THE STABLE MANIFOLD THEOREM AND APPLICATIONS

BENJAMIN Z. CARDINAL

Abstract. This paper presents a proof of the stable manifold theorem, which states that every hyperbolic fixed point has a stable manifold. It also presents an introduction to dynamical systems theory, including brief discussions of λ-contractions, the dynamics of linear maps, and hyperbolicity.

Contents

Date: August 19, 2019.

1. Introduction

The stable manifold theorem is a result from dynamical systems theory. In general, it states that every hyperbolic fixed point has a stable manifold. In this paper a proof of the following statement of the theorem will be presented:

Theorem 1.1. *The Stable Manifold Theorem. Let $A : \mathbb{R}^2 \to \mathbb{R}^2$ be a linear map such that $A = \begin{pmatrix} k_u & 0 \\ 0 & k_s \end{pmatrix}$ with $0 < k_s < 1 < k_u$. Then there exists $\varepsilon > 0$ such that for any C^r map $f : \mathbb{R}^2 \to \mathbb{R}^2$ with $\|Df - A\| \leq \varepsilon$ at some point p, there exists a C^r curve $\gamma : B(p, \varepsilon) \to \mathbb{R}^2$ such that the graph of γ is the set of all points that converge to p.*

In order to present this proof, a bit of knowledge is required. So, Section 2 of this paper covers elementary definitions and ideas from dynamical systems theory. Section 3 covers contractions, and Section 4 covers linear maps and hyperbolicity. Finally, Section 5 presents a proof of the stable manifold theorem.

2. Elementary Definitions and Ideas

Fundamentally, dynamical systems theory is interested in the "eventual or asymptotic behavior of an iterative process" [2]. In simpler words, dynamical systems theory is interested in what happens when we repeat a process many, many times (in the discrete case), or what happens when we let a process run forever (in the continuous case).

But what is a dynamical system? In general, dynamical systems are the objects that arise when a set rule is followed over and over again, like the orbits of planets (which follow the laws of classical mechanics). There are also formal, set theory based definitions, but they go beyond the scope of this paper. Instead, a better way to give some insight into dynamical systems would be to give examples.

Example 2.1. The most basic example of a dynamical system would be inputting a single number into a calculator and hitting one function key over and over again. For instance, you could press the "sin" key. Given the initial number x, this would give the sequence of numbers:

$$x, \sin(x), \sin(\sin(x)), \sin(\sin(\sin(x))), \ldots$$

This is an example of a *discrete dynamical system*, some characteristics of which we can describe. One such characteristic is obvious after computing the first few elements of this sequence: this sequence converges to 0.

Example 2.2. There are also *continuous dynamical systems*, of which a popular example is the most basic model for a population of organisms: an initial population, p_0, that grows proportionally to its size. This can be represented by the differential equation

$$\frac{dp}{dt} = kp.$$

In the above equation, p is the population, t is time, and k is an arbitrary constant. The solution to this equation is $p(t) = p_0 e^{kt}$. As $t \to \infty$, we see that, if $k < 0$, then the population will collapse, and if $0 < k$, then the population will explode.

In these two examples, it was easy to see the long term behavior of the systems in question. However, this is far from the normal case, and to begin to describe the behavior of more complex dynamical systems, some basic definitions are necessary.

First, there are a few definitions and notations that do not pertain exclusively to dynamical systems theory, but to clarify and avoid confusion, they are included.

Notation 2.3. Let $f : \mathbb{R}^n \to \mathbb{R}^n$. The first derivative of f at x is denoted by $f'(x)$, the second derivative by $f''(x)$, and higher order derivatives by $f^{(r)}(x)$. The first derivative may also be represented by Df if it is being treated as the matrix representation of the derivative of a multi-variable function.

Definition 2.4. Let $I \subset \mathbb{R}^n$ be a closed ball and $f : \mathbb{R}^n \to \mathbb{R}^n$. If, for all $x \in I$, $f^{(r)}$ exists and is continuous, then f is of class C^r. If f is of class C^1, then it is *smooth*.

Definition 2.5. A function is a *homeomorphism* if it is a bijection, is continuous, and has a continuous inverse.

Proposition 2.6. *Let $v \in \mathbb{R}^n$, and let L and P be two linear functions with matrix representations A and B, respectively. Then $P \circ L(v) = (B \cdot A)v$.*

Also, unless otherwise stated, we will be working in $\mathbb{R}^n$.

With that out of the way, the following definitions and ideas give us tools to begin to describe the behavior of dynamical systems.

Definition 2.7. A *mapping* or *map*, for short, is a function used to define a dynamical system, where the domain and the co-domain are the same space.

Definition 2.8. To *iterate* a map is to apply the map to a previous result of itself, given an initial input. We will denote the nth iteration of a map f with initial input x as $f^n(x)$. Then, $f^n(x) = f^{n-1}(f(x))$, and $f^0(x) = x$. If f is a homeomorphism, then the inverse of f is f^{-1}, and the nth iteration of f^{-1} is $f^{-n}(x)$.

Definition 2.9. Given a map f and point p, $O^+(p) = \{p, f(p), f^2(p), \ldots\}$ gives the *forward orbit* of p. If f is a homeomorphism, then the *backward orbit* is given by $O^- = \{p, f^{-1}(p), f^{-2}(p), \ldots\}$, and the *full orbit*, $O(p)$, is given by $\{f^n(p) : n \in \mathbb{Z}\}$.

The basic goal of dynamical systems theory is to understand the orbits of maps. Any patterns or repetitions are of particular interest. The following definitions help us to describe patterns, in particular asymptotic behavior.

Definition 2.10. A point x is *periodic* if there exists $n \in \mathbb{N}$ such that $f^n(x) = x$. Also, if a map is periodic, its *prime period* is the least positive n for which $f^n(x) = x$. If a point is periodic and $n = 1$, then it is a fixed point.

Definition 2.11. A point x is *eventually periodic* if x is not periodic, but there exists $m, i \in \mathbb{N}$ such that $f^{n+i}(x) = f^i(x)$ for all $i \geq m$, i.e. $f^i(x)$ is periodic for all $i \geq m$.

Definition 2.12. Let p be a periodic of period n. A point x is *forward asymptotic* to p if $\lim_{i\to\infty} f^{in}(x) = p$. If f has an inverse, then the point x is *backward asymptotic* to p if $\lim_{i\to-\infty} f^{in}(x) = p$. If p is not a periodic point, then asymptotic points can be defined by requiring $\|f^i(x) - f^i(p)\| \to 0$ as $i \to \infty$ (in the forward asymptotic case) or $i \to -\infty$ (in the backward asymptotic case).

Definition 2.13. The set of points forward asymptotic to p is called the *stable set* of p and is denoted by $W^s(p)$. Similarly, the set of points backward asymptotic to p is called the *unstable set* of p and is denoted by $W^u(p)$.

Even with just these basic definitions (and not even all of them) we can begin to show interesting things about dynamical systems, including the proposition below.

Proposition 2.14. *Let $I = [a, b] \subset \mathbb{R}$, and let $f : I \to I$ be continuous. Then f has at least one fixed point in I.*

Proof. Let $g(x) = f(x) - x$. g is continuous on I. There are three cases:

(1) $f(a) \not> a$. Then, by hypothesis, $f(a) \not< a$. So, $f(a) = a$, and a is a fixed point.
(2) $f(b) \not< b$. Then, by hypothesis, $f(b) \not> b$, and $f(b) = b$. Then b is a fixed point.
(3) $f(a) > a$ and $f(b) < b$. Then $g(a) > 0$ and $g(b) < 0$. The Intermediate Value Theorem then says that there exists $c \in I$ such that $f(c) - c = g(c) = 0$. So, $f(c) = c$, and c is a fixed point.

□

3. Contractions

Contractions are not only the simplest kind of asymptotic behavior of discrete-time dynamical systems, they are fundamental to the concept of hyperbolicity and therefore the stable manifold theorem. So, we begin with an introduction to contractions.

Definition 3.1. A map f of a subset X of a Euclidean space is said to be *Lipschitz continuous* with Lipschitz constant λ, or λ-Lipschitz if

$$\|f(x) - f(y)\| \leq \lambda \|x - y\|$$

for any $x, y \in X$, where $\| \cdot \|$ denotes the Euclidean norm.

Definition 3.2. If a map f of a subset X of a Euclidean space is said to be λ-Lipschitz, then f is said to be a *contraction* or a *λ-contraction* if $\lambda < 1$.

As implied by their name, contractions converge to points. More formally,

Theorem 3.3. *(Contraction Principle) Let $X \subset \mathbb{R}^n$ be closed and $f : X \to X$ be a λ-contraction. Then f has a unique fixed point x_0 and $\|f^n(x) - x_0\| \leq \lambda^n \|x - x_0\|$ for every $x \in X$.*

Proof. Let $x, y \in \mathbb{R}^n$ and $n \in \mathbb{N}$. Then, iterate $\|f(y) - f(x)\| \leq \lambda \|x - y\|$. The first step goes as follows:

$$\|f(f(y)) - f(f(x))\| \leq \lambda \|f(y) - f(x)\| \leq \lambda \cdot \lambda \|x - y\| = \lambda^2 \|x - y\|.$$

Repeating n times gives

$$\|f(y)^n - f(x)^n\| \leq \lambda^n \|x - y\|. \tag{3.4}$$

Then, for $m \geq n$, by substituting $f^m(x)$ for $f^n(y)$, the triangle equality gives us

$$\|f^m(x) - f^n(x)\| \leq \sum_{k=0}^{m-n-1} \|f^{n+k+1}(x) - f^{n+k}(x)\| \tag{3.5}$$

$$\leq \sum_{k=0}^{m-n-1} \lambda^{n+k}\|f(x) - x\|. \tag{3.6}$$

Then, because of the fact about partial sums of geometric series,

$$(1-\lambda)\sum_{k=l}^{n-1} \lambda^k = \lambda^l + \lambda^{1+l} + \cdots + \lambda^{n-1} - \lambda^{1+l} - \lambda^{2+l} - \cdots - \lambda^n = \lambda^l - \lambda^n,$$

we get

$$\|f^m(x) - f^n(x)\| \leq \frac{\lambda^n}{1-\lambda}\|f(x) - x\|.$$

Because $\lambda < 1$, $\lambda^n \to 0$ as $n \to \infty$, and $(f^n(x))_{n\in\mathbb{N}}$ is a Cauchy sequence. Then, we have $\lim_{n\to\infty} f^n(x)$ exists and is in X (X is closed). Moreover, because (3.4), this limit is the same for all $x \in X$, so we can denote it by x_0. Then, by the triangle inequality, we have

$$\|f(x_0) - x_0\| \leq \|x_0 - f^n(x)\| + \|f^n(x) - f^{n+1}(x)\| + \|f^{n+1}(x) - f(x_0)\| \tag{3.7}$$

$$\leq (1+\lambda)\|x_0 - f^n(x)\| + \lambda^n\|x - f(x)\|. \tag{3.8}$$

Then, as $n \to \infty$, $\lambda^n \to 0$ and $\|x_0 - f^n(x)\| \to 0$, so $\|f(x_0) - x_0\| = 0$. By substituting x_0 for y in (3.4), we get $\|f^n(x) - x_0\| \leq \lambda^n\|x - x_0\|$. □

However, testing whether or not a map is a λ-contraction can be cumbersome. Below, we show how the derivative can be used to verify the contraction property, but first, a definition:

Definition 3.9. A *convex* set in $\mathbb{R}^n$ is a set C such that for all $a, b \in C$ the line segment with endpoints a, b is entirely contained in C. It is said to be *strictly convex* if for any points a, b in the closure of C the segment from a to b is contained in C, except possibly for one or both endpoints.

Note also that the Euclidean norm $\|\cdot\|$ can be extended to include matrices by defining it as follows:

Definition 3.10. Let A be a matrix. Then $\|A\| = \sup\{\frac{\|Ax\|}{\|x\|} : x \neq 0\} = \sup\{\|Ax\| : \|x\| = 1\}$.

Theorem 3.11. *(The Derivative Test) Let $C \subset \mathbb{R}^n$ be an open, strictly convex set and $\overline{C}$ its closure. Let $f : \overline{C} \to \overline{C}$ be differentiable on C and continuous on $\overline{C}$ with $\|Df\| \leq \lambda < 1$ on C. Then f has a unique fixed point $x_0 \in \overline{C}$ and $\|f^n(x) - x_0\| \leq \lambda^n\|x - x_0\|$ for all $x \in \overline{C}$.*

Proof. Let $x, y \in \overline{C}$. Now, we parameterize the line connecting x and y with $c(t) = x + t(y - x)$ for $t \in [0, 1]$ and let $g(t) = f(c(t))$. By strict convexity, $c((0, 1))$ is contained in C. Then, by the chain rule,

$$\|\frac{d}{dt}g(t)\| = \|Df(c(t))\frac{d}{dt}c(t)\| = \|Df(c(t))(y - x)\| \leq \lambda\|y - x\|. \tag{3.12}$$

Then, by the Mean Value Theorem,

$$\|f(y) - f(x)\| \leq \|g(1) - g(0)\| \leq \|\frac{d}{dt}g(t)\|(1-0) \leq \lambda\|y - x\| \tag{3.13}$$

So, we have $\|f(y) - f(x)\| \leq \lambda\|y - x\|$ which means that f is a λ-contraction. So, by the contraction principle, f has a unique fixed point $x_0 \in \overline{C}$ and $\|f^n(x) - x_0\| \leq \lambda^n\|x - x_0\|$ for all $x \in \overline{C}$. □

Some maps do not necessarily contract on their whole domain, but instead on only a part of it. These are called local contractions. One example, that is particularly useful for our purposes, follows below.

Proposition 3.14. *Let f be a continuously differentiable map with a fixed point x_0 and $\|Df_{x_0}\| < 1$. Then there is a closed neighborhood U (meaning the closure of an open ball containing x_0) such that $f(U) \subset U$ and f is a contraction on U.*

Proof. Let $\eta > 0$. Because Df is continuous, there exists a closed ball $U = \overline{B(x_0, \eta)}$ around x_0 where $\|Df_x\| \leq \lambda < 1$ for all $x \in U$. Fix $x, y \in U$. By the derivative test, $\|f(x) - f(y)\| \leq \lambda\|x - y\|$, and f is a contraction on U. Furthermore, taking $y = x_0$ gives us $\|f(x) - f(x_o)\| \leq \lambda\|x - x_0\| \leq \lambda\eta < \eta$. So, $f(x) \in U$ for all $x \in U$, and $f(U) \subset U$. □

We can use proposition 3.14 to show, given its hypotheses, that not only is there a neighborhood around a fixed point of f on which f is a contraction, but also that any map sufficiently close to f is also a contraction on the same neighborhood.

Corollary 3.15. *Let f be a continuously differentiable map with a fixed point x_0 and $\|Df_{x_0}\| < 1$, and let U be a closed neighborhood of x_0 such that $f(U) \subset U$. Then for all $\varepsilon > 0$, there exists $\delta > 0$ such that if a map g on U satisfies*

$$\|g(x) - f(x)\| \leq \delta$$

and

$$\|Dg(x) - Df(x)\| \leq \delta,$$

then g is a contraction on U, $g(U) \subset U$, and g has its own unique fixed point, $y_0 \in B(x_0, \varepsilon)$.

Proof. Let $\eta > 0$ and $\varepsilon > 0$. The linear map Df_x is continuous and dependent on x. So, there exists a closed ball $U = \overline{B(x_0, \eta)}$ around x_0 where $\|Df_x\| \leq \lambda < 1$ for all $x \in U$. Assume $\eta, \varepsilon < 1$, and let $\delta = \varepsilon\eta(1 - \lambda)/2$. Then, on U,

$$\|Dg\| \leq \|Dg - Df\| + \|Df\| \leq \delta + \lambda \leq \lambda + (1 - \lambda)/2 = (1 + \lambda)/2 < 1.$$

So, by the Derivative Test, g is a contraction on U.

Now, let $\mu = (1 + \lambda)/2$ and $x \in U$. Then $\|x - x_0\| \leq \eta$ and

$$\|g(x) - x_0\| \leq \|g(x) - g(x_0)\| + \|g(x_0) - f(x_0)\| + \|f(x_0) - x_0\| \tag{3.16}$$

$$\leq \mu\|x - x_0\| + \delta + 0 \tag{3.17}$$

$$\leq \mu\eta + \delta \tag{3.18}$$

$$\leq \eta(1 + \lambda)/2 + \eta(1 - \lambda)/2 = \eta, \tag{3.19}$$

which means $g(x) \in U$ and therefore $g(U) \subset U$.

Finally, since g is a contraction on U, it has its own fixed point, denoted here by y_0. Then we have

$$\|x_0 - y_0\| \leq \sum_{n=0}^{\infty} \|g^n(x_0) - g^{n+1}(x_0)\| \tag{3.20}$$

$$\leq \|g(x_0) - x_0\| \sum_{n=0}^{\infty} \mu^n \tag{3.21}$$

$$\leq \frac{\delta}{1-\mu} = \frac{\varepsilon\eta(1-\lambda)}{1-\lambda} < \epsilon, \tag{3.22}$$

so $y_0 \in B(x_0, \varepsilon)$. □

While this corollary does not apply directly to the stable manifold theorem, it is still an interesting application of the contraction principle. It also helps us with the idea that, in the neighborhood of contractions and their fixed points, the behavior of dynamical systems is a lot "nicer," which is a central idea of the stable manifold theorem.

4. Linear Maps and Hyperbolicity

Hyperbolicity, as mentioned before, is essential to the stable manifold theorem (it is, after all, the only hypothesis).With regards to the stable manifold theorem, hyperbolicity describes fixed points. However, in this section, the goal is to get a better sense of what hyperbolicity means in the context of dynamical systems, so we will discuss hyperbolic maps, which are types of linear maps. To begin, we discuss the dynamics of linear maps.

In one dimension, it is fairly straightforward to see how linear maps behave. Given a linear map $x \longmapsto \lambda x$, if $|\lambda| < 1$, then the map contracts to 0. If $|\lambda| > 1$, all nonzero orbits tend to infinity. Finally, if $|\lambda| = 1$, the map is either the identity or $x \longmapsto -x$, for which all orbits have period 2.

For more than one dimension, let us first consider maps in the Euclidean plane. As implied by the definition of a hyperbolic linear map, the asymptotic behavior for maps in the plane depend on the eigenvalues of the matrix representation of the map. There are three possible cases for the eigenvalues: two real eigenvalues, one real eigenvalue, or two complex conjugate eigenvalues. For the scope of this paper, we are only concerned with the first case, where there exists two real eigenvalues.

Proposition 4.1. *A linear map of $\mathbb{R}^2$, $x \longmapsto Ax$, is eventually contracting if all eigenvalues have absolute value less than one.*

Proof. Suppose that, for the eigenvalues λ, $\mu \in \mathbb{R}$, the equations $A\nu = \lambda\nu$ and $A\nu = \mu\omega$ give the nonzero eigenvectors ν and ω. Now, we want to express A in terms of the invariant lines under the transformation, that is, we want to represent A with respect to the basis consisting of ν and ω. From the above equations, we perform a coordinate transform C to get this representation of A, the matrix B.

$$C \cdot A = B = \begin{pmatrix} \lambda & 0 \\ 0 & \mu \end{pmatrix}$$

Then, by considering the diagonal matrix, we see that, if the eigenvalues have absolute value less than one, then similar to the one dimension case, the map is a contraction.

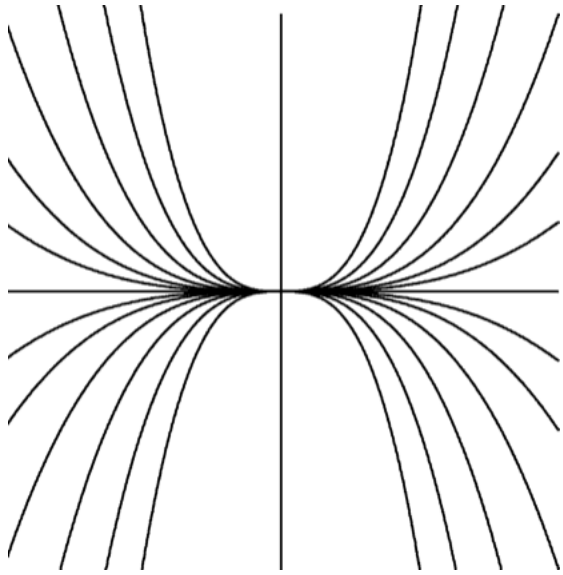

FIGURE 1. A node

With the coordinate change C (which is invertible), it is often the case that $\|Cv\| \neq \|\nu\|$. This would pose a problem, because we use the Euclidean norm in every matrix calculation, particularly with iteration. However, this problem can be solved by defining a new norm $\|\nu\|' := \|C\nu\|$ (this is a norm because C is linear and invertible). So, because, in this case, the calculations that we perform are independent of the choice of norm, no matter the choice of basis, we see that A gives the formula for a contraction. $\square$

Around the origin, we can begin to talk about the way that the points approach their fixed point (the origin) for these linear maps. For instance, let us consider the diagonal matrix A with distinct eigenvalues λ and μ. Then,

$$A^n \begin{pmatrix} x \\ y \end{pmatrix} = \begin{pmatrix} \lambda^n x \\ \mu^n y \end{pmatrix} \tag{4.2}$$

$$= \lambda^n \begin{pmatrix} x \\ (\frac{\mu}{\lambda})^n y \end{pmatrix}. \tag{4.3}$$

Then, let $|\mu| < |\lambda| < 1$. Then, orbits off the y-axis approach the origin at a rate $|\lambda^n|$, in particular, all orbits of points $\begin{pmatrix} x \\ y \end{pmatrix}$ where $x \neq 0$ move along curves preserved by A, which are eventually tangent to the x-axis at 0. In fact, using the rate above, integrating gives the equation $|y| = c|x|^\alpha$, where c is a constant and $\alpha = \log|\mu|/\log|\lambda|$, which gives the formula for these curves. As seen in Figure 1, there are four sets of curves symmetric about the axes, which together form a *node*. In the case that $x, y, \lambda, \mu > 0$, we can verify that the curve is invariant under A by using the fact $\log\mu = \alpha\log\lambda$. Then, we get $\mu = \lambda^\alpha$, which gives

$$c\lambda^{\alpha n}x^\alpha = c(\lambda^n x)^\alpha = \mu^n y = \mu(cx^\alpha).$$

This then implies that $y = cx^\alpha$, which means that y is invariant.

If not all eigenvalues have absolute value less than one, i.e. the map is noncontracting, then there are four other cases.

(1) Nonexpading Maps ($|\lambda| \not> 1, |\mu| \not> 1$). In this case, the map is noncontracting, but it is also nonexpanding. So, there are two possible subcases: the eigenvalues are 1 and -1, or one eigenvalue has absolute value 1 and the other does not have absolute value one. In the first subcase, the map represents a reflection, and the map has an orbit of period-2. In the second subcase, the eigenspace for the eigvenvalue with absolute value one consists of either fixed or period-2 points. Also, any point that is parallel to the eigenspace for the other eigenvalue will approach the eigenspace for the first eigenvalues. On all other points, the map is neither a contraction nor an expansion.

(2) One Expanding and One Neutral ($|\lambda| = 1, |\mu| > 1$). In this case, assume one eigenvalue λ has absolute value one. Then, the other, μ, has absolute value greater than one. Like with the second subcase of nonexpanding maps, the domain decomposes into two subspaces. However, instead of a space on which the map will be a contraction, we get a space on which the map will be an expansion. The other space remains neutral, neither a contraction nor an expansion.

(3) Both Expanding ($|\lambda| > 1, |\mu| > 1$). This case is the opposite of the contracting case, instead of converging along invariant curves, all orbits diverge. The map's inverse will be a contraction.

(4) The Hyperbolic Case ($|\lambda| > 1, |\mu| < 1$). The final case is the *hyperbolic* or *saddle* case.

Definition 4.4. A linear map of $\mathbb{R}^2$ is said to be *hyperbolic* if one eigenvalue is in $(-1, 1)$ and the other is not in $[-1, 1]$.

This definition can be generalized by instead requiring that no eigenvalue have absolute value one, at least one eigenvalue has absolute value less than one, and at least one eigenvalues has absolute value more than one.

As with the contracting case, we can diagonalize a linear map A with eigenvalues $|\lambda| > 1$ and $|\mu| < 1$ to get

$$B^n = \begin{pmatrix} \lambda^n & 0 \\ 0 & \mu^n \end{pmatrix}$$

after iteration. Because $|\lambda| > 1$, we see that all points along the (post-coordinate change) x-axis diverge. Moreover, all points that are off the y-axis have x-coordinates that diverge while their y-coordinate approaches 0. So, all orbits, approach the x-axis. As with the contracting case, all orbits move along the invariant curves $|y| = c|x|^\alpha$, where c is a constant and $\alpha = \log|\mu|/\log|\lambda|$, but in this case, α is negative. So, graphing the curves produces a *saddle*, which resembles the graphs of single-variable hyperbolic functions and gives the name of this case. A picture of this is shown in Figure 2 on the next page.

In this case, $\mathbb{R}^2$ decomposes into two subspaces post-iteration: one contracting (called the *stable subspace*) and one expanding (the *unstable subspace*). This is in contrast with the cases where we saw one subspace that either contracts of expands and one that does neither.

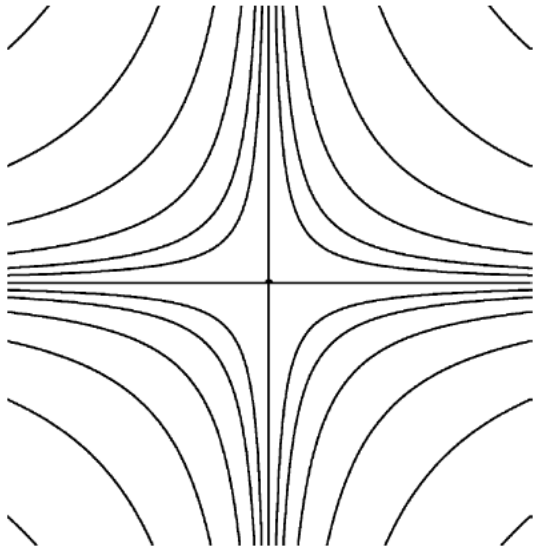

FIGURE 2. A saddle

The general idea about the stable manifold theorem is that specific types of smooth curves always have a small region under which they act like hyperbolic linear maps.

5. The Stable Manifold Theorem

Now, with all this in mind, we discuss the stable manifold theorem. After briefly discussing the ideas of the theorem, we will prove a lemma, which we will then use to prove the theorem as stated below:

Theorem 5.1. *The Stable Manifold Theorem. Let $A : \mathbb{R}^2 \to \mathbb{R}^2$ be a linear map such that $A = \begin{pmatrix} k_u & 0 \\ 0 & k_s \end{pmatrix}$ with $0 < ks < 1 < k_u$. Then there exists $\varepsilon > 0$ such that for any C^r map $f : \mathbb{R}^2 \to \mathbb{R}^2$ with $\|Df - A\| \leq \varepsilon$ at some point p, there exists a C^r curve $\gamma : B(p, \varepsilon) \to \mathbb{R}^2$ such that the graph of γ (the set containing all the ordered pairs $(x, \gamma(x))$) is the set of all points that converge to p.*

As you many have noticed, this version of the theorem only considers functions on $\mathbb{R}^2$. Even though the theorem holds for $\mathbb{R}^n$, we use this version for the sake of simplicity. The same technique applies in $\mathbb{R}^n$, but the arguments are geometrically more clear in $\mathbb{R}^2$.

Now, to discuss the theorem. Instead of a hyperbolic linear map, now consider a map $f : \mathbb{R}^2 \to \mathbb{R}^2$ with a fixed point x_0 such that $Df(x_0)$ is a hyperbolic linear map. So, x_0 is a *hyperbolic fixed point*. The idea is that now, instead of lines given by the eigenvalues, we have curves, or *manifolds* on which points diverge from or converge to x_0, called the *unstable* and *stable* manifolds, respectively. Then, because of Proposition 3.14, there is a small closed ball around x_0 where the orbits of f act as if f is a linear hyperbolic map on the closed region. The same idea, but with divergence, applies to the unstable subspaces.

Finally, before we prove the theorem, we have a definition and a lemma.

Definition 5.2. A *vertical γ-curve* is a graph $c \subset \mathbb{R}^2$ of a C^1 function with x as a function of y whose derivative is less than or equal to a constant γ.

Lemma 5.3. *Suppose $A : \mathbb{R}^2 \to \mathbb{R}^2$ is a linear map given by the matrix $A = \begin{pmatrix} k_u & 0 \\ 0 & k_s \end{pmatrix}$ with $0 < ks < 1 < k_u$. For $\gamma > 0$, $\varepsilon < (k_u - k_s/\gamma + 2 + (1/\gamma))$, and any C^1 map $f : \mathbb{R}^2 \to \mathbb{R}^2$ for which $\|Df - A\| \leq \varepsilon$, the preimage of a γ-curve under f is another γ-curve.*

Proof. For this proof, we define the norm $\|(x, y)\| := |x| + |y|$ on $\mathbb{R}^2$. Also, we will write $f(x, y)$ for $(f_1(x, y), f_2(x, y))$ and D_1 and D_2 for $\partial/\partial x$ and $\partial/\partial y$, respectively, for convenience.

Our goal is to show that for a γ-curve given by $x = c(y)$, we can solve the equation $f_1(x, y) = c(f_2(x, y))$, or, in this case, show that $0 = F(x, y) := f_1(x, y) - c(f_2(x, y)$ for another function $x = g(y)$ with $|Dg| \leq \gamma$, i.e. that g is also a γ-curve

Because the preimage of the γ-curve is non-empty, there exists a and b such that $F(a, b) = 0$. Then, we check that $D_1F \neq 0$. Note that $|D_1f_2| < \varepsilon$ and $|D_2f_1| < \varepsilon$ So, we have

$$|D_1f_1| \geq k_u - |D_1f_1 - k_u| \geq k_u - \varepsilon.$$

Therefore,

$$|D_1F| = |D_1f_1 - Dc \circ f_2D_1f_2| \geq k_u - (1 + \gamma)\varepsilon > 0,$$

and there exists a local solution for a function that gives $x = g(y)$. In order to estimate its derivative, note that

$$|D_2F| = |D_2f_1 - Dc \circ f_2D_2f_2| \leq \gamma(k_s + \varepsilon) - \varepsilon = \gamma k_s - (1 + \gamma)\varepsilon,$$

so, because $(1 + \gamma)(1 + 1/\gamma)\varepsilon < k_u - k_s$ we have that

$$|Dg| = |-\frac{D_2F}{D_1F}| \leq \frac{\gamma k_s - (1 + \gamma)\varepsilon}{k_u - (1 + \gamma)\varepsilon} < \gamma.$$

Now, we show that g is defined on all of $\mathbb{R}$, not just locally. First, note that $|D_2f_2| \leq k_s + \varepsilon$. So, fix $y \in \mathbb{R}$. Then, consider the graph of c over the interval $[a, a+(y-b)]$. The y-coordinates of the preimage extend from b to $b+(y-b)/(k_s+\varepsilon)$, which contains y. So, the preimage of our curve contains an arbitrary point (x, y), and g is globally defined such that $|Dg| \leq \gamma$. □

Note that, above, $\gamma+2+(1/\gamma) \geq 4$,, so $\varepsilon < (k_u < k_s)/4$, and any positive ε does *not* work.

Now, for the theorem.

Proof. For this proof, we will be using the l^∞ space of of bounded sequences, $\{x \in \mathbb{R}^\infty : \|x\|_\infty < +\infty\}$, where $\|x\|_\infty := \sup_{n \in \mathbb{N}_0} |x_n|$. Note that this is a metric space [1, Theorem A.1.14].

Now, let $\varepsilon > 0$ such that we can use Lemma 5.3, with the roles of x and y reversed, i.e., ε is small enough that f preserves horizontal γ-curves for some γ. Then, fix $y \in \mathbb{R}$ and consider a line $L_y := \{(x, y) : x \in \mathbb{R}\}$ and its successive images, $f^n(L_y)$, all of which are horizontal C^r γ-curves. For every $x \in \mathbb{R}$ and $n \in \mathbb{N}$, there exists a unique $z \in \mathbb{R}$ such that $(x, z) \in f^n(L_y)$. So, for all $x, x' \in l^\infty$, there exists unique $y, y' \in \mathbb{R}$ such that $(x', y') \in f^{n-1}(L_y)$ and $f((x', y')) = (x, y)$. Then, for each y, we define the map $\mathscr{F}_y : l^\infty \to l^\infty$ by

$$\mathscr{F}_y(x) = x'.$$

By construction, each sequence of lines above represents a bounded series of x-coordinates and their corresponding y-coordinates where each element x_n of $x \in \mathbb{R}^\infty$ is associated with its own y_n such that $(x_n, y_n) \in f^n(L_y)$.

$\mathscr{F}_y$ acts on each sequence of lines (which represent orbits of f) by dropping the first point of the sequence and re-indexing the sequence. After $n-1$ iterations, we see that a fixed point of $\mathscr{F}_y$ is the sequence of initial x-coordinates of the orbits of f paired with their corresponding y-coordinates. Similarly, every bounded semiorbit of f (one that converges after some $n \in \mathbb{N}$), leads to a fixed point of $\mathscr{F}_y$. As seen in Lemma 5.3, f expands the x-coordinates for each y because entrywise differences between the l^∞-sequences are divided by a factor of $k_u - \varepsilon > 1$, and therefore differences in the sup-norm are also divided by the same factor. So, $\mathscr{F}_y$ is a contraction, and by the Contraction Principle, there is a unique fixed point $(g(y), y) \in l^\infty$, where $g(y)$ is an implicitly defined continuous function dependent on y. So, the graph of $g(y)$ is *the* set of points with bounded positiv semiorbits, i.e., the stable manifold.

Finally, because l^∞ is a normed linear space, we can talk about differentiability. In particular, f and the curves $f^n(L_y)$ are C^r, so we see that $\mathscr{F}_y$, is also C^r. An extension of the Contraction Principle says that the implicit function that gives the fixed points of C^r contractions is itself C^r [1, Theorem 9.2.4], so g is also C^r, as required. □

Acknowledgments

I would like to thank my mentor, Hana Kong, for her continued assistance and patience. Even though this paper was not her area of expertise, she was happy to let me pursue my interests. Moreover, she was ever-willing to provide her help. I would also like to thank Peter May and Daniil Rudenko for organizing a fascinating REU and Apprentice program.

References

[1] Boris Hasselblatt and Anatole Katok. A First Course in Dynamics: with a Panorama of Recent Developments. Cambridge University Press. 2003. *This source was used to enahnce the details in many of the proofs in this paper.*

[2] Robert L. Devaney. An Introduction to Chaotic Dynamical Systems, Second Edition. Benjamin/Cummings. 2003. *This source was used to gain a greater understanding of the elementary ideas of dynamical systems.*