooks, other interactive learning aids, and also links to other websites that contain material we believe will help and speed your understanding. And not only does this approach allow us to make much more sophisticated visualizations available, but it also will permit us to add new and improved material as it becomes available.

¹Its URL is <http://www.ams.org/bookpages/stml-51>.

Chapter 1

Differential Equations and Their Solutions

1.1. First-Order ODE: Existence and Uniqueness

What does the following sentence mean, and what image should it cause you to form in your mind?

Let $V : \mathbf{R}^n \times \mathbf{R} \rightarrow \mathbf{R}^n$ be a time-dependent vector field, and let $x(t)$ be a solution of the differential equation $\frac{dx}{dt} = V(x, t)$ satisfying the initial condition $x(t_0) = x_0$.

Let us consider a seemingly very different question. Suppose you know the wind velocity at every point of space and at all instants of time. A puff of smoke drifts by, and at a certain moment you note the precise location of a particular smoke particle. Can you then predict where that particle will be at all future times?

We will see that when this somewhat vague question is translated appropriately into precise mathematical concepts, it leads to the above “differential equation”, and that the answer to our prediction question translates to the central existence and uniqueness result in the theory of differential equations. (The answer, by the way, turns out to be a qualified “yes”, with several important caveats.)

We interpret “space” to mean the n -dimensional real number space \mathbf{R}^n , so a “point of space” is just an n -tuple $x = (x_1, \dots, x_n)$ of real numbers. If you feel more comfortable thinking $n = 3$, that’s fine

for the moment, but mathematically it makes no difference, and as we shall see later, even when working with real-world, three-dimensional problems, it is often important to make use of higher-dimensional spaces.

On the other hand, an “instant of time” will always be represented by a single real number t . (There are mathematical situations that do require multi-dimensional time, but we shall not meet them here.) Thus, knowing the wind velocity at every point of space and at all instants of time means that we have a function V that associates to each (x, t) in $\mathbf{R}^n \times \mathbf{R}$ a vector $V(x, t)$ in \mathbf{R}^n , the wind velocity at x at time t . We will denote the n components of $V(x, t)$ by $V_1(x, t), \dots, V_n(x, t)$. (We will always assume that V is at least continuous and usually that it is even continuously differentiable.)

How should we model the path taken by a smoke particle? An ideal smoke particle is characterized by the fact that it “goes with the flow”, i.e., it is carried along by the wind. That means that if $x(t) = (x_1(t), \dots, x_n(t))$ is its location at a time t , then its velocity at time t will be the wind velocity at that point and time, namely $V(x(t), t)$. But the velocity of the particle at time t is $x'(t) = (x'_1(t), \dots, x'_n(t))$, where primes denote differentiation with respect to t , i.e., $x' = \frac{dx}{dt} = (\frac{dx_1}{dt}, \dots, \frac{dx_n}{dt})$.

So the path of a smoke particle will be a differentiable curve $x(t)$ in \mathbf{R}^n that satisfies the differential equation $x'(t) = V(x(t), t)$, or $\frac{dx}{dt} = V(x, t)$. If we write this in components, it reads $\frac{dx_i}{dt} = V_i(x_1(t), \dots, x_n(t), t)$, for $i = 1, \dots, n$, and for this reason it is often called a system of differential equations. Finally, if at a time t_0 we observe that the smoke particle is at the point x_0 in \mathbf{R}^n , then the “initial condition” $x(t_0) = x_0$ is also satisfied.

The page devoted to Chapter 1 in the Web Companion contains a QuickTime movie showing the wind field of a time-dependent two-dimensional system and the path traced out by a “smoke particle”. Figure 1.1 shows the direction field and a few such solution curves for an interesting and important one-dimensional ODE called the logistic equation.




Figure 1.1. The logistic equation.

For the logistic equation, the velocity field is given by $V(x, t) = cx(A-x)$. The vertical x -axis represents the size of some quantity, and the horizontal axis is the time, t . This equation models the growth of x in the presence of environmental constraints. The constant A is called the carrying capacity, and $c(A-x)$ is the “growth rate”. Note that the growth rate approaches zero as x approaches the carrying capacity. This equation is discussed in more detail in Section 2.7 on ecological models.

The combination of a differential equation, $\frac{dx}{dt} = V(x, t)$, and an initial condition, $x(t_0) = x_0$, is called an “initial value problem” (IVP), so the above informal prediction question for smoke particles can now be translated into a precise mathematical question: “What

can we say about the existence and uniqueness of solutions to such initial value problems?"

We will discuss this central question in detail below, along with important related questions such as how solutions of an IVP change as we vary the initial condition and the vector field. In order not to over-burden the exposition, we will leave many details of proofs to be worked out by the reader in exercises (with liberal hints). Fully detailed proofs can be found in the appendices and various references.

First let us make precise the definition of a solution of the above initial value problem: it is a differentiable map x of some open interval I containing t_0 into \mathbf{R}^n such that $x(t_0) = x_0$ and $x'(t) = V(x(t), t)$ for all t in I .

We first consider uniqueness. The vector field $V : \mathbf{R}^n \times \mathbf{R} \rightarrow \mathbf{R}^n$ is called *continuously differentiable* (or C^1) if all of its components $V_i(x_1, \dots, x_n, t)$ have continuous first partial derivatives with respect to x_1, \dots, x_n, t , and more generally V is called C^k if all partial derivatives of order k or less of its components exist and are continuous.

1.1.1. Uniqueness Theorem. *Let $V : \mathbf{R}^n \times \mathbf{R} \rightarrow \mathbf{R}^n$ be a C^1 time-dependent vector field on \mathbf{R}^n and let $x^1(t)$ and $x^2(t)$ be two solutions of the differential equation $\frac{dx}{dt} = V(x, t)$ defined on the same interval $I = (a, b)$ and satisfying the same initial condition, i.e., $x^1(t_0) = x^2(t_0)$ for some $t_0 \in I$. Then in fact $x^1(t) = x^2(t)$ for all $t \in I$.*

▷ **Exercise 1–1.** Show that continuity of V is **not** sufficient to guarantee uniqueness for an IVP. Hint: The classic example (with $n = 1$) is the initial value problem $\frac{dx}{dt} = \sqrt{x}$ and $x(0) = 0$. Show that for each $T > 0$, we get a distinct solution $x_T(t)$ of this IVP by defining $x_T(t) = 0$ for $t < T$ and $x_T(t) = \frac{1}{4}(t - T)^2$ for $t \geq T$.

But what about existence?

1.1.2. Local Existence Theorem. *Let $V : \mathbf{R}^n \times \mathbf{R} \rightarrow \mathbf{R}^n$ be a C^1 time-dependent vector field on \mathbf{R}^n . Given $p_0 \in \mathbf{R}^n$ and $t_0 \in \mathbf{R}$, there is a neighborhood O of p_0 and an $\epsilon > 0$ such that for every p*

in O there is a solution $x_p : (t_0 - \epsilon, t_0 + \epsilon) \rightarrow \mathbf{R}^n$ of the differential equation $\frac{dx}{dt} = V(x, t)$ satisfying the initial condition $x_p(t_0) = p$.

The proofs of existence and uniqueness have been greatly simplified over time, but understanding the details still requires nontrivial effort. Here we will sketch some of the most important ideas and constructs that go into the complete proof, but in order not to interrupt the flow of our exposition, we will defer the details to Appendix B. But even if you choose not to study these proofs now, we urge you to do so at some later time. We think you will find that these proofs are so elegant, and the ideas and constructions that enter into them are of such interest in their own right, that studying them is well worth the time and effort it requires.

We begin with a simple but very important reformulation of the ODE initial value problem $x'(s) = V(x(s), s)$ and $x(t_0) = x^0$. Namely, if we integrate both sides of the first of these equations from t_0 to t , we find that $x(t) = x^0 + \int_{t_0}^t V(x(s), s) ds$, and we refer to this equation as the integral form of the initial value problem. Note that by substituting $t = t_0$ in the integral form and by differentiating it, we get back the two original equations, so the integral form and the ODE form are equivalent. This suggests that we make the following definition.

1.1.3. Definition. Associated to each time-dependent vector field V on \mathbf{R}^n and $x^0 \in \mathbf{R}^n$, we define a mapping F^{V, x^0} that transforms a continuous function $x : I \rightarrow \mathbf{R}^n$ (where I is any interval containing t_0) to another such function $F^{V, x^0}(x) : I \rightarrow \mathbf{R}^n$ defined by $F^{V, x^0}(x)(t) = x^0 + \int_{t_0}^t V(x(s), s) ds$.

▷ **Exercise 1–2.** Show that any y of the form $F^{V, x^0}(x)$ satisfies the initial condition $y(t_0) = x^0$, and moreover y is continuously differentiable with derivative $y'(t) = V(x(t), t)$.

Recall that if f is any mapping, then a point in the domain of f such that $f(p) = p$ is called a *fixed point* of f . Thus we can rephrase the integral form of the initial value problem as follows:

1.1.4. Proposition. *A continuous map $x : I \rightarrow \mathbf{R}^n$ is a solution of the initial value problem $x'(t) = V(x(t), t)$, $x(t_0) = x^0$ if and only if x is a fixed point of F^{V, x^0} .*

Now if you have had some experience with fixed-point theorems, that should make your ears perk up a little. Not only are there some very general and powerful results for proving existence and uniqueness of fixed points of maps, but even better, there are nice algorithms for finding fixed points. One such algorithm is the so-called Method of Successive Approximations. (If you are familiar with Newton's Method for finding roots of equations, you will recognize that as a special case of successive approximations.) If we have a set X and a self-mapping $f : X \rightarrow X$, then to apply successive approximations, choose some "initial approximation" x_0 in X and then inductively define a sequence $x_{n+1} = f(x_n)$ of "successive approximations".

▷ **Exercise 1–3.** Suppose that X is a metric space, f is continuous, and that the sequence x_n of "successive approximations" converges to a limit p . Show that p is a fixed point of f .

But is there really any hope that we can use successive approximations to find solutions of ODE initial value problems? Let us try a very simple example. Consider the (time-independent) vector field V on \mathbf{R}^n defined by $V(x, t) = x$. It is easy to check that the unique solution with $x(0) = x^0$ is given by $x(t) = e^t x^0$. Let's try using successive approximations to find a fixed point of F^{V, x^0} . For our initial approximation we choose the constant function $x_0(t) = x^0$, and following the general successive approximation prescription, we define x_n inductively by $x_{n+1} = F^{V, x^0}(x_n)$, i.e., $x_{n+1}(t) = x^0 + \int_0^t x_n(s) ds$.

▷ **Exercise 1–4.** Show by induction that $x_n(t) = P_n(t)x^0$, where $P_n(t)$ is the polynomial of degree n obtained by truncating the power series for e^t (i.e., $\sum_{j=0}^n \frac{1}{j!} t^j$).

That is certainly a hopeful sign, and while one swallow may not make a spring, it should give us hope that a careful analysis of successive approximations might lead to a proof of the existence and uniqueness theorems for an arbitrary vector field V . This is in fact the case, but

we will not give further details here. Instead we refer to Appendix B where you will find a complete proof.

1.1.5. Remark. We give a minor technical point. The argument in Appendix B only gives a local uniqueness theorem. That is, it shows that if $x^1 : (a, b) \rightarrow \mathbf{R}^n$ and $x^2 : (a, b) \rightarrow \mathbf{R}^n$ are two solutions of the same ODE, then if x^1 and x^2 agree at a point, then they also agree in a neighborhood of that point, so that the set of points in (a, b) where they agree is open. But since solutions are by definition continuous, the set of points where x^1 and x^2 agree is also a closed subset of (a, b) , and since intervals are connected, it then follows that x^1 and x^2 agree on all of (a, b) .

1.1.6. Remark. The existence and uniqueness theorems tell us that for a given initial condition x^0 we can solve our initial value problem (uniquely) for a short time interval. The next question we will take up is for just how long we can “follow a smoke particle”. One important thing to notice is the uniformity of the ϵ in the existence theorem—not only do we have a solution for each initial condition, but moreover given any p_0 in \mathbf{R}^n , we can find a **fixed** interval $I = (t_0 - \epsilon, t_0 + \epsilon)$ such that a solution with initial condition p exists on the whole interval I for all initial conditions sufficiently close to p_0 . Still, this may be less than what you had hoped and expected. You may have thought that for each initial condition p in \mathbf{R}^n we should have a solution $x_p : \mathbf{R} \rightarrow \mathbf{R}^n$ of the differential equation with $x_p(t_0) = p$. But such a global existence theorem is too much to expect. For example, taking $n = 1$ again, consider the differential equation $\frac{dx}{dt} = x^2$ with the initial condition $x(0) = x_0$. An easy calculation shows that the unique solution is $x(t) = \frac{x_0}{1 - x_0 t}$. Note that, for each initial condition x_0 , this solution “blows up” at time $T = \frac{1}{x_0}$, and by the Uniqueness Theorem, no solution can exist for a time greater than T .

But, you say, a particle of smoke will never go off to infinity in a finite amount of time! Perhaps the smoke metaphor isn’t so good after all. The answer is that a real, physical wind field has bounded velocity, and it isn’t hard to show that in this case we do indeed have

global existence. You will even prove something a lot stronger in a later exercise.

What **can** be said is that for each initial condition, p , there is a unique “maximal” solution of the differential equation with that initial condition. But before discussing this, we are going to make a simplification and restrict our attention to time-**independent** vector fields (which we shall simply call vector fields). That may sound like a tremendous loss of generality, but in fact it is no loss of generality at all!

▷ **Exercise 1–5.** Let $V(x, t) = (V_1(x, t), \dots, V_n(x, t))$ be a time-dependent vector field in \mathbf{R}^n , and define an associated time independent vector field \tilde{V} in \mathbf{R}^{n+1} by $\tilde{V}(y) = (V_1(y), \dots, V_n(y), 1)$. Show that $y(t) = (x(t), f(t))$ is a solution of the differential equation $\frac{dy}{dt} = \tilde{V}(y)$ if and only if $f(t) = t + c$ and $x(t)$ is a solution of $\frac{dx}{dt} = V(x, t + c)$. Deduce that if $y(t) = (x(t), f(t))$ solves the IVP $\frac{dy}{dt} = \tilde{V}(y)$, $y(t_0) = (x_0, t_0)$, then $x(t)$ is a solution of the IVP $\frac{dx}{dt} = V(x, t)$, $x(t_0) = x_0$.

This may look like a swindle. We don’t seem to have done much besides changing the name of the original time variable t to x_{n+1} and considering it a space variable; that is, we switched to space-time notation. But the real change is in making the velocity an $(n + 1)$ -vector too and setting the last component identically equal to one. In any case this is a true reduction of the time-dependent case to the time-independent case, and as we shall see, that is quite important, since time-independent differential equations have special properties not shared with time-dependent equations that can be used to simplify their study. Time-independent differential equations are usually referred to as *autonomous*, and time-dependent ones as *nonautonomous*. Here is one of the special properties of autonomous systems.

1.1.7. Proposition. *If $x : (a, b) \rightarrow \mathbf{R}^n$ is any solution of the autonomous differentiable equation $\frac{dx}{dt} = V(x)$ and $t_0 \in \mathbf{R}$, then $y : (a + t_0, b + t_0) \rightarrow \mathbf{R}^n$ defined by $y(t) = x(t - t_0)$ is also a solution of the same equation.*

▷ **Exercise 1–6.** Prove the above proposition.

Consequently, when considering the IVP for an autonomous differentiable equation, we can assume that $t_0 = 0$. For if $x(t)$ is a solution with $x(0) = p$, then $x(t - t_0)$ will be a solution with $x(t_0) = p$.

1.1.8. Remark. There is another trick that allows us to reduce the study of higher-order differential equations to the case of first-order equations. We will consider this in detail later, but here is a short preview. Consider the second-order differential equation: $\frac{d^2x}{dt^2} = f(x, \frac{dx}{dt}, t)$. Introduce a new variable v (the velocity) and consider the following related system of first-order equations: $\frac{dx}{dt} = v$ and $\frac{dv}{dt} = f(x, v, t)$. It is pretty obvious there is a close relation between curves $x(t)$ satisfying $x''(t) = f(x(t), x'(t), t)$ and pairs of curves $x(t)$, $v(t)$ satisfying $x'(t) = v(t)$ and $v'(t) = f(x(t), v(t), t)$.

▷ **Exercise 1–7.** Define the notion of an initial value problem for the above second-order differential equation, and write a careful statement of the relation between solutions of this initial value problem and the initial value problem for the related system of first-order differential equations.

We will now look more closely at the uniqueness question for solutions of an initial value problem. The answer is summed up succinctly in the following result.

1.1.9. Maximal Solution Theorem. *Let $\frac{dx}{dt} = V(x)$ be an autonomous differential equation in \mathbf{R}^n and p any point of \mathbf{R}^n . Among all solutions $x(t)$ of the equation that satisfy the initial condition $x(0) = p$, there is a maximum one, σ_p , in the sense that any solution of this IVP is the restriction of σ_p to some interval containing zero.*

▷ **Exercise 1–8.** If you know about connectedness, you should be able to prove this very easily. First, using the local uniqueness theorem, show that any two solutions agree on their overlap, and then define σ_p to be the union of all solutions.

Henceforth whenever we are considering some autonomous differential equation, σ_p will denote this maximal solution curve with initial

condition p . The interval on which σ_p is defined will be denoted by $(\alpha(p), \omega(p))$, where of course $\alpha(p)$ is either $-\infty$ or a negative real number and $\omega(p)$ is either ∞ or a positive real number.

As we have seen, a maximal solution σ_p need **not** be defined on all of \mathbf{R} , and it is important to know just how the solution “blows up” as t approaches a finite endpoint of its interval of definition. *A priori* it might seem that the solution could remain in some bounded region, but it is an important fact that this is impossible—if $\omega(p)$ is finite, then the reason the solution cannot be continued past $\omega(p)$ is simply that it escapes to infinity as t approaches $\omega(p)$.

1.1.10. No Bounded Escape Theorem. *If $\omega(p) < \infty$, then*

$$\lim_{t \rightarrow \omega(p)} \|\sigma_p(t)\| = \infty,$$

and similarly, if $\alpha(p) > -\infty$, then

$$\lim_{t \rightarrow \alpha(p)} \|\sigma_p(t)\| = \infty.$$

▷ **Exercise 1–9.** Prove the No Bounded Escape Theorem. (Hint: If $\lim_{t \rightarrow \omega(p)} \|\sigma_p(t)\| \neq \infty$, then by Bolzano-Weierstrass there would be a sequence t_k converging to $\omega(p)$ from below, such that $\sigma_p(t_k) \rightarrow q$. Then use the local existence theorem around q to show that you could extend the solution beyond $\omega(p)$. Here is where we get to use the fact that there is a neighborhood O of q such that a solution exists with any initial condition q' in O and defined **on the whole interval** $(-\epsilon, \epsilon)$. For k sufficiently large, we will have both $\sigma_p(t_k)$ in O and $t_k > \omega - \epsilon$, which quickly leads to a contradiction.)

Here is another special property of autonomous systems.

▷ **Exercise 1–10.** Show that the images of the σ_p partition \mathbf{R}^n into disjoint smooth curves (the “streamlines” of smoke particles). These curves are referred to as the *orbits* of the ODE. (Hint: If $x(t)$ and $\xi(t)$ are two solutions of the same autonomous ODE and if $x(t_0) = \xi(t_1)$, then show that $x(t_0 + s) = \xi(t_1 + s)$.)

1.1.11. Definition. A C^1 vector field $V : \mathbf{R}^n \rightarrow \mathbf{R}^n$ (and also the autonomous differential equation $\frac{dx}{dt} = V(x)$) is called *complete* if $\alpha(p) = -\infty$ and $\omega(p) = \infty$ for all p in \mathbf{R}^n . In this case, for each $t \in \mathbf{R}$ we define a map $\phi_t : \mathbf{R}^n \rightarrow \mathbf{R}^n$ by $\phi_t(p) = \sigma_p(t)$. The mapping $t \mapsto \phi_t$ is called the *flow* generated by the differential equation $\frac{dx}{dt} = V(x)$.

1.1.12. Remark. Using our smoke particle metaphor, the meaning of ϕ_t can be explained as follows: if a puff of smoke occupies a region U at a given time, then t units of time later it will occupy the region $\phi_t(U)$. Note that ϕ_0 is clearly the identity mapping of \mathbf{R}^n .

▷ **Exercise 1–11.** Show that the ϕ_t satisfy $\phi_{t_1+t_2} = \phi_{t_1} \phi_{t_2}$, so that in particular $\phi_{-t} = \phi_t^{-1}$. In other words, the flow generated by a complete, autonomous vector field is a homomorphism of the additive group of real numbers into the group of bijective self-mappings of \mathbf{R}^n .

In the next section we will see that $(t, p) \mapsto \phi_t(p)$ is jointly continuous, so that the ϕ_t are homeomorphisms of \mathbf{R}^n . Later (in Appendix F) we will also see that if the vector field V is C^r , then $(t, p) \mapsto \phi_t(p)$ is also C^r , so that the flow generated by a complete, autonomous, C^r differential equation $\frac{dx}{dt} = V(x)$ is a homomorphism of \mathbf{R} into the group of C^r diffeomorphisms of \mathbf{R}^n . The branch of mathematics that studies the properties of flows is called *dynamical systems theory*.

• **Example 1–1. Constant Vector Fields.** The simplest examples of autonomous vector fields in \mathbf{R}^n are the constant vector fields $V(x) = v$, where v is some fixed vector in \mathbf{R}^n . The maximal solution curve with initial condition p of $\frac{dx}{dt} = v$ is clearly the linearly parametrized straight line $\sigma_p : \mathbf{R} \rightarrow \mathbf{R}^n$ given by $\sigma_p(t) = p + tv$, and it follows that these vector fields are complete. The corresponding flow ϕ_t is given by $\phi_t(p) = p + tv$, so for obvious reasons these are called constant velocity flows. In words, ϕ_t is translation by the vector tv , and indeed these flows are precisely the one-parameter subgroups of the group of translations of \mathbf{R}^n .

• **Example 1–2. Exponential Growth.** An important complete vector field in \mathbf{R} is the linear map $V(x) = kx$. The maximal solution curves of $\frac{dx}{dt} = kx$ are again easy to write down explicitly, namely

$\sigma_p(t) = e^{kt}p$; i.e., in this case the flow map ϕ_t is just multiplication by e^{kt} .

• **Example 1–3. Harmonic Oscillator.** If we start from the Harmonic Oscillator Equation, $\frac{d^2x}{dt^2} = -x$, and use the trick above to rewrite this second-order equation as a first-order system, we end up with the linear system in \mathbf{R}^2 : $\frac{dx}{dt} = -y$, $\frac{dy}{dt} = x$. In this case the maximal solution curve $\sigma_{(x_0, y_0)}(t)$ can again be given explicitly, namely $\sigma_{(x_0, y_0)}(t) = (x_0 \cos(t) - y_0 \sin(t), x_0 \sin(t) + y_0 \cos(t))$, so that now ϕ_t is rotation in the plane through an angle t . It is interesting to observe that this can be considered a special case of (a slightly generalized form of) the preceding example. Namely, if we identify \mathbf{R}^2 with the complex plane \mathbf{C} in the standard way (i.e., $a+ib := (a, b)$) and write $z = (x, y) = x + iy$, $z_0 = (x_0, y_0) = x_0 + iy_0$, then since $iz = i(x + iy) = -y + ix = (-y, x)$, we can rewrite the above first-order system as $\frac{dz}{dt} = iz$, which has the solution $z(t) = e^{it}z_0$. Of course, multiplication by e^{it} is just rotation through an angle t .

It is very useful to have conditions on a vector field V that will guarantee its completeness.

▷ **Exercise 1–12.** Show that $\|\sigma_p(t) - p\| \leq \int_0^t \|V(\sigma_p(t))\| dt$. Use this and the No Bounded Escape Theorem to show that $\frac{dx}{dt} = V(x)$ is complete provided that V is bounded (i.e., $\sup_{x \in \mathbf{R}^n} \|V(x)\| < \infty$).

▷ **Exercise 1–13.** A vector field V may be complete even if it is not bounded, provided that it doesn't "grow too fast". Let $B(r) = \sup_{\|x\| < r} \|V(x)\|$. Show that if $\int_1^\infty \frac{dr}{B(r)} = \infty$, then V is complete. Hint: How long does it take $\sigma_p(t)$ to get outside a ball of radius R ?

▷ **Exercise 1–14.** If a vector field is not complete, then given any positive ϵ , there exist points p where either $\alpha(p) > -\epsilon$ or $\omega(p) < \epsilon$.

1.2. Euler's Method

Only a few rather special initial value problems can be solved in closed form using standard elementary functions. For the general case it is

necessary to fall back on constructing an approximate solution numerically with the aid of a computer. But what algorithm should we use to program the computer? A natural first guess is successive approximations. But while that is a powerful theoretical tool for studying the general properties of initial value problems (and in particular for proving existence and uniqueness), it does not lead to an efficient algorithm for constructing numerical solutions.

In fact there is no one simple answer to the question of what numerical algorithm to use for solving ODEs, for there is no single method that is “best” in all situations. While there are integration routines (such as the popular fourth-order Runge-Kutta integration) that are fast and accurate when used with many of the equations one meets, there are many situations that require a more sophisticated approach. Indeed, this is still an active area of research, and there are literally dozens of books on the subject. Later, in the chapter on numerical methods, we will introduce you to many of the subtleties of this topic, but here we only want to give you a quick first impression by describing one of the oldest numerical approaches to solving an initial value problem, the so-called “Euler Method”. While rarely an optimal choice, it is intuitive, simple, and effective for some purposes. It is also the prototype for the design and analysis of more sophisticated algorithms. This makes it an excellent place to become familiar with the basic concepts that enter into the numerical integration of ODE.

In what follows we will suppose that $\mathbf{f}(t, \mathbf{y})$ is a C^1 time-dependent vector field on \mathbf{R}^d , t_o in \mathbf{R} and \mathbf{y}_o in \mathbf{R}^d . We will denote by $\sigma(\mathbf{f}, \mathbf{y}_o, t_o, t)$ the solution operator taking this data to the values $\mathbf{y}(t)$ of the maximal solution of the associated initial value problem. By definition, $\mathbf{y}(t)$ is the function defined on a maximal interval $I = [t_o, t_o + T_*)$, with $0 < T_* \leq \infty$, satisfying the differential equation $\frac{d\mathbf{y}}{dt} = \mathbf{f}(t, \mathbf{y})$ and the initial condition $\mathbf{y}(t_o) = \mathbf{y}_o$. The goal in the numerical integration of ODE is to devise effective methods for approximating such a solution $\mathbf{y}(t)$ on an interval $I = [t_o, t_o + T]$ for $T < T_*$. The strategy that many methods use is to discretize the interval I using $N + 1$ equally spaced gridpoints $t_n := t_o + nh, n = 0, \dots, N$ with $h = \frac{T}{N}$ so that $t_0 = t_o$ and $t_N = t_o + T$ and then use some algorithm

to define values $\mathbf{y}_0, \dots, \mathbf{y}_N$ in \mathbf{R}^d , in such a way that when N is large, each \mathbf{y}_n is close to the corresponding $\mathbf{y}(t_n)$. The quantity $\max_{0 \leq n \leq N} \|\mathbf{y}(t_n) - \mathbf{y}_n\|$ is called the *global error* of the algorithm on the interval. If the global error converges to zero as N tends to infinity (for every choice of \mathbf{f} satisfying some Lipschitz condition, t_o , \mathbf{y}_o , and $T < T_*$), then we say that we have a *convergent algorithm*. Euler's Method is a convergent algorithm of this sort.

One common way to construct the algorithm that produces the values y_1, \dots, y_N uses a recursion based on a so-called (one-step) "stepping procedure". This is a discrete approximate solution operator, $\Sigma(\mathbf{f}, \mathbf{y}_n, t_n, h)$, having as inputs

- 1) a time-dependent vector field \mathbf{f} on \mathbf{R}^d ,
- 2) a time t_n in \mathbf{R} ,
- 3) a value \mathbf{y}_n in \mathbf{R}^d corresponding to the initial time, and
- 4) a "time-step" h in \mathbf{R}

and as output a point of \mathbf{R}^d that approximates the solution of the initial value problem $\mathbf{y}' = f(t, \mathbf{y})$, $\mathbf{y}(t_i) = \mathbf{y}_i$ at $t_i + h$ well when h is small. (More precisely, the so-called "local truncation error", $\|\sigma(\mathbf{f}, \mathbf{y}(t_n), t_n, t_n + h) - \Sigma(\mathbf{f}, \mathbf{y}(t_n), t_n, h)\|$, should approach zero at least superlinearly in the time-step h .) Given such a stepping procedure, the approximations \mathbf{y}_n of the $\mathbf{y}(t_n)$ are defined recursively by $\mathbf{y}_{n+1} = \Sigma(\mathbf{f}, \mathbf{y}_n, t_n, h)$. Numerical integration methods that use discrete approximations of derivatives defining the vector field \mathbf{f} to obtain the operator Σ are referred to as finite difference methods.

1.2.1. Remark. Notice that there will be two sources that contribute to the global error, $\|\mathbf{y}(t_n) - \mathbf{y}_n\|$. First, at each stage of the recursion there will be an additional local truncation error added to what has already accumulated up to that point. Moreover, because the recursion uses \mathbf{y}_n rather than $\mathbf{y}(t_n)$, after the first step there will be an additional error that includes accumulated local truncation errors, in addition to amplification or attenuation of these errors by the method. (In practice there is a third source of error, namely machine round-off error from using floating-point arithmetic. Since these are

amplified or attenuated in the same manner as truncation errors, we will often consolidate them and pretend that our computers do precise real number arithmetic, but there are situations where it is important to take it into consideration.)

For Euler's Method the stepping procedure is particularly simple and natural. It is defined by $\Sigma_E(\mathbf{f}, \mathbf{y}_n, t_n, h) := \mathbf{y}_n + h \mathbf{f}(t_n, \mathbf{y}_n)$. It is easy to see why this is a good choice. If as above we denote $\sigma(\mathbf{f}, \mathbf{y}_n, t_n, t)$ by $\mathbf{y}(t)$, then by Taylor's Theorem,

$$\begin{aligned} \mathbf{y}(t_n + h) &= \mathbf{y}(t_n) + h \mathbf{y}'(t_n) + O(h^2) \\ &= \mathbf{y}_n + h \mathbf{f}(t_n, \mathbf{y}_n) + O(h^2) \\ &= \Sigma_E(\mathbf{f}, \mathbf{y}_n, t_n, h) + O(h^2), \end{aligned}$$

so that $\|\sigma(\mathbf{f}, \mathbf{y}_n, t_n, t_n + h) - \Sigma_E(\mathbf{f}, \mathbf{y}_n, t_n, h)\|$, the local truncation error for Euler's Method, does go to zero quadratically in h . When we partition $[t_o, t_o + T]$ into N equal parts, $h = \frac{T}{N}$, each step in the recursion for computing \mathbf{y}_n will contribute a local truncation error that is $O(h^2) = O(\frac{1}{N^2})$. Since there are N steps in the recursion and at each step we add $O(\frac{1}{N^2})$ to the error, this suggests that the global error will be $O(\frac{1}{N})$ and hence will go to zero as N tends to infinity. However, because of the potential amplification of prior errors, this is not a complete proof that Euler's Method is convergent, and we will put off the details of the rigorous argument until the chapter on numerical methods.

▷ **Exercise 1–15.** Show that Euler's Method applied to the initial value problem $\frac{dy}{dt} = y$ with $y(0) = 1$ gives $\lim_{N \rightarrow \infty} (1 + \frac{T}{N})^N = e^T$. For $T = 1$ and $N = 2$, show that the global error is indeed greater than the sum of the two local truncation errors.

1.3. Stationary Points and Closed Orbits

We next describe certain special types of solutions of a differential equation that play an important role in the description and analysis of the global behavior of its flow. For generality we will also consider the case of time-dependent vector fields, but these solutions are really most important in the study of autonomous equations.

If a constant map $\sigma : I \rightarrow \mathbf{R}^n$, $\sigma(t) = p$ for all $t \in I$, is a solution of the equation $\frac{dx}{dt} = V(x, t)$, then $V(p, t) = \sigma'(t) = 0$ for all t , and conversely this implies $\sigma(t) \equiv p$ is a solution. In particular, in the autonomous case, the maximal solution σ_p is a constant map if and only if $V(p) = 0$. Such points p are of course called zeros of the time-independent vector field V , but because of their great importance they have also been given many more aliases, including critical point, singularity, stationary point, rest point, equilibrium point, and fixed point.

A related but more interesting type of solution of $\frac{dx}{dt} = V(x, t)$ is a so-called *closed orbit*, also referred to as a *periodic solution*. To define these, we start with an arbitrary solution σ defined on the whole real line. A real number T is called a *period* of σ if $\sigma(t) = \sigma(t + T)$ for all $t \in \mathbf{R}$, and we will denote by $\mathbf{Per}(\sigma)$ the set of all periods of σ . Of course 0 is always a period of σ , and one possibility is that it is the only period, in which case σ is called a nonperiodic orbit. At the other extreme, σ is a constant solution if and only if every real number is a period of σ .

What other possibilities are there for $\mathbf{Per}(\sigma)$? To answer that, let us look at some obvious properties of the set of periods. First, $\mathbf{Per}(\sigma)$ is clearly a closed subset of \mathbf{R} —this follows from the continuity of σ . Secondly, if T_1 and T_2 are both periods of σ , then $\sigma(t + (T_1 - T_2)) = \sigma((t - T_2) + T_1) = \sigma(t - T_2) = \sigma(t - T_2 + T_2) = \sigma(t)$, so we see that the difference of any two periods is another period. Thus $\mathbf{Per}(\sigma)$ is a closed subgroup of the group of real numbers under addition. But the structure of such groups is well known.

1.3.1. Proposition. *If Γ is a closed subgroup of \mathbf{R} , then either $\Gamma = \mathbf{R}$, or $\Gamma = \{0\}$, or else there is a smallest positive element γ in Γ and Γ consists of all integer multiples of γ .*

▷ **Exercise 1–16.** Prove this proposition. (Hint: If Γ is nontrivial, then the set of positive elements of Γ is nonempty and hence has a greatest lower bound γ which is in Γ since Γ is closed. If $\gamma = 0$, show that Γ is dense in \mathbf{R} and hence it is all of \mathbf{R} . If $\gamma \neq 0$, it is the smallest positive element of Γ . In this case, if $n \in \Gamma$, then dividing n

by γ gives $n = q\gamma + r$ with $0 \leq r < \gamma$. Show that the remainder, r , must be zero.)

A solution σ is called periodic if it is nonconstant and has a nontrivial period, so that by the proposition all its periods are multiples of a smallest positive period γ , called the *prime period* of σ .

A real number T is called a period of the time-dependent vector field V if $V(x, t) = V(x, t + T)$ for all $t \in T$ and $x \in \mathbf{R}$. A repeat of the arguments above show that the set $\mathbf{Per}(V)$ of all periods of V is again a closed subgroup of \mathbf{R} , so again there are three cases: 1) $\mathbf{Per}(V) = \mathbf{R}$, i.e., V is time-independent, 2) $\mathbf{Per}(V) = \{0\}$, i.e., V is nonperiodic, or 3) there is a smallest positive element T_0 of $\mathbf{Per}(V)$ (the prime period of V) and $\mathbf{Per}(V)$ consists of all integer multiples of this prime period.

▷ **Exercise 1–17.** Show that if T is a period of the time-dependent vector field V and σ is a solution of $\frac{dx}{dt} = V(x, t)$, then T is also a period of σ provided there exists a real number t_1 such that $\sigma(t_1) = \sigma(t_1 + T)$. (Hint: Use the uniqueness theorem.)

Note the following corollary: in the autonomous case, if an orbit σ “comes back and meets itself”, i.e., if there are two distinct times t_1 and t_2 such that $\sigma(t_1) = \sigma(t_2)$, then σ is a periodic orbit and $t_2 - t_1$ is a period. For this reason, periodic solutions of autonomous ODEs are also referred to as closed orbits. Another way of stating this same fact is as follows:

1.3.2. Proposition. *Let ϕ_t be the flow generated by a complete, autonomous ODE, $\frac{dx}{dt} = V(x)$. A necessary and sufficient condition for the maximum solution curve σ_p with initial condition p to be periodic with period T is that p be a fixed point of ϕ_T .*

• **Example 1–4.** For the harmonic oscillator system in \mathbf{R}^2 : $\frac{dx}{dt} = -y$, $\frac{dy}{dt} = x$, we have seen that the solution with initial condition (x_0, y_0) is $x(t) = x_0 \cos(t) - y_0 \sin(t)$, $y(t) = x_0 \sin(t) + y_0 \cos(t)$. Clearly the origin is a stationary point, and every other solution is periodic with the same prime period 2π .

1.3.3. Remark. The ODEs modeling many physical systems have periodic orbits, and each such orbit defines a physical “clock” whose natural unit is the prime period of the orbit. We simply choose a configuration of the system that lies on this periodic orbit and tick off the successive recurrences of that configuration to “tell time”. The resolution to which before and after can be distinguished with such a clock is limited to approximately the prime period of the orbit. There seems to be no limit to the benefits of ever more precise chronometry—each time a clock has been constructed with a significantly shorter period, it has opened up new technological possibilities. Humankind has always had a 24-hour period clock provided by the rotation of the earth on its axis, but it was only about four hundred years ago that reasonably accurate clocks were developed with a period in the 1-second range. In recent decades the resolution of clocks has increased dramatically. For example, the fundamental clock period for the computer on which we are writing this text is about 0.4×10^{-9} seconds. The highest resolution (and most accurate) of current clocks is the cesium vapor atomic clocks used by international standards agencies. These have a period of about 10^{-11} seconds (with a drift error of about 1 second in 300,000 years!). This means that if two events occur only one hundred billionth of a second apart, one of these clocks can in principle tell which came first.

1.4. Continuity with Respect to Initial Conditions

We consider next how the maximal solutions σ_p of a first-order ODE $\frac{dx}{dt} = V(x)$ depends on the initial condition p . Eventually we will see that this dependence is as smooth as the vector field V , but as a first step we will content ourselves with proving just continuity. The argument rests on a simple but important general principle called Gronwall’s Inequality.

1.4.1. Gronwall’s Inequality. *Let $u : [0, T) \rightarrow [0, \infty)$ be a continuous, nonnegative, real-valued function and assume that $u(t) \leq U(t) := C + K \int_0^t u(s) ds$ for certain constants $C \geq 0$ and $K > 0$. Then $u(t) \leq Ce^{Kt}$.*

▷ **Exercise 1–18.** Prove Gronwall's Inequality.

Hint: Since $u \leq U$, it is enough to show that $U(t) \leq Ce^{Kt}$, or equivalently that $e^{-Kt}U(t) \leq C$, and since $U(0) = C$, it will suffice to show that $e^{-Kt}U(t)$ is nonincreasing, i.e., that $(e^{-Kt}U(t))' \leq 0$. Since $(e^{-Kt}U(t))' = e^{-Kt}(U'(t) - KU)$ and $U' = Ku$, this just says that $Ke^{-Kt}(u - U) \leq 0$.

1.4.2. Theorem on Continuity w.r.t. Initial Conditions.

Let V be a C^1 vector field on \mathbf{R}^n and let $\sigma_p(t)$ denote the maximal solution curve of $\frac{dx}{dt} = V(x)$ with initial condition p . Then as q tends to p , $\sigma_q(t)$ approaches $\sigma_p(t)$, and the convergence is uniform for t in any bounded interval I on which σ_p is defined.

Proof. We have seen that $\sigma_p(t) = p + \int_0^t V(\sigma_p(s), s) ds$, and it follows that $\|\sigma_p(t) - \sigma_q(t)\| \leq \|p - q\| + \int_0^t \|V(\sigma_p(s), s) - V(\sigma_q(s), s)\| ds$. On the other hand, it is proved in Appendix A that on any bounded set (and in particular on a bounded neighborhood of $\sigma_p(I) \times I$) V satisfies a Lipschitz condition $\|V(x, t) - V(y, t)\| \leq K \|x - y\|$, so it follows that $\|\sigma_p(t) - \sigma_q(t)\| \leq \|p - q\| + K \int_0^t \|\sigma_p(s) - \sigma_q(s)\| ds$. It now follows from Gronwall's Inequality that $\|\sigma_p(t) - \sigma_q(t)\| \leq \|p - q\| e^{Kt}$. ■

1.4.3. Remark. For the differential equation $\frac{dx}{dt} = kx$, the maximal solution is $\sigma_p(t) = e^{kt}p$, so $\|\sigma_p(t) - \sigma_q(t)\| = e^{kt} \|p - q\|$. Thus if k is positive, then any two solutions diverge from each other exponentially fast, while if k is negative, all solutions approach the origin (and hence each other) exponentially fast.

But continuity with respect to initial conditions is not the whole story.

1.4.4. Theorem on Smoothness w.r.t. Initial Conditions.

Let V be a C^r vector field on \mathbf{R}^n , $r \geq 1$, and let $\sigma_p(t)$ denote the maximal solution curve of $\frac{dx}{dt} = V(x)$ with initial condition p . Then the map $(p, t) \mapsto \sigma_p(t)$ is C^r .

The proof of this theorem is one of the most difficult in elementary ODE theory, and we have deferred it to Appendix F.

Let $V : \mathbf{R}^n \times \mathbf{R}^k \rightarrow \mathbf{R}^n$ be a smooth function. Then to each α in \mathbf{R}^k we can associate a vector field $V(\cdot, \alpha)$ on \mathbf{R}^n , defined by $x \mapsto V(x, \alpha)$. For this reason it is customary to consider V as a “vector field on \mathbf{R}^n depending on a parameter α in \mathbf{R}^k ”. It is often important to know how solutions of $\frac{dx}{dt} = V(x, \alpha)$ depend on the parameter α , and this is answered by the following theorem.

1.4.5. Theorem on Smoothness w.r.t. Parameters. *Let $V : \mathbf{R}^n \times \mathbf{R}^k \rightarrow \mathbf{R}^n$ be a C^r map, $r > 1$, and let σ_p^α denote the maximum solution curve of $\frac{dx}{dt} = V(x, \alpha)$ with initial condition p . Then the map $(p, \alpha, t) \mapsto \sigma_p^\alpha(t)$ is C^r .*

▷ **Exercise 1–19.** Deduce this from the Theorem on Smoothness w.r.t. Initial Conditions. Hint: This is another one of those cute reduction arguments that this subject is full of. The idea is to consider the vector field \tilde{V} on $\mathbf{R}^n \times \mathbf{R}^k$ defined by $\tilde{V}(x, \alpha) = (V(x, \alpha), 0)$ and to note that its maximal solution with initial condition (p, α) is $t \mapsto (\sigma_p^\alpha(t), \alpha)$.

You may have noticed an ambiguity inherent in our use of σ_p to denote the maximal solution curve with initial condition p of a vector field V . After all, this maximal solution clearly depends on V as well as on p , so let us now be more careful and denote it by σ_p^V . Of course, this immediately raises the question of just how σ_p^V depends on V . If V changes just a little, does it follow that σ_p^V also does not change by much? If we return to our smoke particle in the wind metaphor, then this seems reasonable; if we make a tiny perturbation of the direction and speed of the wind at every point, it seems that the path of a smoke particle should not be grossly different. This intuition is correct, and all that is required to prove it is another tricky application of Gronwall’s Inequality.

1.4.6. Theorem on the Continuity of σ_p^V w.r.t. V . *Let V be a C^1 time-dependent vector field on \mathbf{R}^n and let K be a Lipschitz constant for V , in the sense that $\|V(x, t) - V(y, t)\| \leq K \|x - y\|$ for all x, y , and t . If W is another C^1 time-dependent vector field on \mathbf{R}^n such that $\|V(x, t) - W(x, t)\| \leq \epsilon$ for all x and t , then $\|\sigma_p^V(t) - \sigma_p^W(t)\| \leq \frac{\epsilon}{K} (e^{Kt} - 1)$.*

▷ **Exercise 1–20.** Prove the above theorem. Hint: If we define $u(t) = \|\sigma_p^V(t) - \sigma_p^W(t)\| + \frac{\epsilon}{K}$, then the conclusion may be written as $u(t) \leq \frac{\epsilon}{K}e^{Kt}$, which follows from Gronwall’s Inequality provided we can prove $u(t) \leq \frac{\epsilon}{K} + K \int_0^t u(s) ds$. To show that, start from $u(t) - \frac{\epsilon}{K} = \|\sigma_p^V(t) - \sigma_p^W(t)\| \leq \int_0^t \|V(\sigma_p^V(s)) - W(\sigma_p^W(s))\| ds$ and use

$$\begin{aligned} \|V(\sigma_p^V(s)) - W(\sigma_p^W(s))\| &\leq \|V(\sigma_p^V(s)) - V(\sigma_p^W(s))\| \\ &\quad + \|V(\sigma_p^W(s)) - W(\sigma_p^W(s))\| \\ &\leq (Ku(s) - \epsilon) + \epsilon = Ku(s). \end{aligned}$$

1.5. Chaos—Or a Butterfly Spoils Laplace’s Dream

L’état présent du système de la Nature est évidemment une suite de ce qu’elle était au moment précédent et, si nous concevons une intelligence qui, pour un instant donné, embrasse tous les rapports des êtres de cet Univers, elle pourra déterminer pour un temps quelconque pris dans le passé ou dans l’avenir la position respective, les motions et généralement toutes les affections de ces êtres. . .

—Pierre Simon de Laplace, 1773¹

The so-called “scientific method” is a loosely defined iterative process of experimentation, induction, and deduction with the goal of deriving general “laws” for describing various aspects of reality. Prediction plays a central role in this enterprise. During the period of discovery and research, comparing experiments against predictions helps eliminate erroneous preliminary versions of a theory and conversely can provide confirming evidence when a theory is correct. And when a theory finally has been validated, its predictive power can lead to valuable new technologies. In the physical sciences, the laws frequently take the form of differential equations (of just the sort we have been

¹The current state of Nature is evidently a consequence of what it was in the preceding moment, and if we conceive of an intelligence that at a given moment knows the relations of all things of this Universe, it could then tell the positions, motions and effects of all of these entities at any past or future time. . .

considering) that model the time-evolution of various real-world processes. So it should not be surprising that the sort of issues that we have just been discussing have important practical and philosophical ramifications when it comes to evaluating and interpreting the predictive power of such laws, and indeed some of the above theorems were developed for just such reasons.

At first glance, it might appear that theory supports Laplace's ringing deterministic manifesto quoted above. But if we examine matters with more care, it becomes evident that, while making dependable predictions might be possible for a god who could calculate with infinite precision and who knew the laws with perfect accuracy, for any lesser beings there are severe problems not only in practice but even in principle.

First let us look at the positive side of things. In order to make reliable predictions based on a differential equation $\frac{dx}{dt} = V(x)$, at least the following two conditions must be satisfied:

- 1) There should be a unique solution for each initial condition, and it should be defined for all $t \in \mathbf{R}$.
- 2) This solution should depend continuously on the initial condition and also on the vector field V .

Initial value problems that satisfy these two conditions are often referred to as "well-posed" problems.

The importance of the first condition is obvious, and we will not say more about it. The second is perhaps less obvious, but nevertheless equally important. The point is that even if we know the initial conditions with perfect accuracy (which we usually do not), the finite precision of machine representation of numbers as well as round-off and truncation errors in computer algorithms would introduce small errors. So if arbitrarily small differences in initial conditions resulted in wildly different solutions, then prediction would be impossible.

Similarly we do not in practice ever know the vector field V perfectly. For example, in the problem of predicting the motions of the planets, it is not just their mutual positions that determine the force law V , but also the positions of all their moons and of the great multitude of asteroids and comets that inhabit the solar system. If the

tiny force on Jupiter caused by a small asteroid had a significant effect on its motion, then predicting the planetary orbits would be an impossible task.

In the preceding section we saw that complete, C^1 vector fields do give rise to a well-posed initial value problem, so Laplace seems to be on solid ground. Nevertheless, even though the initial value problems that arise in real-world applications may be technically well-posed in the above sense, they often behave as if they were ill-posed. For a class of examples that turns up frequently—the so-called chaotic systems—predictability is only an unachievable theoretical ideal. While their short-term behavior is predictable, on longer time-scales prediction becomes, for practical purposes, impossible. This may seem paradoxical at first; if we have an algorithm for predicting accurately for ten seconds, then should not repeating it with that first prediction as a new initial condition provide an accurate prediction for twenty seconds? Unfortunately, a hallmark feature of chaotic systems, called “sensitive dependence on initial conditions”, defeats this strategy.

Let us consider an initial value problem $\frac{dx}{dt} = V(x)$, $x(0) = p_0$ and see how things go wrong for a chaotic system when we try to compute $\sigma_{p_0}(t)$ for large t . Suppose that p_1 is very close to p_0 , say $\|p_0 - p_1\| < \delta$, and let us compare $\sigma_{p_1}(t)$ and $\sigma_{p_0}(t)$. Continuity with respect to initial conditions tells us that for δ small enough $\sigma_{p_1}(t)$ at least initially will not diverge too far from $\sigma_{p_0}(t)$. In fact, for a chaotic system, a typical behavior—when p_0 is near a so-called “strange attractor”—is for $\sigma_{p_1}(t)$ to at first “track” $\sigma_{p_0}(t)$ in the sense that $\|\sigma_{p_0}(t) - \sigma_{p_1}(t)\|$ initially stays nearly constant or even decreases—so in particular the motions of $\sigma_{p_0}(t)$ and $\sigma_{p_1}(t)$ are highly correlated. But then, suddenly, there will be a period during which $\sigma_{p_1}(t)$ starts to veer off in a different direction, following which $\|\sigma_{p_0}(t) - \sigma_{p_1}(t)\|$ will grow exponentially fast for a while. Soon they will be far apart, and although their distance remains bounded, from that time forward their motions become completely uncorrelated. If we make δ smaller, then we can guarantee that $\sigma_{p_1}(t)$ will track $\sigma_{p_0}(t)$ for a longer period, but (and this is the essence of sensitive dependence on initial conditions) **no matter how small we make δ , the veering away and loss of correlation will always occur**. The reason this is relevant

is that when we try to compute $\sigma_{p_0}(t)$, there will always be some tiny error in the initial condition, and in addition there will be systematic rounding, discretization, and truncation errors in our numerical integration process, so we are always in essence computing $\sigma_{p_1}(t)$ for some p_1 near p_0 rather than computing $\sigma_{p_0}(t)$ itself. The important thing to remember is that even the most miniscule of deviations will get enormously amplified after the loss of correlation occurs.

While there is no mathematical proof of the fact, it is generally believed that the fluid mechanics equations that govern the evolution of weather are chaotic. The betting is that accurate weather predictions more than two weeks in advance will never be feasible, no matter how much computing power we throw at the problem. As the meteorologist Edward Lorenz once put it, "... the flap of a butterfly's wings in Brazil can set off a tornado in Texas." This metaphor has caught on, and you will often hear sensitive dependence on initial conditions referred to as the "butterfly effect".

In Figure 1.2 we show a representation of the so-called "Lorenz attractor". This shows up in an ODE that Lorenz was studying as a highly over-simplified meteorological model. The Web Companion has a QuickTime Movie made with 3D-XplorMath that shows the Lorenz attractor being generated in real time. What is visible from the movie (and not in the static figure) is how two points of the orbit that are initially very close will moments later be far apart, on different "wings" of the attractor. (By the way, the fact that the Lorenz attractor resembles a butterfly is totally serendipitous!)




Figure 1.2. The Lorenz attractor.

Another strange attractor, shown in Figure 1.3, appears in an ODE called the Rikitake Two-Disk Dynamo. Like the Lorenz system, the Rikitake ODE was invented to model an important real-world phenomenon, namely the Earth’s geomagnetic field. The flipping back and forth between attractor “wings” in this case corresponds to the flipping of the Earth’s North and South Magnetic Poles that has long been known from the geologic record.




Figure 1.3. The Rikitake attractor.

Fortunately, even though systems that exhibit sensitive dependence on initial conditions do not permit long-time a priori prediction, it does not follow that such systems cannot be used to **control** processes that go on over long time periods. For example, when NASA sends a space-probe to a distant planet, the procedure is to look at all initial conditions and times that end up at the appropriate point on the given planet and then among these optimize for some variable (such as the transit time or payload weight). Of course they are using the prediction that with this choice of time and initial condition the probe will end up on the planet, but they realize that this prediction is only a first approximation. After lift-off, the current position and velocity of the probe is measured at intervals small enough to assure only small deviation from the previous predicted values. Then, these actual position and velocity are compared with the desired values and a “mid-course correction” is programmed that will bring the actual values back in line with the desired values. The equations governing

a space probe are not actually chaotic, but this same sort of controllability has also been proved rigorously for certain chaotic systems.

Experiment. Balance a broomstick vertically as best you can and let it fall. Repeat this many times, each time measuring the angle it makes with a fixed direction. You will see that the angles are randomly distributed around a circle, suggesting sensitive dependence on initial conditions (even though this system is not technically chaotic). Now place the broomstick on your fingertip and try to control it in a nearly upright position by making rapid slight finger motions—most people know almost instinctively how to do this. It is also instructive to note that you can make small rapid back-and-forth motions with your finger in a pre-planned direction, adding small perturbations as required to maintain the broomstick in approximate balance. (It is a fact that this actually serves to stabilize the control problem.)

We hope you have asked yourself an obvious question. If the weather is too chaotic to predict, can we perhaps nevertheless control it? After all, if a tiny butterfly can really perturb things enough to cause a storm a week later, it should not be beyond the power of humans to sense the effects of this perturbation while it is still small enough to counteract. (Of course this is not an entirely new idea—people have been seeding clouds to produce rain for decades. But the real challenge is to learn enough about how large weather systems evolve to be able to guide their development effectively with available amounts of energy.)

▷ **Exercise 1–21.** Learn how to control the weather. Hint: It could easily take you a lifetime to complete this exercise, but if you succeed, it will have been a life well spent.

1.5.1. Further Notes on Chaos. The study of chaotic systems is a relatively new field of mathematics, and even the “correct” definition of chaos is still a matter of some debate. In fact chaos should probably be thought of more as a “syndrome”—a related collection of symptoms—than as a precisely defined concept. We have concentrated here on one particular symptom of chaotic systems, their sensitive dependence on initial conditions, but there are others that

are closely related and equally as important, such as having a positive “Lyapounov Exponent”, the existence of so-called “strange attractors”, “homoclinic tangles”, and “horseshoe maps”. These latter concepts are quite technical, and we will not attempt to define or describe them here (but see the references below).

In recent years chaos theory and the related areas of dynamical systems and nonlinear science, have been the focus of enormous excitement and enthusiasm, giving rise to a large and still rapidly growing literature consisting of literally hundreds of books, some technical and specialized and others directed at the lay public. Two of the best nontechnical expositions are David Ruelle’s “Chance and Chaos” and James Gleick’s “Chaos: Making a New Science”. For an excellent introduction at a more mathematically sophisticated level see the collection of articles in “Chaos and Fractals: The Mathematics Behind the Computer Graphics”, edited by Robert Devaney and Linda Keen. Other technical treatment we can recommend are Steven Strogatz’ “Nonlinear Dynamics and Chaos”, Hubbard and West’s “Differential Equations: A Dynamical Systems Approach”, Robert Devaney’s “A First Course in Chaotic Dynamical Systems”, and Tom Mullin’s “The Nature of Chaos”.

1.6. Analytic ODE and Their Solutions

Until now we have worked entirely in the real domain, but we can equally well consider complex-valued differential equations. Of course we should be precise about how to interpret this concept, and in fact there are several different interpretations with different levels of interest and sophistication. Using the most superficial generalization, it seems as if there is nothing really new—since we can identify \mathbf{C} with \mathbf{R}^2 , a smooth vector field on \mathbf{C}^n is just a smooth vector field on \mathbf{R}^{2n} . But even here there are some advantages in using a complex approach. Recall the important two-dimensional real linear system $\frac{dx}{dt} = -y$, $\frac{dy}{dt} = x$, mentioned earlier, that arises when we reduce the harmonic oscillator equation $\frac{d^2x}{dt^2} = -x$ to a first-order system. We saw that if we regard \mathbf{R}^2 as \mathbf{C} and write $z = x + iy$ as usual, then our system becomes $\frac{dz}{dt} = iz$, so the solution with initial condition

z_0 is evidently $z(t) = e^{it}z_0$, and we recover the usual solution of the harmonic oscillator by taking the real part of this complex solution.

But if you have had a standard course on complex function theory, then you can probably guess what the really important generalization should be. First of all, we should replace the time t by a complex variable τ , demand that the vector field V that occurs on the right-hand side of our equation $\frac{dz}{d\tau} = V(z)$ is an analytic function of z , and look for analytic solutions $z(\tau)$.

To simplify the notation, we will consider the case of a single equation, but everything works equally well for a system of equations $\frac{dz_i}{d\tau} = V_i(z_1, \dots, z_n)$. We shall also assume that V is an entire function (i.e., defined and analytic on all of \mathbf{C}), but the generalization to the case that V is only defined in some simply connected region $\Omega \subset \mathbf{C}$ presents little extra difficulty.

Let us write $H(B_r, \mathbf{C})$ for the space of continuous, complex-valued functions defined on B_r (the closed disk of radius r in \mathbf{C}) that are analytic in the interior. Just as in the real case, we can define the map $F = F^{V, z_0}$ of $H(B_r, \mathbf{C})$ into itself by $F(\zeta)(\tau) = z_0 + \int_0^\tau V(\zeta(\sigma)) d\sigma$. Note that by Cauchy's Theorem the integral is well-defined, independent of the path joining 0 to τ , and since the indefinite integral of an analytic function is again analytic, F does indeed map $H(B_r, \mathbf{C})$ to itself. Clearly $F(\zeta)(0) = z_0$ and $\frac{d}{d\tau}F(\zeta)(\tau) = V(\zeta(\tau))$, so $\zeta \in H(B_r, \mathbf{C})$ satisfies the initial value problem $\frac{dz}{d\tau} = V(z)$, $z(0) = z_0$ if and only if it is a fixed point of F^{V, z_0} . The fact that a uniform limit of a sequence of analytic functions is again analytic implies that $H(B_r, \mathbf{C})$ is a complete metric space in the metric $\rho(\zeta_1, \zeta_2) = \|\zeta_1 - \zeta_2\|_\infty$ given by the "sup" norm, $\|\zeta\|_\infty = \sup_{\tau \in B_r} |\zeta(\tau)|$. We now have all the ingredients required to extend to this new setting the same Banach Contraction Principle argument used in Appendix B to prove the existence and uniqueness theorem in the real case. It follows that given $z \in \mathbf{C}$, there is a neighborhood O of z and a positive ϵ such that for each $z_0 \in O$ there is a unique $\zeta_{z_0} \in H(B_\epsilon, \mathbf{C})$ that solves the initial value problem $\frac{dz}{d\tau} = V(z)$, $z(0) = z_0$. And the proof in Appendix F that solutions vary smoothly with the initial condition generalizes to show that ζ_{z_0} is holomorphic in the initial condition z_0 .

Now let us consider the case of a real ODE, $\frac{dx}{dt} = V(x)$, but assume that the vector field $V : \mathbf{R}^n \rightarrow \mathbf{R}^n$ is analytic. This means simply that each component $V_i(x_1, \dots, x_n)$ is given by a convergent power series. Then these same power series extend the definition of V to an analytic map of \mathbf{C}^n to itself, and we are back to the situation above. (In fact, this is just the special case of what we considered above when the coefficients of the power series are all real.) Of course, if we consider only the solutions of this “complexified” ODE whose initial conditions z_0 are real and also restrict the time parameter τ to real values, then we get the solutions of the original real equation. So what we learn from this excursion to \mathbf{C}^n and back is that when the right-hand side of the ODE $\frac{dx}{dt} = V(x)$ is an analytic function of x , then the solutions are also analytic functions of the time and the initial conditions.

This complexification trick is already useful in the simple case that the vector field V is linear, i.e., when $V_i(x) = \sum_j^n A_{ij}x_j$ for some $n \times n$ real matrix A . The reason is that the characteristic polynomial of A , $P(\lambda) = \det(A - \lambda I)$, always factors into linear factors over \mathbf{C} , but not necessarily over \mathbf{R} . In particular, if P has distinct roots, then it is diagonalizable over \mathbf{C} and it is trivial to write down the solutions of the IVP in an eigenbasis. We will explore this in detail in Chapter 2 on linear ODEs.

1.7. Invariance Properties of Flows

In this section we suppose that V is some complete vector field on \mathbf{R}^n and that ϕ_t is the flow on \mathbf{R}^n that it generates. For many purposes it is important to know what things are “preserved” (i.e., left fixed or “invariant”) under a flow.

For example, the function $F : \mathbf{R}^n \rightarrow \mathbf{R}$ is said to be invariant under the flow (or to be a “constant of the motion”) if $F \circ \phi_t = F$ for all t . Note that this just means that each solution curve σ_p lies on the level surface $F = F(p)$ of the function F . (In particular, in case $n = 2$, where the level “surfaces” are level curves, the solution curves will in general be entire connected components of these curves.)

▷ **Exercise 1–22.** Show that a differentiable function F is a constant of the motion if and only if its directional derivative at any point x in the direction $V(x)$ is zero, i.e., $\sum_k \frac{\partial F(x)}{\partial x_k} V_k(x) = 0$.

The flow is called *isometric* (or distance preserving) if for all points p, q in \mathbf{R}^n and all times t , $\|\phi_t(p) - \phi_t(q)\| = \|p - q\|$, and it is called *volume preserving* if for all open sets O of \mathbf{R}^n , the volume of $\phi_t(O)$ equals the volume of O .

Given a linear map $B : \mathbf{R}^n \rightarrow \mathbf{R}^n$, we get a bilinear map $\hat{B} : \mathbf{R}^n \times \mathbf{R}^n \rightarrow \mathbf{R}$ by $\hat{B}(u, v) = \langle Bu, v \rangle$, where $\langle Bu, v \rangle$ is just the inner product (or dot product) of Bu and v . We say that the flow preserves the bilinear form \hat{B} if $\hat{B}((D\phi_t)_x(u), (D\phi_t)_x(v)) = \hat{B}(u, v)$ for all u, v in \mathbf{R}^n and all x in \mathbf{R}^n .

Here, the linear map $D(\phi_t)_x : \mathbf{R}^n \rightarrow \mathbf{R}^n$ is the differential of ϕ_t at x ; i.e., if the components of $\phi_t(x)$ are $\Phi_i(x, t)$, then the matrix of $D(\phi_t)_x$ is just the Jacobian matrix $\frac{\partial \Phi_i(x, t)}{\partial x_j}$. We will denote the determinant of this latter matrix (the Jacobian determinant) by $J(x, t)$. We note that because $\phi_0(x) = x$, $\frac{\partial \Phi_i(x, 0)}{\partial x_j}$ is the identity matrix, and it follows that $J(x, 0) = 1$.

▷ **Exercise 1–23.** Since, by definition, $t \mapsto \phi_t(x)$ is a solution of $\frac{dx}{dt} = V(x)$, $\frac{\partial \Phi_i(x, t)}{\partial t} = V_i(\phi_t(x))$. Using this, deduce that $\frac{\partial}{\partial t} \frac{\partial \Phi_i(x, t)}{\partial x_j} = \sum_k \frac{\partial V_i(\phi_t(x))}{\partial x_k} \frac{\partial \Phi_k(x, t)}{\partial x_j}$ and in particular that $\left(\frac{\partial}{\partial t}\right)_{t=0} \frac{\partial \Phi_i(x, t)}{\partial x_j} = \frac{\partial V_i(x)}{\partial x_j}$.

▷ **Exercise 1–24.** We define a scalar function $\mathbf{div}(V)$, the divergence of V , by $\mathbf{div}(V) := \sum_i \frac{\partial V_i}{\partial x_i}$. Using the formula for the derivative of a determinant, show that $\left(\frac{\partial}{\partial t}\right)_{t=0} J(x, t) = \mathbf{div}(V)(x)$.

▷ **Exercise 1–25.** Now, using the “change of variable formula” for an n -dimensional integral, you should be able to show that the flow generated by V is volume preserving if and only if $\mathbf{div}(V)$ is identically zero. Hint: You will need to use the group property, $\phi_{t+s} = \phi_t \circ \phi_s$.

▷ **Exercise 1–26.** Let B_{ij} denote the matrix of the linear map B . Show that a flow preserves \hat{B} if and only if $\sum_k \left(B_{ik} \frac{\partial V_k}{\partial x_j} + \frac{\partial V_k}{\partial x_i} B_{kj} \right) = 0$. Show that the flow is isometric if and only if it preserves \hat{I} (i.e.,

the inner product) and hence if and only if the matrix $\frac{\partial V_i}{\partial x_j}$ is everywhere skew-symmetric. Show that isometric flows are also measure preserving.

▷ **Exercise 1–27.** Show that the translation flows generated by constant vector fields are isometric and also that the flow generated by a linear vector field $V(x) = Ax$ is isometric if and only if A is skew-adjoint. Conversely show that if $V(x)$ is a vector field generating a one-parameter group of isometries of \mathbf{R}^n , then $V(x) = v + Ax$, where v is a point of \mathbf{R}^n and A is a skew-adjoint linear map of \mathbf{R}^n . Hint: Show that $\frac{\partial^2 V_i}{\partial x_j \partial x_k}$ vanishes identically.