Chapter 3

Existence and Uniqueness

> *An intellect which at a certain moment would know all forces that set nature in motion, and all positions of all items of which nature is composed, if this intellect were also vast enough to submit these data to analysis, it would embrace in a single formula the movements of the greatest bodies of the universe and those of the tiniest atom; for such an intellect nothing would be uncertain and the future just like the past would be present before its eyes*. (Pierre-Simon Laplace, *Essai philosophique sur les probabilities,* 1814)

The goal of this chapter is to prove the fundamental theorems of existence and uniqueness for solutions of ordinary differential equations (ODEs). As Laplace most eloquently stated, if one knows precisely the initial condition for the system of ODEs that describe the dynamics of a closed universe, it is possible—in principle—to construct the solution. The analysis in this chapter will also lead to a review of some fundamental mathematical machinery, such as the contraction-mapping theorem. We will find this theorem of use in many more exotic locales in later chapters.

The hypotheses of the existence theorem reveal some surprising requirements on the vector field for the solution of an ODE to exist and be unique. The theorem also makes clear that solutions of differential equations need not exist for all time, but only over limited intervals, even when the vector field is perfectly well behaved.

3.1 Set and Topological Preliminaries

Some of the basic notions from topology are essential in the study of dynamical systems, so we pause for a moment to collect some notation and recall a few of the ideas from set theory and topology that will be needed. Some common mathematical notation will be often used:

- ▷ $\mathbb{R}$ is the real line, and $\mathbb{R}^+ = \{x \in \mathbb{R} : x > 0\}$.[13]
- ▷ $\mathbb{R}^n$ is n-dimensional Euclidean space.

[13] The notation $\{a : b\}$ means the set of all a such that b holds. So, for example, $\{x \in \mathbb{R} : |x| < 1\}$ is the set of all real numbers between minus one and plus one.

▷ $\mathbb{Z}$ is the set of all integers.

▷ $\mathbb{N}$ is the set of natural numbers (the nonnegative integers including zero).

The Euclidean norm is denoted by $|x|$. A solid ball of radius r around a point x_o is the *closed* set

$$B_r(x_o) = \left\{x \in \mathbb{R}^n : |x - x_o| \le r\right\}. \tag{3.1}$$

We will be dealing primarily with differential equations on $\mathbb{R}^n$. The slightly more general case of "manifolds" is based on this analysis, since a manifold is a space that, locally, looks like Euclidean space.[14] Some common manifolds are

▷ $\mathbb{S}^d = \left\{(x_1, x_2, \ldots, x_{d+1}) : x_1^2 + x_2^2 + \cdots + x_{d+1}^2 = 1\right\}$ is the d-dimensional sphere; it is the boundary of a unit ball in $d + 1$ dimensions;

▷ $\mathbb{T}^d$ is the d-dimensional torus; and

▷ $\mathbb{S}^1 = \mathbb{T}^1$ is the circle.

Note that the "common sphere" embedded in three-dimensional space is denoted $\mathbb{S}^2$, the two-sphere, since it is a two-dimensional set. Additional notations include

▷ $\in$, an *element* of a set;

▷ $\subset$, a *subset*;

▷ $\cap$, intersection; and

▷ $\cup$, *union*.

For example, $3 \in \{5, 3, 2\}$, $\{0, 1, 2\} \cap \{2, 1\} = \{1, 2\}$, $\bigcup_{j=3}^{10} \{n \in \mathbb{N} : \ n < j\} = \{0, 1, 2, 3, 4, 5, 6, 7, 8, 9\}$, and $\bigcap_{j>0} \{n \in \mathbb{N} : \ n < j\} = \{0\}$. The qualifier symbols are denoted

▷ $\exists$, meaning *there exists*, and

▷ $\forall$, meaning *for all*.

A topological space is characterized by a collection of *open sets*. For Euclidean space the basic open sets are the open balls, $\{x : |x - x_o| < r\}$. By definition, a union of any number of open sets is declared open, as is the intersection of any finite number of open sets. Similarly, the basic closed sets are the closed balls $B_r(x_o)$. By definition, the intersection of any number of closed sets is closed, as well as the union of finitely many closed sets. The word *neighborhood* is used to denote some arbitrary set that encloses a designated point:

▷ *neighborhood*: N is a neighborhood of a point x if N contains an open set containing x.

Note that a neighborhood can be open or closed, but it must contain some open set. This excludes calling the set $\{x\}$ a neighborhood of x; however, for any $r > 0$, the closed ball $B_r(x)$ is a neighborhood of x. Often, we think of neighborhoods as being "small" sets in some sense, but this is not a requirement.

[14]Manifolds will be discussed more completely in Chapter 5.

Convergence

Sequences are ordered lists; for example, $S = \left\{s_j \in \mathbb{R}^n : j \in \mathbb{N}\right\}$. A sequence is *convergent* if it approaches a fixed value, s^*, i.e., if $\left|s_j - s^*\right| \to 0$ as $j \to \infty$. Formally, we say that the sequence S converges if for every $\varepsilon > 0$ there is an $N(\varepsilon)$ such that whenever $n > N(\varepsilon)$, then $|s_n - s^*| < \varepsilon$.

More generally a point x_o is called a *limit point* of the sequence x_j if there is a subsequence $\left\{s_{k_i} : \ k_i \in \mathbb{N}, \ \ k_j \to \infty \ \text{ as } \ j \to \infty\right\}$ that converges to x^*. For example, the sequence $\left\{(-1)^j : j \in \mathbb{N}\right\}$ has both 1 and -1 as limit points. With this notion we can formally define a

> ▷ *closed set*: A set S is closed if it includes all of its limit points; that is, if s^* is a limit point of some sequence in S, then $s^* \in S$.

The *closure* of a set S, denoted $\bar{S}$, is the union of the set and the limit points of every sequence in S.

The boundary of a set S is denoted ∂S. Consequently $\partial B_1(0) = \mathbb{S}^{n-1}$ is the unit sphere. A set is *bounded* if it is contained in some ball $B_r(0)$; otherwise, it is *unbounded*. A set that is both closed and bounded is called a

> ▷ *compact set*: A closed and bounded set in a finite-dimensional space is compact.

One of the basic theorems of topology states that every compact set, $C \subset \mathbb{R}^n$, can be covered by a finite number of balls: $C \subset \bigcup_{i=1}^{N} B_{r_i}(x_i)$.[15] Another important result is the next theorem.

Theorem 3.1 (Bolzano–Weierstrass). *Suppose every element of a sequence is contained in a compact set. Then the sequence has at least one limit point.*

Uniform Convergence

If a sequence depends upon a parameter—the elements of the sequence are functions, say, $f_n(x)$—then there is another notion of convergence that is important, that of

> ▷ *uniform convergence*: A sequence $\{f_n(x) : n \in \mathbb{N}, x \in E\}$ converges uniformly if for every $\varepsilon > 0$ there is an $N(\varepsilon)$ that can be chosen independently of x, such that whenever $n > N(\varepsilon)$, then $|f_n(x) - f^*(x)| < \varepsilon$ for all $x \in E$.

Uniformity of convergence will be especially important to help prove that limits of continuous functions are continuous. Recall that a continuous function $f \in C^0(E)$ is one for which for every $x \in E$ and every $\varepsilon > 0$, there is a $\delta(\varepsilon, x)$ such that $|f(y) - f(x)| < \varepsilon$ for all $y \in B_\delta(x)$. Here we allowed the distance δ to depend on both the accuracy ε and the choice of point x. An important consequence of uniform convergence is the next lemma.

Lemma 3.2. *The limit of a uniformly convergent sequence of continuous functions is continuous.*

[15] Indeed, this is usually taken as the more general definition of compact: a set for which *every* open cover has a finite subcover.

Proof. Let $u(x)$ denote the limit of $u_n(x)$; we must show that there is a $\delta(\varepsilon, x)$ such that $|u(y) - u(x)| < \varepsilon$ whenever $y \in B_\delta(x)$. Insert four new terms that sum to zero into this norm:

$$\begin{aligned}|u(y) - u(x)| &= |u(y) - u_n(y) + u_n(y) - u_n(x) + u_n(x) - u(x)| \\ &\le |u(y) - u_n(y)| + |u_n(y) - u_n(x)| + |u_n(x) - u(x)| .\end{aligned}$$

Since by assumption u_n converges uniformly, then for any $x \in E$ and any $\varepsilon/3$ there is a given N such that $|u_n(x) - u(x)| < \varepsilon/3$ whenever $n > N$. Moreover, since u_n is continuous for any fixed n, there is a $\delta(\varepsilon, x)$ such that $|u_n(x) - u_n(y)| < \varepsilon/3$ for each $y \in B_\delta(x)$. As a consequence,

$$|u(y) - u(x)| < \frac{\varepsilon}{3} + \frac{\varepsilon}{3} + \frac{\varepsilon}{3} = \varepsilon,$$

so u is continuous. □

There is also a uniform version of continuity:

▷ *uniform continuity*: A function f is uniformly continuous on E if for every $x \in E$ and every $\varepsilon > 0$, there is a $\delta(\varepsilon)$, independent of x, such that $|f(y) - f(x)| < \varepsilon$ for all $y \in B_\delta(x)$.

It is not too hard to show that when E is a compact set, then every continuous function on E is also uniformly continuous (see Exercise 2).

A generalization of Lemma 3.2 is easily obtained: if each of the elements of a convergent sequence is uniformly continuous, then the limit is also uniformly continuous.

3.2 Function Space Preliminaries

A function $f : D \to R$ is a map from its domain D to its range R; that is, given any point $x \in D$, there is a unique point $y \in R$, denoted $y = f(x)$. In our applications the domain is often a subset of Euclidean space, $E \subset \mathbb{R}^n$, and the range is $\mathbb{R}^n$; in this case, $f : E \to \mathbb{R}^n$ is given by n components $f_i(x_1, x_2, \ldots, x_n)$, $i = 1, 2, \ldots, n$. The set of functions denoted $C(E)$ or $C^0(E)$ consists of those functions on the domain E whose components are continuous. Colloquially we say "f is C^0" if it is a member of this set. If it is necessary to distinguish different ranges, the set of continuous functions from D to R is denoted $C^0(D, R)$; the second argument is often omitted if it is obvious. For a function $f : E \to \mathbb{R}^n$, the derivative at point x is written $Df(x) : \mathbb{R}^n \to \mathbb{R}^n$; it is defined to be a linear operator given by the *Jacobian matrix*

$$Df(x) \equiv \begin{pmatrix} \frac{\partial f_1}{\partial x_1} & \frac{\partial f_1}{\partial x_2} & \cdots & \frac{\partial f_1}{\partial x_n} \\ \frac{\partial f_2}{\partial x_1} & \frac{\partial f_2}{\partial x_2} & \cdots & \frac{\partial f_2}{\partial x_n} \\ \vdots & \vdots & \vdots & \vdots \\ \frac{\partial f_n}{\partial x_1} & \frac{\partial f_n}{\partial x_2} & \cdots & \frac{\partial f_n}{\partial x_n} \end{pmatrix}. \tag{3.2}$$

A function is $C^1(E)$—continuously differentiable—if the elements of $Df(x)$ are continuous on the open set E. Colloquially we will say that f is *smooth* when it is a C^1 function of its arguments.

Spaces of functions, like $C(E)$ and $C^1(E)$, are examples of infinite dimensional linear spaces, or *vector spaces*. Just as for ordinary vectors (recall §2.1), linearity means that whenever f and $g \in C(E)$, then so is $c_1 f + c_2 g$ for any (real) scalars c_1 and c_2. Much of our theoretical analysis will depend upon convergence properties of sequences of functions in some such space. To talk about convergence it is necessary to define a *norm* on the space; such norms will be denoted by $\|f\|$ to distinguish them from the finite dimensional Euclidean norm $|x|$. We already met one such norm, the operator norm, in (2.23). For continuous functions, the *supremum* or *sup-norm*, defined by

$$\|f\| \equiv \sup_{x \in E} |f(x)| \tag{3.3}$$

will often be used. For example, if $E = \mathbb{R}$, and $f = \tanh(x)$, then $\|f\| = 1$. Other norms include the L_p norms,

$$\|f\|_p = \left(\int_E |f(x)|^p \, dx \right)^{1/p} ,$$

but these will not have much application in this book. This formula becomes the sup-norm in the limit $p \to \infty$, which is why the sup-norm is also called the L_∞ norm and is often denoted $\|f\|_\infty$.

Metric Spaces

A normed space is an example of a metric space. A metric is a distance function $\rho(f, g)$ that takes as arguments two elements of the space and returns a real number, the "distance" between f and g. A metric must satisfy the three properties

1. $\rho(f, g) \geq 0$, and $\rho(f, g) = 0$ only when $f \equiv g$ (positivity),
2. $\rho(f, g) = \rho(g, f)$ (symmetry), and
3. $\rho(f, h) \leq \rho(f, g) + \rho(g, h)$ (triangle inequality).

Associated with any norm $\|f\|$ is a metric defined by $\rho(f, g) = \|f - g\|$. Therefore, a normed vector space is also a metric space; however, metric spaces need not be vector spaces, since in a metric space there is not necessarily a linear structure.

A sequence of functions f_n that are elements of a metric space X is said to *converge* to f^* if $\rho(f_n, f^*) \to 0$ as $n \to \infty$. Since the distance $\rho(f_n, f^*)$ is simply a number, the usual definition of limit can be used for this convergence. Note that the norm (3.3) bounds the Euclidean distance: if we use

$$\rho(f, g) = \|f - g\|_\infty , \text{ then } |f(x) - g(x)| \leq \rho(f, g).$$

Thus, convergence of a sequence of functions f_n in norm implies that the sequence of points $f_n(x)$ converges *uniformly*.

Another notion often used to discuss convergence is that of

▷ *Cauchy sequence*: Given a metric space X with metric ρ, a sequence $f_n \in X$ is Cauchy if, for every $\varepsilon > 0$ there is an $N(\varepsilon)$ such that whenever $m, n \geq N(\varepsilon)$, then $\rho(f_n, f_m) < \varepsilon$.

Informally, a Cauchy sequence satisfies

$$\rho(f_n, f_m) \to 0 \text{ as } m, n \to \infty,$$

where m and n approach infinity independently. One advantage of this idea is that the value of the limit of a sequence need not be known in order to check if it is Cauchy.

It is easy to see that every convergent sequence is a Cauchy sequence. However, it is not necessarily true that every Cauchy sequence converges.

Example: Consider the sequence of functions $f_n(x) = \sin(nx)\big/n \in C[0, \pi]$, the continuous functions on the interval $[0, \pi]$. This sequence converges to $f^* = 0$ in the sup norm because

$$\|f_n - 0\| = \frac{1}{n} \to 0.$$

The sequence is also Cauchy because

$$\|f_m - f_n\| \leq \frac{1}{n} + \frac{1}{m} \leq \frac{2}{N} < \frac{3}{N} \quad \forall m, n \geq N.$$

Thus for any ε, we may choose $N(\varepsilon) = 3\big/\varepsilon$ so that the difference is smaller than ε. ∎

Example: Consider the sequence $f_n = \sum_{j=1}^{n} \frac{x^j}{j}$ of functions in $C(-1, 1)$. Assuming that $m > n$, then

$$\|f_m - f_n\| = \left\| \sum_{j=n+1}^{m} \frac{x^j}{j} \right\| = \sum_{j=n+1}^{m} \frac{1}{j} \geq \int_n^m \frac{dy}{y} = \ln\left(\frac{m}{n}\right),$$

since the supremum of $\left|x^j\right|$ on $(-1, 1)$ is 1. This does not go to zero for m and n arbitrarily large but otherwise independent. For example, selecting $m = 2N$ and $n = N$ gives a difference larger than $\ln 2$. Consequently, the sequence is not Cauchy.

Note that for any fixed $x \in (-1, 1)$ this sequence converges to the function $-\ln(1-x)$; however, it does not converge uniformly since the number of terms needed to obtain an accuracy ε depends upon x. Thus in the sense of our function space norm, the sequence does not converge on $C(-1, 1)$. ∎

A space X that is nicely behaved with respect to Cauchy sequences is called a

▷ *complete space*: A normed space X is complete if every Cauchy sequence in X converges to an element of X.

For the case of linear spaces a complete space is called a

▹ *Banach space*: A complete normed linear space is a Banach space.

Some spaces, like a closed interval with the Euclidean norm, are complete, and some, like an open interval, are not. The space $C(E)$ with the L_∞ norm is complete.[16] However, the continuous functions are not complete in the L_2-norm.

Example: Let $f_n \in C[-1, 1]$ be the sequence

$$f_n = \begin{cases} 1, & x \le 0, \\ \dfrac{1}{1+nx}, & x > 0. \end{cases} \tag{3.4}$$

With the L_2-norm, this sequence limits to the function $f = \begin{cases} 1, & x \le 0 \\ 0, & x > 0 \end{cases}$ because

$$\| f_n - f \|_2 = \left(\int_0^1 \frac{dx}{(1+nx)^2} \right)^{1/2} = \frac{1}{\sqrt{1+n}} \underset{n\to\infty}{\to} 0.$$

Note that the limit, however, is not in $C[-1, 1]$. In the L_2-norm, the sequence is also a Cauchy sequence:

$$\begin{aligned} \| f_n - f_m \|_2^2 &= \int_0^1 \left(\frac{1}{1+nx} - \frac{1}{1+mx} \right)^2 dx \le \int_0^1 \left[\left(\frac{1}{1+nx} \right)^2 + \left(\frac{1}{1+mx} \right)^2 \right] dx \\ &= \frac{1}{1+n} + \frac{1}{1+m} \le \frac{2}{N}, \end{aligned}$$

for any $n, m \ge N$—of course every convergent sequence is Cauchy. As a consequence, the L_2-norm is not complete on the space $C[-1, 1]$. ∎

Example: Now consider the sequence (3.4) with the sup-norm. In this case the sequence does not converge to f, since

$$\| f_n - f \| = \max \left(|1 - 1|, \sup_{x \in (0,1]} \left| \frac{1}{1+nx} \right| \right) = \max\{0, 1\} = 1.$$

Accordingly, the very definition of convergence can depend upon the choice of norm. Moreover, this sequence is not Cauchy in the sup-norm:

$$\| f_n - f_m \| = \sup_{x \in [0,1]} \left| \frac{1}{1+nx} - \frac{1}{1+mx} \right| = \sup_{x \in [0,1]} \left| \frac{m-n}{(1+nx)(1+mx)} x \right|.$$

Differentiation of this expression shows that its maximum occurs at $x = (mn)^{-1/2}$ and has the value $\| f_n - f_m \| = \left| \frac{\sqrt{m}-\sqrt{n}}{\sqrt{m}+\sqrt{n}} \right|$ that does not approach zero for all $m, n \ge N \to \infty$. For example, $\| f_{4N} - f_N \| = \frac{1}{3}$. This proves that the sequence is not Cauchy. ∎

[16] The nontrivial proof is given in (Friedman, 1982) and (Guenther and Lee, 1996).

Since complete spaces are so important, it is worthwhile to note that given one such space we can construct more of them by taking subsets, as in the next lemma.

Lemma 3.3. *A closed subset of a complete metric space is complete.*

Proof. To see this, first note that if $f_j \in Y \in X$ is a Cauchy sequence on a complete space X, then $f_j \to f^* \in X$. Moreover, since f is a limit point of the sequence f_j, and a closed set Y includes all of its limit points, then $f \in Y$. □

The issues that we have discussed are rather subtle and worthy of a second look—see Exercise 1.

Contraction Maps

We have already used the concept of an operator, or map, $T : X \to X$, from a metric space to itself in Chapter 2: an $n \times n$ matrix is a map from $\mathbb{R}^n$ to itself. We will have many more occasions to use maps in our study of dynamical systems, including the proof of the existence and uniqueness theorem in §3.3. This proof will rely heavily on what is perhaps the most important theorem in all of analysis, the fixed-point theorem of Stefan Banach (1922).

Theorem 3.4 (Contraction Mapping). *Let $T : X \to X$ be a map on a complete metric space X. If T is a contraction, i.e., if for all $f, g \in X$, there exists a constant $c < 1$ such that*

$$\rho\,(T(f), T(g)) \le c\rho(f, g), \tag{3.5}$$

then T has a unique fixed point, $f^* = T(f^*) \in X$.

Proof. The result will be obtained iteratively. Choose an arbitrary $f_o \in X$. Define the sequence $f_{n+1} = T(f_n)$. We wish to show that f_n is a Cauchy sequence. Applying (3.5) repeatedly yields

$$\rho(f_{n+1}, f_n) = \rho\,(T(f_n), T(f_{n-1})) \le c\rho(f_n, f_{n-1}) \le c^2\rho(f_{n-1}, f_{n-2}) \le \cdots \le c^n\rho(f_1, f_o).$$

Therefore, for any integers $m > n$, the triangle inequality implies that

$$\rho(f_m, f_n) \le \sum_{i=n}^{m-1} \rho(f_{i+1}, f_i) \le \sum_{i=n}^{m-1} c^i\rho(f_1, f_0) = \frac{1 - c^{m-n}}{1 - c} c^n \rho(f_1, f_0) \le Kc^n,$$

where $K = \rho(f_1, f_o)\big/(1 - c)$. Since $c < 1$, then for any ε there is an N such that for all $m, n \ge N$, $\rho(f_m, f_n) \le Kc^N < \varepsilon$. This implies that the sequence f_n is Cauchy and, since X is complete, that the sequence converges.

The limit, f^*, is a fixed point of T. Indeed, suppose that N is large enough so that $\rho(f_n, f^*) < \varepsilon$ for all $n > N$, then

$$\begin{aligned}\rho(T(f^*), f^*) &\le \rho(T(f^*), f_{n+1}) + \rho(f_{n+1}, f^*)\\ &= \rho(T(f^*), T(f_n)) + \rho(f_{n+1}, f^*) < (c + 1)\varepsilon.\end{aligned}$$

Because this is true for any ε, the distance is zero and $T(f^*) = f^*$.

Finally, we show that the fixed point is unique. Suppose to the contrary that there are two fixed points $f \neq g$. Then, $\rho(f, g) = \rho(T(f), T(g)) \leq c\rho(f, g)$. Since $c < 1$, this is impossible unless $\rho(f, g) = 0$, but this contradicts the assumption $f \neq g$; thus, the fixed point is unique. □

Example: Consider the space $C^0(\mathbb{S})$ of continuous functions on the circle with circumference one, i.e., continuous functions that are periodic with period one: $f(x + 1) = f(x)$. For any $f \in C^0(\mathbb{S})$ define the operator

$$T(f)(x) = \frac{1}{2} f(2x).$$

Note that $T(f) \in C^0(\mathbb{S})$, and, using the sup-norm, that $\|T(f) - T(g)\| = 1/2\, \|f - g\|$; therefore, T is a contraction map on $C^0(\mathbb{S})$. What is its fixed point? According to the theorem, any initial function will converge to the fixed point under iteration. For example, let $f_o(x) = \sin(2\pi x)$. Then $f_1(x) = 1/2 \sin(4\pi x)$, and after n steps, $f_n = \frac{1}{2^n} \sin(2^{n+1}\pi x)$. A previous example showed that this sequence converges to $f^* = 0$ in the sup-norm. In conclusion, $f^* = 0$ is the unique fixed point. ■

Example: As a slightly more interesting example, consider the same function space but let

$$T(f)(x) = \cos(2\pi x) + \frac{1}{2} f(2x). \tag{3.6}$$

Note that T decreases the sup-norm by a factor of $1/2$ as before, so it is still contracting. For example, the sequence starting with the function $f_o(x) = \sin(2\pi x)$ is

$$\begin{aligned}
f_1(x) &= \cos(2\pi x) + \frac{1}{2} \sin(4\pi x), \\
f_2(x) &= \cos(2\pi x) + \frac{1}{2} \cos(4\pi x) + \frac{1}{4} \sin(8\pi x), \\
f_j(x) &= \sum_{n=0}^{j-1} \frac{\cos(2^{n+1}\pi x)}{2^n} + \frac{1}{2^j} \sin(2^{j+1}\pi x).
\end{aligned}$$

The last term goes to zero in the sup-norm, and by the contraction-mapping theorem, the result is guaranteed to be unique and continuous. The fixed point is not an elementary function but is easy to graph; see Figure 3.1; it is an example of a Weierstrass function (Falconer 1990). ■

Lipschitz Functions

Another ingredient that we will need in the existence and uniqueness theorem is a notion that is stronger than continuity but slightly less stringent than differentiability:

> ▷ *Lipschitz*: Let E be an open subset of $\mathbb{R}^n$. A function $f : E \to \mathbb{R}^n$ is Lipschitz if for all $x, y \in E$, there is a K such that
>
> $$|f(x) - f(y)| \leq K\, |x - y|\,. \tag{3.7}$$

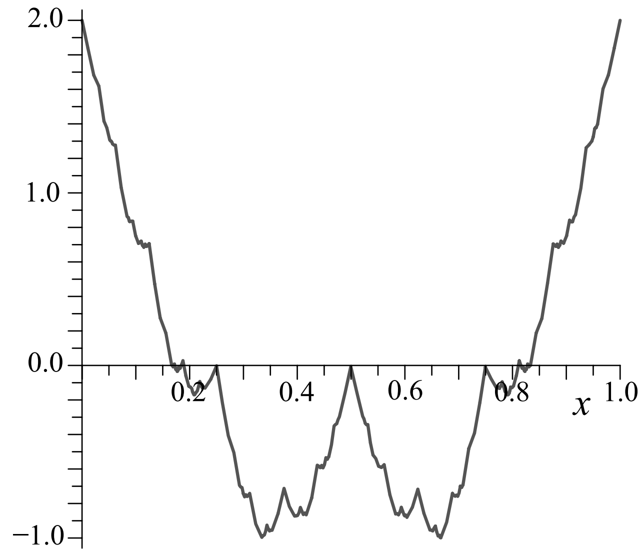

Figure 3.1. *The fixed point of the operator* (3.6).

The smallest such constant K is called the Lipschitz constant for f on E; it has the geometrical interpretation that the slope of the chord between the two points $(x, f(x))$ and $(y, f(y))$ is at most K in absolute value.

The Lipschitz property implies more than continuity, but less than differentiability.

Lemma 3.5. *A Lipschitz function is uniformly continuous.*

Proof. Choose any x. Then for every ε, there is a $\delta = \varepsilon \big/ K$ such that whenever $|x - y| \leq \delta$, $|f(x) - f(y)| \leq \varepsilon$. This is the definition of continuity. The continuity is uniform because δ is chosen independently of x. □

If the open set E is unbounded, then the assumption that f is Lipschitz is often too strong. For example, $f = x^2$ is not Lipschitz on $\mathbb{R}$, even though it is Lipschitz on every bounded interval (a, b). A weaker notion is

▷ *locally Lipschitz*: f is locally Lipschitz on an open set E if for every point $x \in E$, there is a neighborhood N such that f is Lipschitz on N. The Lipschitz constant can vary with the point and indeed become arbitrarily large.

Every differentiable function is locally Lipschitz.

Lemma 3.6. *Let f be a C^1 function on a compact, convex set A. Then f satisfies a Lipschitz condition on A with Lipschitz constant $K = \max_{x \in A} \|Df\|$.*

$$\begin{tikzcd}[column sep=small]
 & C^1(E) \arrow[dr, phantom, sloped, "\subset"] & & & \\
 & & \text{Locally Lipschitz}(E) \arrow[dr, phantom, sloped, "\subset"] & & \\
\text{Lipschitz}(E) \arrow[uurr, phantom, sloped, "\subset"] \arrow[dr, phantom, sloped, "\subset"] & & & & C^0(E) \\
 & & \text{Uniformly } C^0(E) \arrow[urr, phantom, sloped, "\subset"] & &
\end{tikzcd}$$

Figure 3.2. *Inclusion relations for Lipschitz functions.*

Proof. Since A is convex, the points on a line between two points $x, y \in A$, are also in A. Accordingly, when $0 \le s \le 1$, $\xi(s) = x + s(y - x) \in A$. Therefore

$$f(y) - f(x) = \int_0^1 \frac{d}{ds}\left(f(\xi(s))\right) ds = \int_0^1 Df\left(\xi(s)\right)(y - x) ds,$$

which amounts to the mean value theorem. Since A is compact and the norm of the Jacobian $\|Df\|$ is continuous, it has a maximum value K, as defined in the lemma. Thus

$$|f(y) - f(x)| \le \int_0^1 \|Df(\xi(s))\| \; |y - x| \, ds \le K\,|y - x|\,. \tag{3.8}$$

This is exactly the promised condition. □

Corollary 3.7. *If f is C^1 on an open set E, then it is locally Lipschitz.*

Proof. For any $x \in E$, there is an ε such that $B_\varepsilon(x) \subset E$. Since $B_\varepsilon(x)$ is compact and convex, the previous lemma applies. □

Finally, the lemma can be generalized to arbitrary compact sets.

Corollary 3.8. *Let $E \subset \mathbb{R}^n$ be an open set and $A \subset E$ be compact. Then if f is locally Lipschitz on E, it is Lipschitz on A.*

Proof. Every compact set can be covered by finitely many balls $B_j = B_{s_j}(x_j)$, $j = 1, \ldots, N$. The previous lemma implies that f satisfies a Lipschitz condition on each ball with constant K_j. Since there are finitely many elements in the cover, f satisfies a Lipschitz condition on A with Lipschitz constant $K = \max_{j \in [1,N]}(K_j)$. □

Some of the relationships between continuous, Lipschitz, and smooth functions are summarized in Figure 3.2.

Example: The function $f(x) = |x|$ is continuous and Lipschitz on $\mathbb{R}$. It is obviously C^1 on $\mathbb{R}^+$ and $\mathbb{R}^-$, and if x and y have the same sign, then $|f(x) - f(y)| = |x - y|$. So the only thing to be checked is the Lipschitz condition when the points have the opposite sign. Although this is obvious geometrically, let us be formal: let $x > 0 > y$; then $|f(x) - f(y)| = ||x| - |y|| \le x + |y| = |x - y|$. So f is Lipschitz with $K = 1$.

However, the function $f(x) = |x|^{1/2}$ is not Lipschitz on any interval containing 0. For example, choosing $x = 4\varepsilon$, $y = -\varepsilon$, we then have

$$|f(x) - f(y)| = \sqrt{|x|} - \sqrt{|y|} = \sqrt{\varepsilon} = \frac{\varepsilon}{\sqrt{\varepsilon}} = \frac{1}{\sqrt{\varepsilon}} \frac{4\varepsilon - (-\varepsilon)}{5} = \frac{1}{5\sqrt{\varepsilon}} |x - y|,$$

so that as ε becomes small, the needed value of $K \to \infty$. ∎

All these formal definitions have been given to provide us with the tools to prove that solutions to certain ODEs exist and, if the initial values are given, are unique. We are finally ready to begin this analysis.

3.3 Existence and Uniqueness Theorem

Before we can begin to study properties of the solutions of differential equations, we must discover if there *are* solutions in the first place: do solutions exist? The foundation of the theory of differential equations is the theorem proved by the French analyst Charles Emile Picard in 1890 and the Finnish topologist Ernst Leonard Lindelöf in 1894 that guarantees the existence of solutions for the initial value problem

$$\dot{x} = f(x), \quad x(t_o) = x_o \tag{3.9}$$

for a solution $x : \mathbb{R} \to \mathbb{R}^n$ and vector field $f : \mathbb{R}^n \to \mathbb{R}^n$. We were able to avoid this discussion in Chapter 2 because linear differential equations can be solved explicitly. Since this is not the case for more general ODEs, we must now be more careful.

The main tool that we will use in developing the theory is the reformation of the differential equation as an integral equation. Formally integrating the ODE in (3.9) with respect to t yields

$$x(t) = x_o + \int_{t_0}^{t} f(x(\tau))d\tau. \tag{3.10}$$

This equation, while correct, is actually not a "solution" for $x(t)$ since in order to find x, the integral on the right-hand side must be computed—but this requires knowing x.

Begin by imagining that (3.10) can be solved to find a function

$$x : J \to \mathbb{R}^n$$

on some time interval $J = [t_o - a, t_o + a]$. Since the integral in (3.10) does not require differentiability of $x(\tau)$, we will only assume that it is continuous. However, given that such a solution $x(t)$ to (3.10) exists, then it is actually a solution to the ODE (3.9).

Lemma 3.9. *Suppose $f \in C^k(E, \mathbb{R}^n)$ for $k \geq 0$ and $x \in C^0(J, E)$ is a solution of* (3.10). *Then $x \in C^{k+1}(J, E)$ and is a solution to* (3.9).

Proof. First note that if x solves (3.10), then $x(t_o) = x_o$. Since $x \in C^0(J)$, the integrand $f(x(\tau))$ is also continuous and so the right-hand side of (3.10), being the integral of a continuous function, is C^1; consequently, the left-hand side of (3.10), $x(t)$, is also differentiable. By the fundamental theorem of calculus, the derivative of the right-hand side is