Introduction to Analysis I&II: Mathematics 401-402

Notes on Fleming's Text, plus other comments and ideas.

Kevin R. Vixie
Center for Geometric Analysis and Data
and
Department of Mathematics
Washington State University

January 16, 2020

Contents

1 Preface **1**

2 Basic notions in inner product spaces **3**

2.1 Section 1.1 . . . 3
2.2 Section 1.2 . . . 3
2.3 Section 1.3 . . . 4
2.4 Section 1.4 . . . 5
2.5 Section 1.5 . . . 5
2.5.1 Definition of Convex . . . 5
2.5.2 Supporting Hyperplanes, Half Spaces and Close Convex Sets . . . 5
2.5.3 Convex Combinations . . . 6

3 Topology in Metric spaces **9**

3.1 Functions . . . 9
3.2 Limits and Continuity of Transformations . . . 9
3.3 Sequences in $\mathbb{R}^n$ . . . 10
3.4 Bolzano-Weierstrass Theorem . . . 12
3.5 Relative Neighborhoods, Continuous Transformations . . . 14
3.6 Topological Spaces . . . 15
3.7 Connectedness . . . 15
3.8 Compactness . . . 15
3.9 Metric Spaces . . . 16
3.10 Spaces of Continuous functions . . . 17
3.11 Noneuclidean Norms on $\mathbb{R}^n$ . . . 18

4 Differentiation **21**

4.1 Directional and Partial Derivatives . . . 21
4.2 Linear Spaces and Functions . . . 21
4.3 Differentiable Functions . . . 22
4.4 Functions of Class C^q . . . 26
4.5 Local Extrema . . . 27
4.5.1 Assuming only that f is in C^2 . . . 28
4.5.2 A Slightly More Geometric Approach . . . 30
4.5.3 A $o(|h|^2)$ Approach . . . 30
4.5.4 When $Q(X_0,) = M(X_0)$ has non-zero determinant (and when it doesn't) . . . 31
4.6 Convex and Concave Functions . . . 32

5 Properties of Differentiable Functions **35**
5.1 Linear Transformations . . . 35
5.2 Affine Transformations . . . 35
5.3 Differentiable Transformations . . . 35
5.4 Compositions . . . 36
5.5 Inverse Function Theorem . . . 37
5.6 Implicit Function Theorem . . . 40
5.7 Manifolds . . . 43
5.7.1 Embedded manifolds are enough . . . 43
5.7.2 Definitions . . . 44
5.7.3 Intersections . . . 45
5.7.4 Comments . . . 46

6 Measure and Integration **47**
6.1 Riemann vs Lebesgue . . . 47
6.2 Iterated Integrals . . . 48
6.3 When Integrals Diverge. . . . 51
6.4 Transforming Volumes . . . 53
6.5 Regions Bounded by Simple Closed Curves . . . 54
6.6 The Area Formula . . . 57

Appendix **59**

Chapter 1

Preface

I begain writing these notes when the only text for the class was Fleming's book, but now I also use Lindstrom's book as needed. I will keep updating and adding to these notes as we go along.

These notes will not replace Fleming's book – it is always assumed you have read and (mostly) understood Fleming's book, using Lindstrom to illuminate when you need to. I use these notes as a way to share anything I want to share in writing.

Note 1: I use $\mathbb{R}^n$ where Fleming uses E^n.
Note 2: I use my own words without necessarily telling you they come from some other set of notes I have written – I honestly think that anti-plagiarism thing has been taken too far. It is appropriate and only right to very carefully acknowledge when you are quoting other people's writings/ideas, but it is nonsensical to get excited about how I use my own writing, because, frankly, it is my own writing and I can use it any way I like!

Chapter 2

Basic notions in inner product spaces

The first chapter in Fleming is a chapter that gets you thinking about things as an analyst and goes over some of the important tools and ideas we will use throughout the course.

2.1 Section 1.1

Sup (and Inf) It is an **axiom** in our system that every set in $\mathbb{R}$ that is bounded has a smallest or least upper bound (also known as the supremum): if $B < \infty$ and for every $x \in E$, $x \leq B$, there there is a least number s such that $x \leq s$ for all $x \in E$ and if there is another $t \in \mathbb{R}$ such that $x \leq t$ for all $x \in E$, then $s \leq t$.

Archimedean Property If $\epsilon > 0$ and $x > 0$, then there is a positive integer m such that $x < m\epsilon$. Existence of least upper bounds can be used to prove this.

Example if $x < y$, there is a rational number m/k, m and k integers, such that $x < m/k < y$. Proof for the case $0 < x < y$. (1) There is a $k \in \mathbb{Z}$ such that $2 < k(y - x)$. This implies that the gap between kx and ky is greater than 1 and there must be an integer m strictly between them: $0 < kx < m < ky$. (2) divide this result by k to get what we want, $0 < x < m/k < y$.

2.2 Section 1.2

The key ideas in this section are the idea of an inner product and the idea of the norm or 2-norm for vectors in $\mathbb{R}^n$. There are also some very important inequalities: Cauchy's inequality and the Triangle inequality.

Proof of Cauchy's Formula:

1. We define the norm of x to be $|x| = \sqrt{x_1^2 + \cdots + x_n^2}$. Note that $|x| = \sqrt{x \cdot x}$.
2. Note that
$$|x \cdot y| \leq |x|\ |y| \Leftrightarrow \left| \frac{x}{|x|} \cdot \frac{y}{|y|} \right| \leq 1$$
3. As a result, we can simply prove that for all x,y such that $|x| = |y| = 1$, we have that $|x \cdot y| \leq 1$.

4. Next we observe that for any real numbers a,b, $0 \leq (a \pm b)^2$ which implies that $|ab| \leq \frac{a^2}{2} + \frac{b^2}{2}$. Now we use this to prove that when $|x| = |y| = 1$, $|x \cdot y| \leq 1$:

5. Since $|x| = \sum_{i=1}^{n} x_i^2 = 1$ and $|x| = \sum_{i=1}^{n} x_i^2 = 1$.

$$\begin{aligned} |x \cdot y| &= |x_1 y_1 + \cdots + x_n y_n| \\ &\leq \sum_{i=1}^{n} |x_i y_i| \\ &\leq \sum_{i=1}^{n} \left(\frac{x_i^2}{2} + \frac{y_i^2}{2} \right) \\ &= 1 \end{aligned}$$

Problem 4 in section 1.2 Problem 4 in section 1.2 illustrates a common use of the Cauchy (and the more general Hölder's) inequality: I give a different hint here than Fleming gives.

Exercise 2.2.1. Show that $\sum_{i=11}^{n} |x_i| \leq \sqrt{n}|x|$.

1. What vector v satisfies $x \cdot v = \sum_{i=1}^{n} |x_i|$?
2. Apply Cauchy's inequality to x and v.

Orthogonal: i,e when $\langle \mathbf{x}, \mathbf{y} \rangle = 0$ The idea of orthogonality and orthonormal is very important: you should sketch examples in 2 and 3 dimensions.

Exercise 2.2.2. Suppose that $x, y \in \mathbb{R}^n$ are unit (column) vectors, so the transposes of each, x^T and y^T, are row vectors. Let I_n be the identity operator represented by the matrix with ones down the diagonal and zeros everywhere else. **Show that $(I_n - xx^T)y = y - \langle x, y \rangle x$ is orthogonal to x. More generally, Show that $I_n - xx^T$ is the matrix that projects any vector in $\mathbb{R}^n$ onto the n-1-dimensional subspace orthogonal to x.**

2.3 Section 1.3

Key ideas:

1. Given two points x and y, $\alpha x + (1 - \alpha)y$, $0 \leq \alpha \leq 1$ is the line segment from x to y.

2. Let $w \in \mathbb{R}^n$, and $c \in \mathbb{R}$. The set $\{x | x \in \mathbb{R}^n \text{ and } w \cdot x = c\}$ is called a hyperplanes and it divides $\mathbb{R}^n$ into two halfspaces. Note that choosing w such that $|w| = 1$ gets us all possible hyperplanes.

   Exercise 2.3.1. Prove it – that choosing w such that $|w| = 1$ gets us all possible hyperplanes. That is, if there is a hyperplane of the form $\{x | x \in \mathbb{R}^n \text{ and } w \cdot x = c\}$ where $|w| \neq 1$, then there is a u such that $|u| = 1$ and $\{x | x \in \mathbb{R}^n \text{ and } w \cdot x = c\} = \{x | x \in \mathbb{R}^n \text{ and } u \cdot x = b\}$ for some $b \in \mathbb{R}$ and $u \in \mathbb{R}^n$, $|u| = 1$.

3. E is convex if $x, y \in E \Rightarrow \ (\alpha x + (1 - \alpha)y) \in E \ \ \forall \ \ \alpha \in [0, 1]$

4. A closed convex sets E equals the intersection of all half-spaces containing E.

   Exercise 2.3.2. See if you can convince yourself of that fact: that closed convex sets equal the intersection of all half-spaces containing them. (We will return to this in section 1.5)

2.4 Section 1.4

Key ideas:

- Open balls $B(x, r)$, Interior points, Open sets $\mathcal{O}$
- $\bigcup_{\alpha \in \mathcal{A}} \mathcal{O}_\alpha$, $\bigcap_{i=1}^n \mathcal{O}_i$, $\bigcup_{i=1}^n \mathcal{C}_i$, $\bigcap_{\alpha \in \mathcal{A}} \mathcal{C}_\alpha$
- $\text{int}(E)$, $\text{clos}(E)$, limit points, cluter points
- $\emptyset \subseteq \text{int}(E) \subseteq E \subseteq \text{clos}(E) \subseteq X$
- Examples

Exercise 2.4.1. If E is a nonempty open set that does not equal the whole space X, then $\emptyset \subsetneq \text{int}(E) = E \subsetneq \text{clos}(E) \subsetneq X$. Let $0 = \subseteq$, and $1 =$ '='. We can represent the inclusion relation by $(0, 1, 0, 0)$ – we will call this 4-tuple the inclusion string. Using subsets of $X = \mathbb{R}$, see how many of the 16 possibilities for the inclusion string (x, x, x, x) you can find using different subsets.

2.5 Section 1.5

2.5.1 Definition of Convex

A set E is convex if the line segment joining any two points in E is also in E. We can express this in more than one way:

1. If $x, y \in E$, then the line segment between x and y is also in E
2. $x, y \in E \Rightarrow (\alpha x + (1 - \alpha)y) \quad \in E$ for all $0 \leq \alpha \leq 1$.
3. $x, y \in E \Rightarrow (\alpha(x - y) + y) \quad \in E$ for all $0 \leq \alpha \leq 1$.

The last way of expressing it makes it easy to see that the analytic expression involving α is gives us the line segment from y to x as α goes from 0 to 1, as long as you remember that $x - y$ is the vector from y to x.

2.5.2 Supporting Hyperplanes, Half Spaces and Close Convex Sets

If $h = \{x | x \cdot w = c\}$ is a hyperplane then there are two associated (closed) halfspaces: $H = \{x | x \cdot w \leq c\}$ and $H = \{x | x \cdot w \geq c\}$. A supporting hyperplane h_K of a closed convex set K is a hyperplane that intersects K with the additional property that K is entirely contained in either one or the other of the two halfspaces associated with h_k.

Intuitively, one can see that if we rotate everything appropriately, then K will sit on and above h_K but will not stick below h_K. In this case, h_K is supporting K.

If K is a closed convex set, then it equals the intersection of all the closed halfspaces containing K:

$$K = \bigcap_{H \in \mathcal{H}_K} H$$

where $\mathcal{H}_K$ are all the halfspaces containing K.

Proof.
Since every $K \subset H$ for every $H \in \mathcal{H}_K$, we have that $K \subset \bigcap_{H\in\mathcal{H}_K} H$. Therefore, we simply need to show $K \supset \bigcap_{H\in\mathcal{H}_K} H$ – i.e. that if x is not in K then $x \notin \bigcap_{H\in\mathcal{H}_K} H$.

Choose any point x not in K. Since K is a closed set, there is a closest point $k \in K$ and $d(x,k) = \inf_{y\in K} d(x,y) > 0$. (This follows from the fact $d(x,y)$ is a continuous function in y on K and a theorem we will meet in Chapter 2.) Now consider the hyperplane $h_k \equiv \{y \in \mathbb{R}^n | \langle y-k, x-k\rangle = 0\}$. Clearly, $k \in h_k$. Now, by a geometric argument, we prove that $K \subset \{y \in \mathbb{R}^n | \langle y-k, x-k\rangle \leq 0\}$. This will show that $x \notin \bigcap_{H\in\mathcal{H}_K} H$.

Suppose that there is a point $a \in K$ such that $\langle a-k, x-k\rangle > 0$. Because a and k are in K and K is convex, we have that the entire line segment between a and k are also in K. As a result, the point $\langle x-k, \frac{a-k}{|a-k|}\rangle \frac{a-k}{|a-k|} + k$ is closer to x than k is (Prove it: draw a careful picture!). That is a contradiction. Therefore there is no point $a \in K$ such that $\langle a-k, x-k\rangle > 0$. □

For every $k \in \partial K$, there is a $w \in \mathbb{S}^{n-1}$ such that $K \subset \{y | \langle y-k, w\rangle \leq 0\}$. That is, every point in the boundary of K, ∂K, is contained in a supporting hyperplane of K.

Proof.
Since k is in the boundary of K, there is a sequence $\{x_i\}_{i=1}^{\infty} \subset K^c$, such that $|x_i - k| \underset{i\to\infty}{\to} 0$. By the proof above, for each $x_i \in K^c$, there is a $k_i \in K$ such that $|x_i - k_i| \leq |x_i - k|$ is the distance from x_i to K and $K \subset \{y | \langle y-k_i, x_i-k_i\rangle \leq 0\}$. Defining $v_i \equiv \frac{x_i-k_i}{|x_i-k_i|}$, we have that $\{v_i\}_{i=1}^{\infty} \subset \mathbb{S}^{n-1}$ must have a cluster point $w \in \mathbb{S}^{n-1}$. Note that

$$\{y | \langle y-k_i, x_i-k_i\rangle \leq 0\} = \{y | \langle y-k_i, \frac{x_i-k_i}{|x_i-k_i|}\rangle \leq 0\}$$

Suppose that $p \in K$. We have that $\langle p-k_i, \frac{x_i-k_i}{|x_i-k_i|}\rangle \leq 0$ for all i, $k_i \to k$, and $\frac{x_i-k_i}{|x_i-k_i|} \to w$. Because $\langle a,b\rangle$ is a continuous function in a and b, this implies that $\langle p-k, w\rangle \leq 0\}$ and we conclude that $K \subset \{y | \langle y-k, w\rangle \leq 0\}$. □

2.5.3 Convex Combinations

Suppose that x is the convex combination of $\{x_i\}_{i=1}^m \in \mathbb{R}^n$ and $m \geq n+2$. In other words, suppose that $x = \sum_{i=1}^m \alpha_i x_i$ where $\{x_i\}_{i=1}^m \in \mathbb{R}^n$, $m \geq n+2$, $0 \leq \alpha_i \leq 1$ for all i and $\sum_i \alpha_i = 1$. Then there is a convex combination of at most $n+1$ of $\{x_i\}_{i=1}^m$ that also equals x.

Proof.
Define m to be the number of nonzero α_i's. We assume that at least $n+2$ of the α_i's are non-zero, otherwise we are done. Let M be the matrix whose m columns are the n-vectors $\{x_i\}_{i=1}^m$ Now add another row of ones to get a matrix $\hat{M}$ with $n+1$ rows and m columns. Since $m \geq n+2$, $\hat{M}$ has a non-zero null vector $(\beta_1, ..., \beta_m)$. By design $\sum_i \beta_i = 0$ and $0 = \sum_{i=1}^m \beta_i x_i$. In fact $\sum_i t\beta_i = 0$ and $0 = \sum_{i=1}^m t\beta_i x_i$ for all real t. This implies that $x = \sum_{i=1}^m (\alpha_i + t\beta_i)x_i$ and $\sum_{i=1}^m (\alpha_i + t\beta_i) = 1$ for all t. We can choose t small enough that at least one of the $(\alpha_i + t\beta_i)$ is zero and all of the $(\alpha_i + t\beta_i)$ are non-negative. □

Suppose that E is a set in $\mathbb{R}^n$, x is a convex combination of $n+1$ points from E, and E has at most n connected components. Then there are n points in E that also be combined in a convex combination to get x.

Proof.
We suppose that $x = \sum_{i=1}^{m} \alpha_i x_i$ where $0 \leq \alpha_i \leq 1$ for all i and $\sum_i \alpha_i = 1$. We assume that we cannot find n points in E such that x is the convex combination of those n points otherwise we are done. This implies that $0 < \alpha_i$ for all i.

Suppose first we are in $\mathbb{R}^2$ and we have 3 points. Then the ray through x from any of those points x_1, x_2, or x_3 must not contain any points in E otherwise x is the combination of a point in one of those rays and the opposite x_i. Since x is also not in E, these rays, together with x partition E into at least 3 non-empty open subsets. See figure 2.1

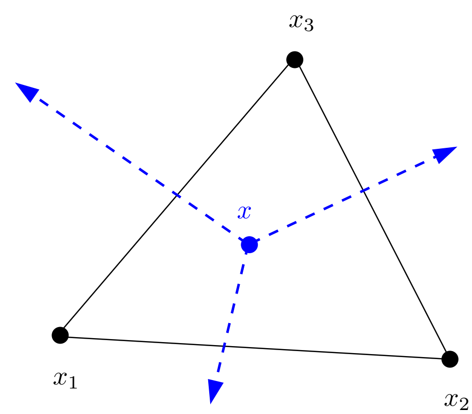

Figure 2.1: Proof in the case $n = 2$

Suppose that we are in $\mathbb{R}^n$ and $n > 2$. The argument we just used generalizes to any dimension, but visualizing takes a bit of effort.

1. First we note that the $n + 1$ points define a non-degenerate simplex S (i.e. a simplex with nonempty interior) in $\mathbb{R}^n$. Otherwise the n+1 points are contained in an n-1 dimensional plane and our previous argument gives us the result we are trying to prove.

2. This means that all subsimplices are non-degenerate as well.

3. The $n + 1$ points $\{x_i\}_{i=1}^{n+1}$ have the following property: if we chose one of these points x_k and we consider the complement $X_k \equiv \{x_i\}_{i \neq k}$, then the line through x_k and x hits the interior of the simplex S_k generated by X_k.

4. Now generate the $n + 1$ cones C_k with vertex x, consisting of all the infinite rays starting at x, which contain the antipodal points to points in S_k. The boundary of C_k are all the rays whose antipodal points are in a subsimplex of dimension less than or equal to $n - 1$ plus the vertex. If the vertex x is in E, we are done. If any point in E is in a ray that is antipodal to a subsimplex of dimension $n - 1$ or less, then x is the convex combination of at most n points in E.

5. We have that $x_k \in C_k$ for all k.

6. Therefore, if there is no way to express x as a convex combination of n or fewer points in E, then E has at least $n + 1$ connected components.

□

Chapter 3

Topology in Metric spaces

I am following the text fairly closely, though I am often giving different proofs and examples.

3.1 Functions

Learn to think about various sets in connection to functions and their properties:

- If $f : \mathbb{R}^n \to \mathbb{R}^k$, the set $\{(x, y)|y = f(x)\} \subset \mathbb{R}^{n+k} = \mathbb{R}^n \times \mathbb{R}^k$ is the graph of a function.
- inverse image of a set $F \subset \mathbb{R}^k$ in the range of $f : \mathbb{R}^n \to \mathbb{R}^k$, is the set $E = f^{-1}(F)$.
- Inverse images of any f preserves unions and intersections of subsets of the co-domain:

  * $f^{-1}(A \cup B) = f^{-1}(A) \cup f^{-1}(B)$
  * $f^{-1}(A \cap B) = f^{-1}(A) \cap f^{-1}(B)$

  even though, if f is not one to one, it is always the case that there are sets A and B in the domain of f such that:

  * $f(A \cup B) \neq f(A) \cup f(B)$
  * $f(A \cap B) \neq f(A) \cap f(B)$

3.2 Limits and Continuity of Transformations

There are three ideas here:

- if x_i is a sequence of values in $\mathbb{R}^n$, then if there is a point $x_0 \in \mathbb{R}^n$ such that the sequence eventually enters any small ball about x_0 and never leaves that ball, we say that x_0 is the limit of the sequence. More succinctly, if we choose any $\epsilon > 0$ (no matter how small!), there is an N_ϵ such that $x_i \in B(x_0, \epsilon)$ whenever $i > N_\epsilon$.

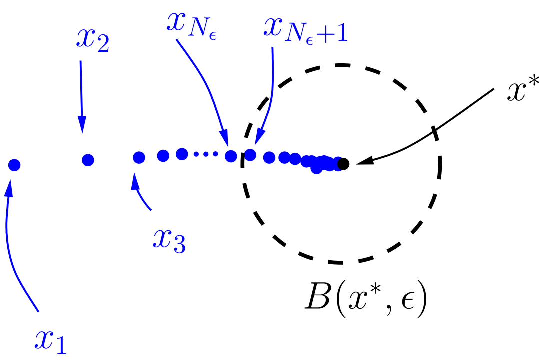

- if x is a continuous variable (like a point in $\mathbb{R}^n$), $f : \mathbb{R}^n \to \mathbb{R}^k$ and $x_0 \in \mathbb{R}^n$, then we say that $f(x) \to y_0 \in \mathbb{R}^k$ as $x \to x_0$ if, no matter how small we choose $\epsilon > 0$, we can chose $\eta > 0$ so that $x \in B(x_0, \eta)$ implies that $f(x) \in B(y_0, \epsilon)$

- If, no matter how small we choose $\epsilon > 0$, we can chose $\eta > 0$ so that $x \in B(x_0, \eta)$ implies that $f(x) \in B(f(x_0), \epsilon)$, we say that f **is continuous at** x_0. If f is continuous at every point in its domain, we simply say that f **is continuous** or f is continuous everywhere.

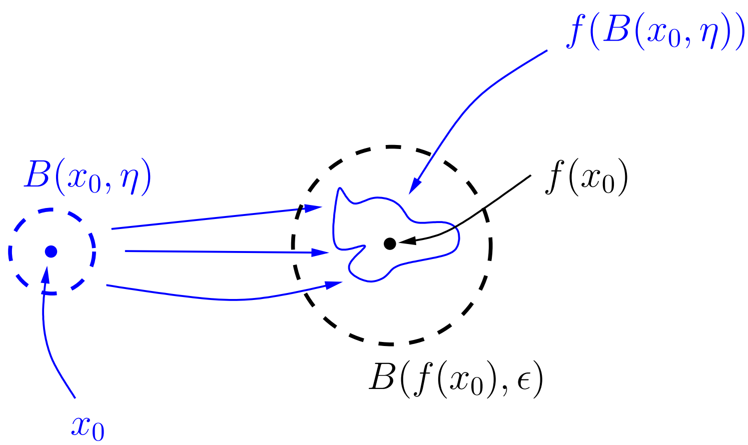

3.3 Sequences in $\mathbb{R}^n$

I recommend Burn's book [2] for a nice, problem driven exploration of sequences. (Earlier editions of he book are fine.)

- As mentioned in the previous section, a sequence $\{x_i\}_{i=1}^{\infty} \subset \mathbb{R}^k$ converges to $x^* \in \mathbb{R}^k$ if for any $\epsilon > 0$ there is an N_ϵ such that $i > N_\epsilon$ implies that $|x_i - x^*| < \epsilon$.

- A sequence is Cauchy if for any $\epsilon > 0$ there is an N_ϵ such that $m, n > N_\epsilon$ implies that $|x_n - x_m| < \epsilon$.

- In the text, it is shown that a sequence in $\mathbb{R}^k$ converges if and only if it is a Cauchy Sequence.

- The infinite sum of a sequence $\sum_{i=1}^{\infty} x_i$ (also called a series) is defined to be the limit f the partial sums $\lim_{m\to\infty} S_m$ where $S_m \equiv \sum_{i=1}^{m} x_i$.
- Define $S_m^n \equiv \sum_{i=m}^{n} x_i$. Then the m-tail of an infinite sum $\sum_{i=1}^{\infty} x_i$ is the sum $T_m = \sum_{i=m}^{\infty} x_i \equiv \lim_{n\to\infty} S_m^n$ where $S_m^n \equiv \sum_{i=m}^{n} x_i$.
- An infinite sum/series $\sum_{i=1}^{\infty} x_i = S$ if the sequence S_m converges to S. An infinite sum $\sum_{i=1}^{\infty} x_i$ converges to some point in $\mathbb{R}^k$ if $|T_m| \to 0$. Prove: $|T_m| \to 0$ if and only if S_m is a Cauchy Sequence.
- A series $\sum_{i=1}^{\infty} x_i$ is absolutely convergent if $\sum_{i=1}^{\infty} |x_i|$ is finite or equivalently $\sum_{i=1}^{\infty} |x_i|$ is convergent or equivalently $\sum_{i=1}^{\infty} |x_i| < infty$. Absolutely convergent implies convergent, but convergent does not imply absolutely convergent.
- $\sum_{i=0}^{\infty} x^i = \frac{1}{1-x}$ for any $x \in \mathbb{R}$ such that $-1 < x < 1$.
- Prove: $|x_i| < y_i$ and $\sum_{i=1}^{\infty} y_i < \infty$, $\Rightarrow \sum_{i=1}^{\infty} x_i$ is absolutely convergent.
- The previous bullet can be used to prove convergence by comparing series with known series and even integral.
- Show: since we know that $\int_1^{\infty} \frac{1}{x^p} dx < \infty$ when $p > 1$, this implies that $\sum_1^{\infty} \frac{1}{i^p} < \infty$.
- Show that $\sum_{i=1}^{\infty} a_i$ is convergent if $\lim_{i\to\infty} \left|\frac{a_{i+1}}{a_i}\right| < \rho < 1$. Hint: compare with $\sum_k \rho^k$. Note: I am abbreviating $\sum_{k=1}^{\infty} \rho^k$ by $\sum_k \rho^k$.

Exercise 3.3.1. Show that if

1. $X \equiv \{x_i\}_{i=1}^{\infty} \subset \mathbb{R}$
2. $|X| = \infty$: i.e. there are an infinite number of points in X (unlike, e.g., the infinite sequence $x_i = (-1)^i$ which has only two points, -1 and 1 in it.)
3. $\hat{x} \equiv \inf\{x \mid x \in X\}$,
4. and $\hat{x} \notin X$: in other words $\hat{x} \neq x_i$ for any i,

then there is a subsequence of X, x_{i_k} converging to $\hat{x}$: i.e

$$\lim_{k\to\infty} x_{i_k} = \hat{x}$$

Solution to Exercise 3.3.1:
We will do this in steps:

1. Draw examples to enable you to go between the pictures behind the proof and the written proof: this will take some effort, but you will be rewarded with understanding. See as an example, Figure 3.1 illustrating Step (5).
2. First: note that for every $\epsilon > 0$, there is some element of the sequence x_p such that $\hat{x} < x_p < \hat{x} + \epsilon$. Otherwise, there would be some $\epsilon > 0$ for which this is false, impling that $\hat{x} + \epsilon \leq x_i$ for all i. And that implies that $\hat{x}$ is NOT the greatst lower bound for X!

3. At this point we are almost there, we just needed a nicely ordered sequence of such points. That is what the rest of the proof does: it extracts the subsequence we need, using Step (2) several times.

4. Define $H_j = \{x_i\}_{i=1}^{j}$. These sets are just the first j elements for the sequence, $j = 1, 2, 3, ...$

5. Define $m_j = \{\text{minimum element of } H_j\}$. Notice that m_j is nonincreasing.

6. Define $i_1 = 1$. Since $\hat{x} \neq x_i$ for any i, this implies that $\epsilon \equiv x_{i_1} - \hat{x} > 0$. Thus, by Step (2), there will be a smallest $j < \infty$, where $m_j < i_1$. We define i_2 to be that j and note that $x_{i_2} = m_{i_2}$. (Why? Because it is at i_2 that the minumum jups down, implying that it was the i_2th element x_{i_2} that forced the minimum down.)

7. repeating this step, we get that i_{k+1} is the smallest $j > i_k$ such that $m_j < x_{i_k}$. Again $x_{i_{k+1}} = m_{i_{k+1}}$

8. Again, using Step (2), we know that m_j's converge to $\hat{x}$. (Otherwise $m_i > \hat{x} + \epsilon$ for all i and some $\epsilon > 0$. But Step (2) says that for some p, $x_p < \hat{x} + \epsilon$, which tells us that $m_p < \hat{x} + \epsilon$, which contradicts the first inequality,$m_i > \hat{x} + \epsilon$ for all i. Thus the m_j's converge to $\hat{x}$.)

9. So we know that the x_{i_k}'s, which are all distinct by design, converge monotonically to $\hat{x}$

□

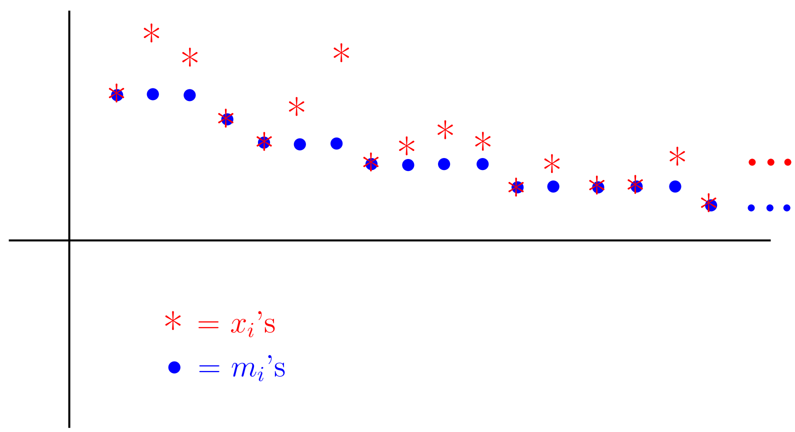

Figure 3.1: Example illustrating the m_j's.

3.4 Bolzano-Weierstrass Theorem

Proof 1: Here is my outline of a somewhat different path to the BW Theorem.

Part I Suppose that E is a bounded infinite set in $\mathbb{R}$. Let s_0 be the supremum of E. If s_0 is an accumulation point of E we are done. Otherwise s_0 is an isolated point in E, in which case we define $E_1 = E\{s_0\}$ and define s_1 to be the supremum of E_1. If this process stops in a finite number of steps, we have found an accumulation point. If it continues on an infinite number of times, each of the s_i are distinct, strictly decreasing elements in E, bounded below because E is a bounded set. Therefore the infimum of $\{s_i\}_{i=1}^{\infty}$ is an accumulation point of $\{s_i\}_{i=1}^{\infty}$ and therefore of E.

Part II This can then be extended to $\mathbb{R}^n$ by going coordinate by coordinate.

- Suppose that E is an infinite bounded set in $\mathbb{R}^n$. Let $\{x_i\}_{i=1}^{\infty}$ be any infinite sequence of distinct points in E. Because E is infinite, we know we can choose such an infinite sequence. Let $x_i = (x_i^1, x_i^2, x_i^3, ..., x_i^n)$. We now note that while each of the x_i are distinct, it might not be the case that each of the x_i are distinct. *As a result, for this step, we introduce the notion of* **accumulation point for sequences** *to: a point y^* is an accumulation point of a sequence $\{y_i\}$ if for every $\epsilon > 0$, $|y^* - y_i| < \epsilon$ for an infinite number of the i's.* We note that the argument for part one goes through exactly as before.
- Now use the fact that the first step of this proof tells us that $\{x_i^1\}_{i=1}^{\infty} \subset \mathbb{R}$ has an accumulation point and therefore a subsequence $\lim_{k\to\infty} x_{i_k}^1 \to \hat{x}^1$.
- Now define $x_k = x_{i_k}$. Since all the x_i's are distinct (in $\mathbb{R}^n$) so are the x_k's. We have that the first coordinate of the x_k's converges to $\hat{x}^1$.
- Now repeat, looking at the second coordinate $\{x_k^2\}_{k=1}^{\infty} \subset \mathbb{R}$. After n iterations of this procedure, we end up with $\hat{x} = (\hat{x}^1, \hat{x}^2, ..., \hat{x}^n)$, an accumulation point of E.

Proof 2: Here is something closer to the book,

Infinite Pigeon Hole Principle The set we start, E_0, is bounded and infinite – i.e. there are an infinite number of distinct points in the set. We want to conclude that there is a point x^* that is an accumulation point of E_0.

1. Any bounded set $E_0 \subset R^n$ can be enclosed in some closed cube of side length $R < \infty$, as long as R is big enough. Label this cube C_0.
2. If we cut the cube in half along each coordinate direction, we get 2^n closed subcubes, each having edge-length $\frac{R}{2}$.
3. Suppose that each of the subcubes had only a finite number of points from E_0 in them and say that the largest number contained in any of them is N. Then there would be at most $2^n N$ in E. That is a contradition, so it must be that at least one of the subcubes of edgelength $\frac{R}{2}$ has an infinite number of points in it. Label this subcube C_1
4. Define $E_1 = C_1 \cap E_0$.
5. Repeat: subdivide C_1 to get 2^n cubes of edge length $\frac{R}{2^2}$, at least one of which has a infinite number of points from E_1 in it. Label that cube C_2 and $E_2 = C_2 \cap E_1$.
6. Continuing this way we get a nested sequence of closed cubes with edge lengths

$$\{R, \frac{R}{2}, \frac{R}{2^2}, \frac{R}{2^3}, \frac{R}{R^4}, ...\}$$

.

7. We know that an infinite sequence of closed sets whose diameters are converging to zero has a non-empty intersection consisting of exactly one point. Lets call that point c^*:

$$c^* \equiv \bigcap_{i=1}^{\infty} C_i$$

.

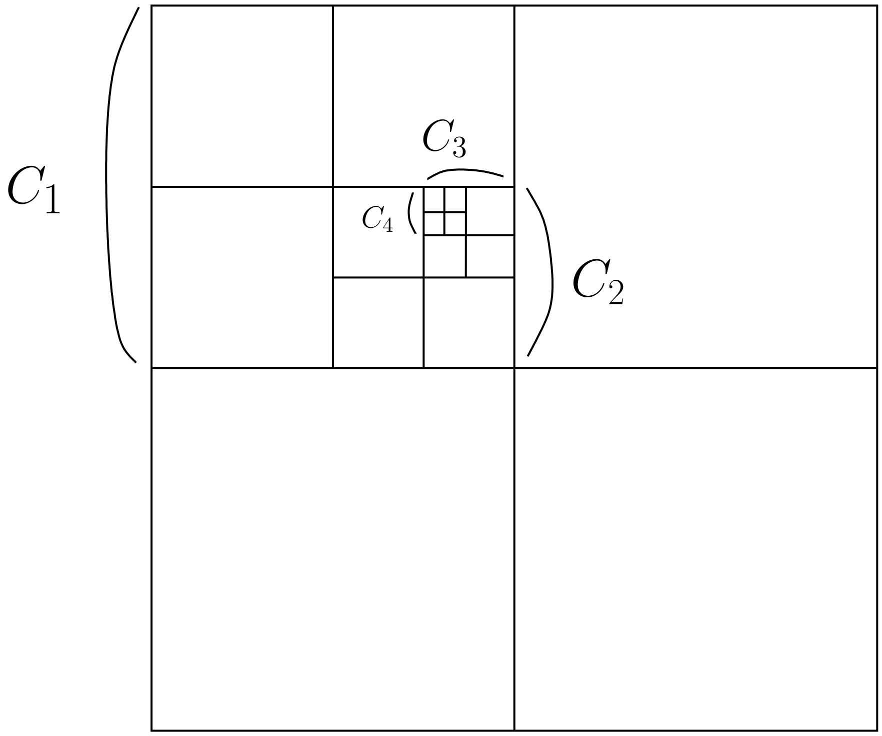

Cauchy Sequences Converge Strictly speaking, we have already shown the BW theorem in the book since c^* satisfies the property that every neighborhood of c^* contains an infinite number of points from E_0. Why? Because if you choose the neighborhood $B(c^*, \epsilon)$, then, as long as I choose C_k, where $\sqrt{n}\frac{R}{2^k} < \epsilon$, I know that $C_k \subset B(c^*, \epsilon)$ and by construction, $|C_k \cap E_0| = \infty$, where $|A|$ for a set A is the number of points in the set A. (This follows from the facts that c^* is in C_i for all i and $\text{diam}(C_i) = \sqrt{n}\frac{R}{2^i}$.)

But we can also construct a sequence $\{x_i\}_{i=1}^{\infty}$ of distinct points in E_0 that converges to c^*. We do that now.

1. pick any point in E_0 and label it x_0, now chose any point in $E_1 \setminus \{x_0\}$ and label it x_1.
2. Keep doing this to get, at the $(k+1)$th step, $x_k \in E_k \setminus \{x_0, x_1, ..., x_{k-1}\}$.
3. Notice that we can always do that since all the E_i's are infinite and removing at most a finite number of points does not create an empty set!
4. Notice that $x_i \in E_i \subset C_i$. This implies that $\{x_i\}_{i=0}^{\infty}$ is a Cauchy sequence with a limit x^*.
5. Since the tail of $\{x_i\}_{i=0}^{\infty}$ is in every C_i, and each of these are closed, $x^* \in C_i$ for all i. This implies that $x^* \in \bigcap_{i=0}^{\infty} C_i = c^*$.
6. Therefore $x_i \to c^*$.

3.5 Relative Neighborhoods, Continuous Transformations

3.6 Topological Spaces

3.7 Connectedness

3.8 Compactness

Here is a slightly different proof of the fact that if $E \subset \mathbb{R}^n$ is closed and bounded, then every open cover of E has a finite subcover, i.e. E is compact.

1. Suppose there is no finite subcover of the open cover $\mathcal{U}$.

2. First note that the collection of open balls with rational radius, centered on points in $\mathbb{R}^n$ which have rational coordinates, is a countable collection. (Prove it!) Name this collection $\mathcal{O} = \{O_i\}_{i=1}^{\infty}$

3. If $\mathcal{U}$ is an open cover of E, then for every point $x \in E$, there is an open set $U \in \mathcal{U}$ containing x.

4. There is some open ball with radius r_x, such that $B(x, r_x) \subset U$.

5. There is a point $p \in \mathbb{R}^n$ with rational coordinates such that $|p - x| < \frac{r_x}{4}$.

6. Choose a rational number q, $\frac{r_x}{4} < q < \frac{r_x}{2}$. Observe that
$$x \in B(p, q) \subset B(x, r_x) \subset U$$
and that $B(p, q) = O_i \in \mathcal{O}$ for some i. Restating, we have that for any $x \in E$ there is a $U \in \mathcal{U}$ and an $O_i \in \mathcal{O}$ such that:
$$x \in O_i \subset U \in \mathcal{U}.$$

7. Let $\hat{\mathcal{O}}$ be the subset of $\mathcal{O}$ containing all the O_i's needed in the previous step to cover all the x's in E.

8. For every O_i in $\hat{\mathcal{O}}$ which is a subset of some $U \in \mathcal{U}$, choose one element $U \in \mathcal{U}$ that contains O_i and label it U_i. Call this subcollection of $\mathcal{U}$, $\hat{\mathcal{U}}$, and note that $\hat{\mathcal{U}}$ is countable and infinite, so we can relabel $\hat{\mathcal{U}}$ to get $\hat{\mathcal{U}} = \{U_i\}_{i=1}^{N}$ with $N \leq \infty$.

9. By step 6 we have that $E \subset \bigcup_{U_i \in \hat{\mathcal{U}}} U_i$. This implies that $N = \infty$.

10. Because there is no finite subcover of $\mathcal{U}$, there is no finite subcover of $\hat{\mathcal{U}}$.

11. Define $C_m = E \setminus (\bigcup_{i=1}^{m} U_i)$.

12. Since E is closed, $E \subset B(0, R)$ for some $R < \infty$, and the U_i are open, we have that each of the C_m are closed, $C_1 \supset C_2 \supset C_3 \supset \cdots$ and all the $C_i \subset B(0, R)$.

13. Note that none of the C_i are empty (otherwise $\bigcup_{i=1}^{m} U_i$ would cover E and E would have a finite subcover), yet $\bigcap_{i=1}^{\infty} C_i = \emptyset$, contradicting corollary 2 on page 45 pf Fleming.

14. Thus there must be a finite subset of elements of $\mathcal{U}$ that also cover E.

3.9 Metric Spaces

A metric space is a set X with a distance function $d(\cdot,\cdot)$ that assigns to the pair of points $x, y \in X$ a distance $d(x, y)$ such that:

1. $d(x, y) = 0$ if and only if $x = y$.
2. $d(x, y) = d(y, x)$ for all $x, y \in X$
3. $d(x, z) = d(x, y) + d(y, z)$ for all $x, y, z \in X$

Note that $\mathbb{R}^n$ with the distance $d(x, y) = |x - y|$ is a metric space.

While topological spaces can be very wild, metric spaces are much better behaved. But there can be strangeness that metric spaces exhibit when they do not have the extra structure that $\mathbb{R}^n$ has. And $\mathbb{R}^n$ has a great deal more structure. For example:

1. $\mathbb{R}^n$ is a vector space, the **points in $\mathbb{R}^n$ are vectors** so we can add them and multiply them by scalars (real numbers).
2. This allows us to have **linear maps**, which are simple, and to talk about functions which are well approximated by those liner maps (differentiable functions).
3. in $\mathbb{R}^n$, the distance comes from a **norm** which has a special relation to the vector space structure.
4. in $\mathbb{R}^n$, the norm comes from an **inner product** which allows us to talk about angles and orthogonality.

As a result, some of the intuition you have for spaces is due to those really nice features, and when you want to know what is true in metric spaces, you have to shed some of those preconceptions.

For example, in a metric space with no extra structure, it makes no sense to ask about directions because there are no directions. On the other hand, a straight line in $\mathbb{R}^n$ is the shortest distance between any two points lying on that straight line. So we can attempt to define straight lines by thinking about shortest paths between two points in X.

Definition 3.9.1 (paths in metric spaces)**.** *A path in a metric space from $a \in X$ to $b \in X$ is a continuous map from $\gamma : [0, 1] \subset \mathbb{R} \to X$ such that $\gamma(0) = a$ and $\gamma(1) = b$.*

Definition 3.9.2 (lengths of paths in metric spaces)**.** *let γ be a path in a metric space. The length of γ, $l(\gamma)$, is $sup_P \sum_{i=0}^{N-1} d(\gamma(x_i, x_{i+1})$, where the P ranges over all partitions of $[0, 1]$: $x_0 = a < x_1 < x_2 < \cdots < x_{N-1} < x_N = b$ with $N < \infty$.*

Definition 3.9.3 (geodesics in metric spaces)**.** *Suppose that X is a path connected space. Then the geodesic distance from a to b is $d_g(a, b) \equiv inf_{\gamma\in\Gamma} l(\gamma)$ where Γ is the collection of all paths that start at a and go to b.*

There of course are questions like, "Is there a path from a to b whose length equals $d_g(a, b)$?". When the metric space X is complete (i.e. when every Cauchy sequence in X converges to a point in X) and the space is locally compact (every point in X has an open neighborhood whose closure is compact), it turns out that there is always a geodesic connecting any two points in the space.

Remark 3.9.1 (path metric spaces)**.** *To be a bit more efficient, we should not worry about connectedness or path connectedness and simply define the distance from a to b to be infinite if either (1) there is not path from a to b or (2) the length of all paths from a to b is infinite.*

There is a very nice book by Burago, Burago and Ivanov titled "A Course in Metric Geometry" which I recommend very highly. It is not fast reading, but it is very interesting and is aimed at those that want to learn this on their own. You can find a pdf here:
http://www.math.psu.edu/petrunin/papers/alexandrov/bbi.pdf. If you want to study this, you should buy the hard copy. I recommend using the pdf to decide. See also
http://www.pdmi.ras.ru/ svivanov/papers/bbi-errata.pdf.

Exercise 3.9.1. Define the g-length of a path between two points x, y in $\mathbb{R}^2$ to be the infimum of the lengths of paths between x and y where the length of a path $\gamma : [0,1] \to \mathbb{R}^2$ is defined to be $\int_0^1 g(\gamma(t))|\dot{\gamma}(t)|\,dt$. Describe the geodesics (shortest paths) when:

1) $g = \chi_{B(0,1)}$.
2) $g = \chi_{\mathbb{R}^1 \setminus B(0,1)}$.
3) $g = \chi_E$ where E is the union of three disjoint closed and bounded sets.

3.10 Spaces of Continuous functions

In this section, the main idea is the "uniform norm" and the fact that this norm in spaces of functions gives us completeness for both the space of bounded functions on S, $\mathcal{B}(S)$ and the space of continuous and bounded functions on S, $\mathcal{C}(S)$.

The ideas are simple:

1. The uniform norm on functions from $S \to \mathbb{R}^k$:
$$||f|| \equiv \sup_{x \in S} |f(x)|$$
is the main ideas here – it makes everything work for us.

2. If f_i is a Cauchy sequence in the uniform norm, then we know that for every $x \in S$, the sequence $\{f_i(x)\}_{i=1}^{\infty} \subset \mathbb{R}^k$ is a Cauchy Sequence converging to some $f^*(x) \in \mathbb{R}^k$.

3. When each of the functions $\{f_i\}_{i=1}^{\infty}$ is bounded, the function f^* is also bounded.

4. When each of the functions $\{f_i\}_{i=1}^{\infty}$ is bounded and continuous, then the limit f^* is also bounded and continuous.

5. Note that$\{f_i\}_{i=1}^{\infty}$ is a sequence of functions from S to $\mathbb{R}^k$, while for any $x \in S$, $\{f_i(x)\}_{i=1}^{\infty}$ is a sequence in $\mathbb{R}^k$.

Understanding why 1-5 are true one is of the missions in this section. Another goal is to get you used to thinking of vector spaces of functions and closed subspaces of function vector spaces.

Exercise 3.10.1. Extra Credit: Show that the set of all 3rd order polynomials is a vector space on the set $S = [0,1] \subset \mathbb{R}$ is a complete subspace of $\mathcal{C}(S)$, bounded continuous functions. In other words, show that a Cauchy sequence (in the uniform norm) of functions of the form

$$\{a_i + b_i x + c_i x^2 + d_i x^3\}_{i=1}^{\infty}$$

converge to some other function of the form $f^*(x) = a^* + b^*x + c^*x^2 + d^*x^3$.
Hint:

1. Show that
$$M \equiv \begin{bmatrix} 1 & x_1 & x_1^2 & x_1^3 \\ 1 & x_2 & x_2^2 & x_2^3 \\ 1 & x_3 & x_3^2 & x_3^3 \\ 1 & x_4 & x_4^2 & x_4^3 \end{bmatrix}$$
where the $x_i \in [0,1]$ are all distinct, is a nonsingular matrix – there are no, nontrivial null vectors.

2. This implies that the singular values of M are all non-zero. let σ_4 be the smallest singular value.

3. If $v \equiv (a,b,c,d)$ and $|v| > \delta$ then $|Mv| > \delta\sigma_4$.

4. Show that if $||f-g|| < \epsilon$ then $|v| = |(a,b,c,d)| < \frac{2\epsilon}{\sigma_4}$.

5. Use this to show that if $||f_i - f_j|| < \epsilon$ then,
$$|(a_i - a_j, b_i - b_j, c_i - c_j, d_i - d_j)| < \frac{2\epsilon}{\sigma_4}.$$

6. Show that if $|(a_i - a_j, b_i - b_j, c_i - c_j, d_i - d_j)| < \epsilon$ then $||f_i - f_j|| < 4\epsilon$.

7. Conclude that $\{f_i\}_{i=1}^{\infty}$ is Cauchy if and only if $\{(a_i, b_i, c_i, d_i)\}_{i=1}^{\infty}$ is Cauchy, that
$$(a^*, b^*, c^*, d^*) \equiv \lim_{i\to\infty}(a_i, b_i, c_i, d_i)$$
exists and $|(a_i - a^*, b_i - b^*, c_i - c^*, d_i - d^*)| \to 0$ implies that $||f_i - f^*|| \to 0$.

3.11 Noneuclidean Norms on $\mathbb{R}^n$

You know the Euclidean norm in $\mathbb{R}^n$: $|x| = \sqrt{x_1^2 + x_2^2 + \cdots + x_n^2}$. Any norm in $\mathbb{R}^n$ that is not this norm is called a non-Euclidean norm. The two most frequently used non-Euclidean norms are the L_1 and L_∞ norms,
$$|x|_1 \equiv \sum_{i=1}^{n} |x_i|$$
and
$$|x|_\infty \equiv \sup_{i\in\{1,2,\cdots,n\}} |x_i|.$$

Actually though there are an infinite number of non-Euclidean norms – one for every convex subset of $\mathbb{R}^n$ that is also symmetric with respect to the origin:

Definition 3.11.1. *E is* **symmetric with respect to the origin** *if $x \in E$ implies that $-x \in E$.*

Exercise 3.11.1. Extra Credit: See if you can find a way to create a one-to-one correspondence between origin symmetric, convex, compact subsets, of $\mathbb{R}^n$ having non-empty interior and norms in $\mathbb{R}^n$. Hint: Let $|x|_2$ denote the Euclidean norm. For any $v \in \mathbb{R}^n$ such that $|v|_2 \neq 0$, define the ray $R_v \equiv \{x \in \mathbb{R}^n | x = tv, 0 < t < \infty\}$. Notice that

1. $$\mathbb{R}^n = \bigcup_{v \in \partial B(0,1)} R_v \cup \{0\},$$

2. if we know a norm $|x|$ at one point of a ray R_v you know it at every point of R_v and

3. for any $\alpha \in [0,1]$,
$$|\alpha x + (1-\alpha)y| \leq \alpha|x| + (1-\alpha)|y|.$$

Chapter 4

Differentiation

This chapter is a somewhat gentle, but determined introduction to differentiation from the perspective of a geometric analyst.

4.1 Directional and Partial Derivatives

Pretty much what you have seen before.

4.2 Linear Spaces and Functions

This section is a reviews linear functions because derivatives are, in fact just linear functions which approximate the function we are differentiating. So you need to have an instinctive grasp of linear maps. And thus, there is an entire (though short) section reviewing them.

Remark 4.2.1 (Dual Spaces)**.** *You are already aware that the dual space of $\mathbb{R}^n$ looks a whole lot like $\mathbb{R}^n$. In fact, if you think of $\mathbb{R}^n$ as the space of column vectors of length n, then the dual space is the space of row vectors of length n (alternatively, it is the space of $1 \times n$ matrices). This is coincidental: not all dual spaces V^* look just like the original space V. We now give an example of a case in which the dual space looks very different.*

Example 4.2.1 (a dual space V^* can look different than V)**.** *We begin with a sequence of definitions that you should try to grasp intuitively. Think about them, draw some pictures, etc.*

Measures: *functions μ which map any subset of a space X to non-negative real numbers or ∞. They must also satisfy rules like $\mu(\cup_i A_i) \leq \sum_i \mu(A_i)$ and $\mu(\emptyset) = 0$. This is often called an* outer measure*. We will often refer to the value $\mu(A)$ as the* volume of A *or* measure of A*.*

Measurable sets: *sets E such that for all $A \subset X$, $\mu(A) = \mu(A \cap E) + \mu(A \cap E^c)$. These are the only sets we really pay attention to, in order to avoid things like the Banach-Tarski Paradox.*

Radon measures: *are measures (1) that assign finite volumes to compact sets, (2) such that we can approximate the measure or volume of any set E by the measures of open sets containing E and by the measure of closed sets contained in E, and (3) such that the class of measurable sets includes all open sets in $\mathbb{R}^n$.*

Signed Radon Measure: *a measure μ that can be expressed as $\mu(A) = \mu_p(A) - \mu_n(A)$, where μ_p and μ_n are Radon measures.*

Compactly Supported Functions: *f is compactly supported if $\{x \mid f(x) \neq 0\} \subset B(0, R)$ for some $R < \infty$. More precisely, the closure of $\{x \mid f(x) \neq 0\}$, called the support of f, is a compact set. In this case, we say that f has* compact support.

The space $C_c(\mathbb{R}^n; \mathbb{R})$: *Let $C_c(\mathbb{R}^n; \mathbb{R}) \equiv \{f : R^n \to \mathbb{R} \mid f$ is continuous, with compact support$\}$*

The result: *it turns out that the set of continuous linear functionals on $C_c(\mathbb{R}^n; \mathbb{R})$ – the dual space of $C_c(\mathbb{R}^n; \mathbb{R})$ – is the set of signed Radon Measures on R^n. Every continuous linear functional $L \in (C_c(\mathbb{R}^n; \mathbb{R}))^*$ is given by $L(f) \equiv \int f d\mu$ for some signed radon measure μ.*

The moral of the story: *The measures in $(C_c(\mathbb{R}^n; \mathbb{R}))^*$ do not look at all like the functions in $C_c(\mathbb{R}^n; \mathbb{R})$! Thus, even though in vector spaces like $\mathbb{R}^n$ the dual space is essentially the same as $\mathbb{R}^n$, this is not always the case!*

Note that, in our case, the precise definition of continuous linear functional is that for every compact set K, then there is a C_K such that $L(f) \leq C_K |f|_{\sup}$ for all f such that support of f is in K, where $|f|_{\sup}$ is the sup*-norm of f, i.e. $|f|_{\sup} \equiv \max_{x \in \mathbb{R}^n} |f(x)|$.*

Remark 4.2.2 (Inner Product Spaces are Nicer)**.** *If you are in a Hilbert space - a complete, normed vector space where your norm comes from an inner product*

$$|x| = \sqrt{\langle x, x \rangle},$$

then it is the case that your space and its dual are essentially the same: there is an isometric isomophism that connects them. That is as normed spaces, the space and its dual are indistinguishable. This is the message of the Riesz Representation Theorem that is first encountered in more advanced analysis classes: for any $w \in H^*$, where H is a Hilbert space, there is a vector $v_w \in H$ such that $w(x) = \langle v_w, x \rangle$ for all $x \in H$.

4.3 Differentiable Functions

Here we see that a function $f : \mathbb{R}^n \to \mathbb{R}^m$ is differentiable at $x \in \mathbb{R}^n$ if it is well approximated at x by a linear function:

$$f(x + h) = f(x) + L_x(h) + g(h)$$

and

$$\lim_{|h| \to 0} \frac{|g(h)|}{|h|} = 0.$$

In other words $(\Delta_x f)(h) \equiv f(x + h) - f(x)$ is well approximated by $L_x(h)$, where L_x is a linear function from $\mathbb{R}^n$ to $\mathbb{R}^m$:

$$f(x + h) - f(x) = L_x(h) + g(h)$$

or equivalently

$$\lim_{|h| \to 0} \frac{(f(x + h) - f(x)) - L_x(h)}{|h|} \to 0.$$

There is a geometric version of this: f is differentiable if there is a linear function L_x such that the graph of $(\Delta_x f)(h)$ and the graph of $L_x(h)$ can be contained in cones that can be made as narrow as you like, as long as you are willing to zoom into the point $(x, f(x)) \in \mathbb{R}^{n+m}$ enough. We can see this more precisely without much trouble.

1. For $h \in \mathbb{R}^n$, $(\hat{x}, f(\hat{x})) + (h, L(h))$ is a point on the shifted linear subspace of $\mathbb{R}^{n+m}$ that approximates the graph of f, $\{(x+h, f(x+h)) \mid x \in \mathbb{R}^n\}$.

2. The error vector $(0, g(h))$, has norm $|g(h)| \leq \alpha(|h|)|h|$ where
$$\alpha(s) \equiv \sup_{\{v||v|\leq s\}} \frac{|g(v)|}{|v|}.$$
We note that $\alpha(s) \to 0$ monotonically as $s \to 0$.

3. This says that the point $(x, f(x)) = (\hat{x}+h, f(\hat{x}+h))$ is in the ball centered at $P_h \equiv (\hat{x}+h, f(\hat{x}) + L(h))$ with radius $|g(h)| \leq \alpha(|h|)|h|$.

4. Choose an $\epsilon > 0$. Let δ be any δ that makes $\alpha(\delta) \leq \epsilon$ true. Then for any $h \in B(0,\delta)$ we have that $(\hat{x}+h, f(\hat{x}+h)) \in B(P_h, \epsilon|h|)$.

5. Let $x \in \mathbb{R}^k$, T be a linear subspace of $\mathbb{R}^k$, and $\theta < \frac{\pi}{2}$. Define the cone about T, with vertex at x, and angle θ to be:
$$C_k(x, T, \theta) \equiv \{y \in \mathbb{R}^k | \frac{y-x}{|y-x|} \cdot \frac{v-x}{|v-x|} \geq \cos(\theta) \text{ for some } v \in T + x\}$$

6. The union of the balls in Step 4 is a subset of $C_{n+m}((\hat{x}, f(\hat{x})), T_{(\hat{x}, f(\hat{x}))}, 2\arcsin(\epsilon))$. NOTE: we could do even better since the distance of the P_h from the vertex of the cone is larger than $|h|$, implying that a cone that is skinnier in some places would also work. But the statement is true as it stands.

7. Define
$$\text{Cyl}(\hat{x}, \delta) = B(\hat{x}, \delta) \times \mathbb{R}^m \subset \mathbb{R}^{n+m}.$$

8. We conclude that
$$\{(x,y) \in \mathbb{R}^{n+m} | y = f(x)\} \cap \text{Cyl}(\hat{x}, \delta) \subset C_{n+m}\left((\hat{x}, f(\hat{x})), T_{(\hat{x}, f(\hat{x}))}, 2\arcsin(\epsilon)\right) \cap \text{Cyl}(\hat{x}, \delta).$$

9. One can use what we have derived above to get a slightly different statement: for any $\epsilon > 0$ and for any point $(x, f(x))$ in the graph of f, $G_f \equiv \{(x,y) | x \in \mathbb{R}^n, y = f(x)\}$, where f is differentiable, because $B((x, f(x)), \delta) \subset \text{Cyl}(x, \delta)$, we can conclude that
$$B((x, f(x)), \delta) \cap G_f \subset B((x, f(x)), \delta) \cap C_{n+m}((x, f(x)), T_{(x, f(x))}, 2\arcsin(\epsilon)))$$
where again, δ has been chosen so that $\alpha(\delta) \leq \epsilon$.

Figures 4.1 – 4.4 illustrate this. The first figure shows what happens in the simplest case of $f : \mathbb{R}^1 \to \mathbb{R}^1$ while the last three illustrate the case of $f : \mathbb{R}^2 \to \mathbb{R}^1$.

Exercise 4.3.1. Suppose that $g : \mathbb{R} \to \mathbb{R}$ and $f : \mathbb{R} \to \mathbb{R}$ are differentiable at x_0. Use the fact that differentiability of a function c(x) is equivalent to the existence of a scalar $a = \frac{dc}{dx}$ such that $c(x_0 + h) = c(x_0) + ah + e(h)$ where the function $e(h) \sim o(h)$ (i.e. $e(h)$ is "little o of h"), to prove that fg is differentiable at x_0. (Note: your proof will end up establishing the product rule.

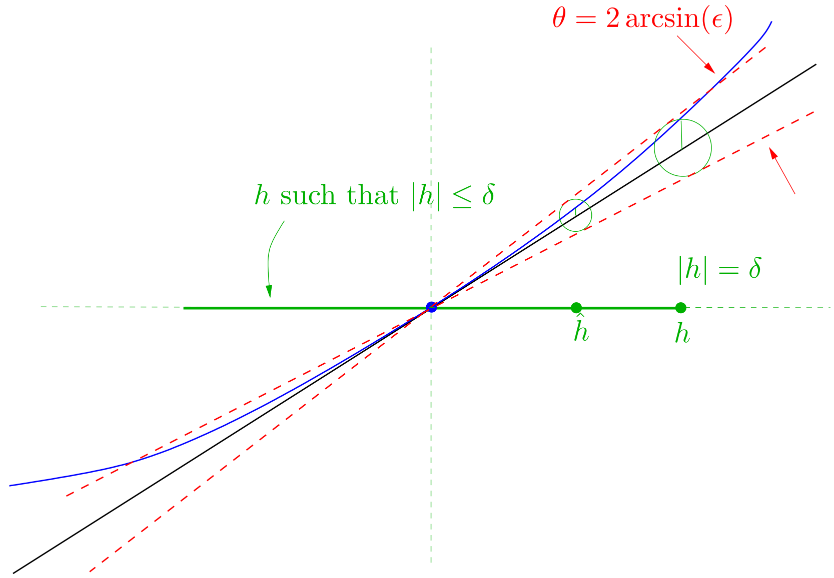

Figure 4.1: **One dimensional case**: Here the vertical axis is $(\Delta_x f)(h) = f(x+h) - f(x)$ and the horizontal axis is the h axis. The cone (in red), centered on the linear approximation to the graph (**black** line), that contains the graph of the function (in blue) can be made as small as we like by choosing δ small enough that $\alpha(\delta)$ is as small as we like. But since $\sin(\theta/2) = \alpha(\delta) \leq \epsilon$ implies that $\theta \leq 2\arcsin(\epsilon)$ and $\lim_{s\to 0}\arcsin(s) = 0$, we get that we can make the angle as small as we like by choosing ϵ small enough.

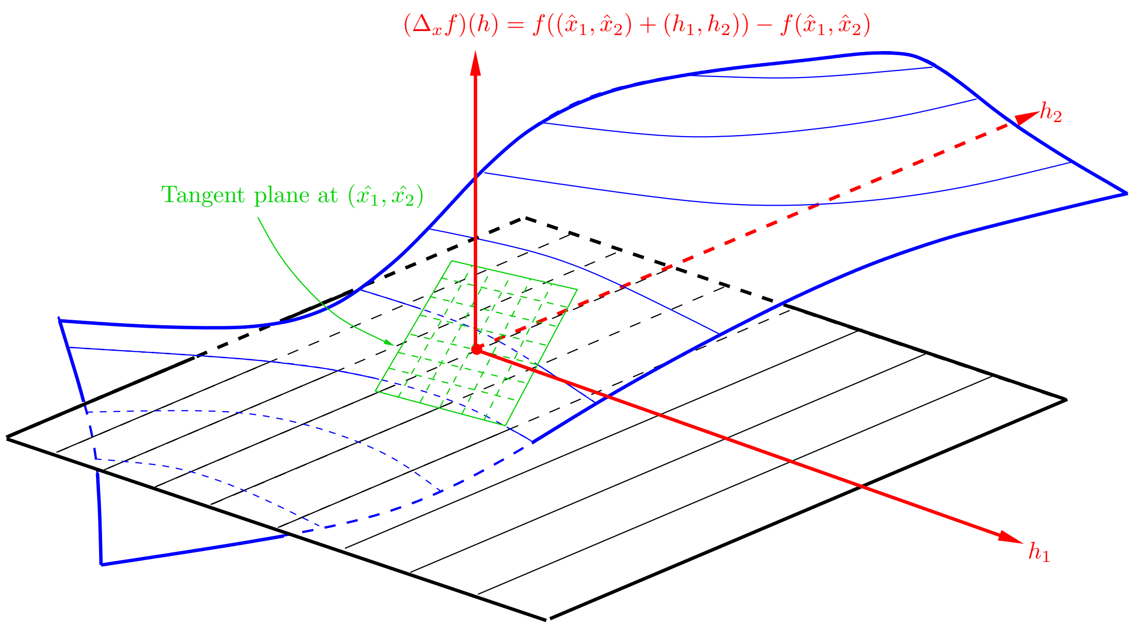

Figure 4.2: **higher dimensional case:** Here a function f mapping from $\mathbb{R}^2$ to $\mathbb{R}$ is depicted. The point at which the linear approximation is being calculated is $\hat{x} = (\hat{x}_1, \hat{x}_2)$.

h_2

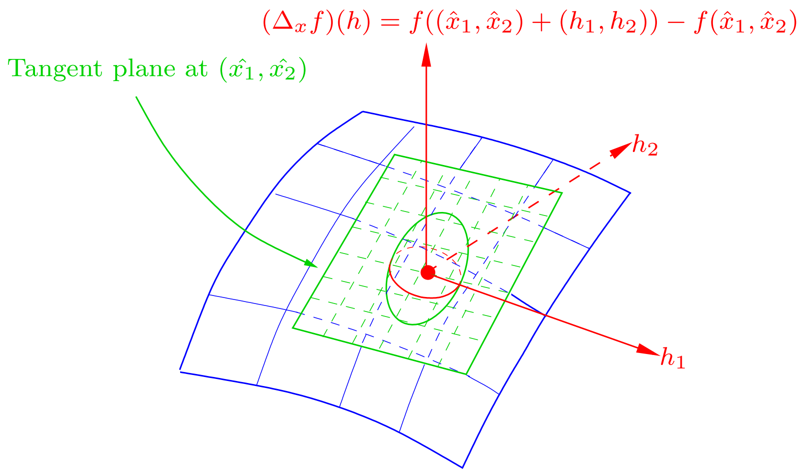

Figure 4.3: **higher dimensional case:** Zooming in a bit. The red ball in the (h_1, h_2)-plane projects to the green ellipsoid in the tangent plane.

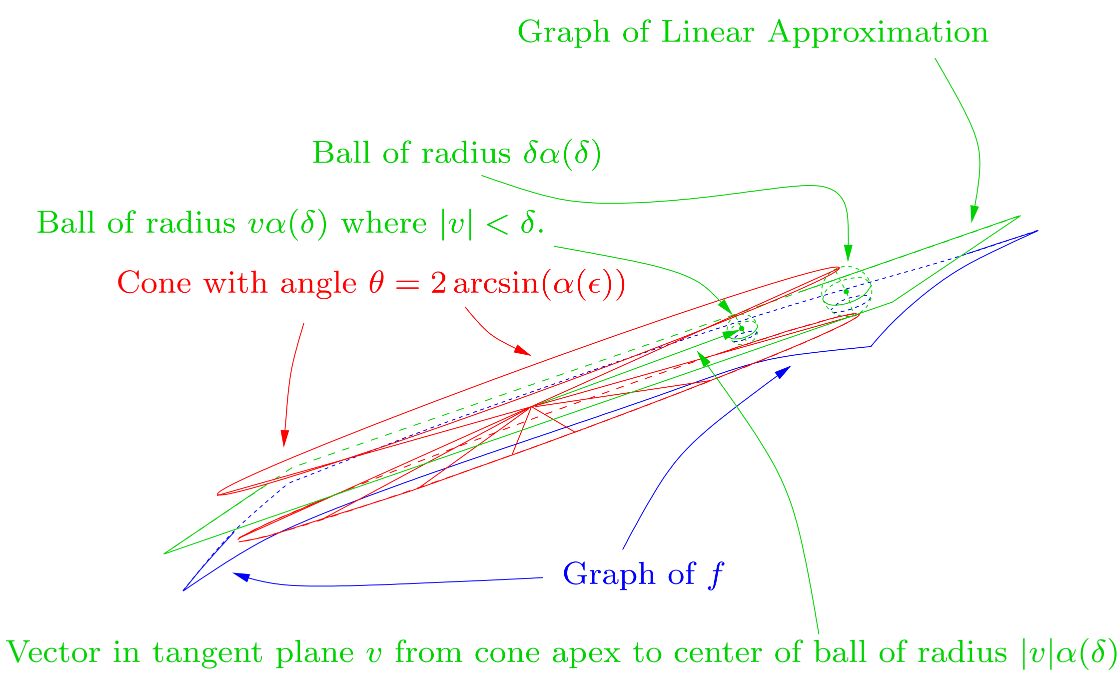

Figure 4.4: **higher dimensional case:** The cone that contains the graph of $L(h) + f(\hat{x})$ and the graph of f.

4.4 Functions of Class C^q

A function that is continuous and differentiable everywhere satisfies the *mean value theorem*:

Theorem 4.4.1 (Mean Value Theorem)**.** *Let $f : \mathbb{R} \to \mathbb{R}$ be continuous and differentiable everywhere. Suppose that $a < b$. Then there is a point $c \in (a, b)$ such that:*

$$f'(c) = \frac{f(b) - f(a)}{(b - a)}$$

Since we have gone over the proof in class, and there is a proof you have read in the appendix (A.2), I won't go over the proof again here. Instead, I will present another result that does not require differentiability, only continuity. First though, we define *upper supporting* and *lower supporting*:

Definition 4.4.1. *Suppose that $f : [a, b] \to \mathbb{R}$. We say that the line $y = f(c) + \alpha(x - c)$ is an* upper supporting *line if $f(c) + \alpha(x - c) \geq f(x)$ for all $x \in [a, b]$ and some $c \in [a, b]$. Likewise, we say the line $y = f(c) + \alpha(x - c)$ is* lower supporting *if $f(c) + \alpha(x - c) \leq f(x)$ for all $x \in [a, b]$ and some $c \in [a, b]$.*

Theorem 4.4.2. *Suppose that $a < b$ and that $f : \mathbb{R} \to \mathbb{R}$ continuous for $x \in [a, b]$. Then there are points $c_1, c_2 \in [a, b]$ such that:*

1. $y = f(c_1) + \left[\frac{f(b)-f(a)}{b-a}\right](x - c_1)$ *is* lower supporting
2. $y = f(c_2) + \left[\frac{f(b)-f(a)}{b-a}\right](x - c_2)$ *is* upper supporting

Proof. The idea of the proof is very similar to the proof of the usual mean value theorem.

1. Suppose that $g(a) = g(b)$, $a < b$. Since g is continuous, g attains a maximum and minimum values, M and m, on the compact set $[a, b]$.
2. Define c_1 to be any point such that $f(c_1) = m$ and c_2 to be any point such that $f(c_2) = M$.
3. Note that the lines $y = g(c_1) + 0(x - c_1)$ and $y = g(c_2) + 0(x - c_2)$ are (respectively) lower and upper supporting lines for f in $[a, b]$
4. Define $g(x) = f(x) - \left[\frac{f(b)-f(a)}{b-a}\right] x$.
5. Note that $g(a) = g(b)$ and thus there are points c_1 and c_2 in $[a, b]$ such that $g(c_1)+0(x-c_1) \leq g(x)$ for all $x \in [a, b]$ and $g(c_2) + 0(x - c_2) \geq g(x)$ for all $x \in [a, b]$.
6. Since $g(x) = f(x) - \left[\frac{f(b)-f(a)}{b-a}\right] x$, those inequalities – $g(c_1) + 0(x - c_1) \leq g(x)$ and $g(c_2) + 0(x - c_2) \geq g(x)$ – translate into the inequalities we started out to prove.

□

4.5 Local Extrema

When the determinant of the second order derivative matrix at x_0 is non-zero, the local behavior of functions $f : \mathbb{R}^n \to \mathbb{R}$ is determined by the first and second derivatives at x_0. In this section, we explore a few ways to prove this. We begin by assuming f has more derivatives than we actually end up needing.

Assuming that a function $f : \mathbb{R}^n \to \mathbb{R}$ has bounded 3rd derivatives, we can show that a non-singular Hessian (the matrix of second derivatives) determines the behavior of the function in the neighborhood of a critical point – a point where the derivative ∇f is 0. (I showed this in class on Monday November 28.)

Some notation:

$$
\begin{array}{lcl}
x, x_0 & = & \text{points in } \mathbb{R}^n \\
h & = & (h^1, h^2, ..., h^n), \text{ a difference vector in } \mathbb{R}^n \\
f_i & = & \frac{\partial f}{\partial x_i} \\
\nabla f & = & (f_1, f_2, ..., f_n) \text{ the gradient vector of } f \\
f_{i,j} & \equiv & \frac{\partial^2 f}{\partial x_j \partial x_i}, \text{ the second order partial derivatives of } f \\
f_{i,j}(x) & = & \text{the second order partial derivative evaluated at } x \\
Q(x,) & = & \text{the matrix } M = M(x) \text{ where } m_{i,j} = f_{i,j}(x) \\
Q(x,h) & \equiv & \langle h, M(x)h \rangle = \sum_{i,j=1}^n f_{i,j}(x) h^i h^j \text{ is the quadratic form containing} \\
 & & \text{the second order information about } f \text{ at the point } x \\
f_{i,j,k} & \equiv & \frac{\partial^3 f}{\partial x_k \partial x_j \partial x_i}, \text{ the third order partial derivatives of } f
\end{array}
$$

Now:

1. Symmetric matrices are diagonalizable by an orthogonal change of basis that is generated by the eigenbasis.

2. Therefore, choosing a point x_0 to study the function f at, we can change the basis so that the matrix $M(x_0)$ is diagonal. I.e.

$$
M(x_0) = \begin{bmatrix} \lambda_1 & 0 & 0 & \dots & 0 \\ 0 & \lambda_2 & 0 & \dots & 0 \\ 0 & 0 & \lambda_3 & \dots & 0 \\ \vdots & \vdots & \vdots & & \vdots \\ 0 & 0 & 0 & \dots & \lambda_n \end{bmatrix}
$$

where the λ_i are the eigenvalues of $M(x_0)$.

3. We conclude that

$$
Q(x,h) = \langle h, M(x)h \rangle = \sum_{i=1}^n \lambda_i (h^i)^2
$$

4. $|h^i| \le |h|$ for all i.

5. if we assume that $\lambda_i > 0$ for all i, then $\lambda_i > \alpha > 0$ for all i for some $\alpha > 0$. This implies that $Q(x,h) > \alpha |h|^2$

6. Choosing an x_0 where $\nabla f(x_0) = 0$, shifting the coordinate system by $(x_0, f(x_0))$ so that we now have that f(0) = 0, and rotating the basis (changing the basis using the eigenbasis of $M(x_0)$), we get that the Taylor series expansion for f at the point 0 is:

$$f(0+h) = \frac{1}{2}Q(0,h) + \frac{1}{6}\sum_{i,j,k=1}^{n} f_{i.j.k}(c)h^i h^j h^k$$

where c is a point on the line segment between 0 and h.

7. Now, since we are assuming that the third order derivatives are bounded in a neighborhood of 0 (x_0 before the shift) by some constant K, we can bound the remainder term:

$$\frac{1}{6}\sum_{i,j,k=1}^{n} f_{i,j,k}(c)h^i h^j h^k \leq \frac{1}{6}n^3 K|h|^3 = C|h|^3$$

where $0 < C < \infty$.

8. We therefore conclude that if all the eigenvalues of the second order matrix are positive at 0 – in Fleming's notation $Q(0,) > 0$ – we have that in some neighborhood of 0

$$f(0+h) \geq \frac{\alpha}{2}|h|^2 - C|h|^3$$

.

9. We conclude that 0 (i.e. x_0) is a local minimum of f. (Use item number 8 to prove this!)

Exercise 4.5.1. Suppose that f has bounded third partial derivatives in some neighborhood of the origin. Show that if the eigenvalues of $Q(x_0,) = M(x_0)$ are all negative, then x_0 is a local maximum.

4.5.1 Assuming only that f is in C^2

Fleming approaches this problem a bit differently, requiring only that the function to has continuous partial derivatives up to order 2 – i.e. $f \in C^2$. Let $\partial B(0,1) \equiv \{h \mid |h| = 1\}$. We need the following results:

Lemma 4.5.1. *$H_1(x) \equiv \min_{h\in\partial B(0,1)} \sum_{i,j} f_{i,j}(x)h^i h^j$ is a continuous function of x.*

Proof. Suppose that $H_1(x)$ is not continuous. Then we can find a sequence of $x_k \to x^*$ such that $H_1(x^*) \neq \lim_{k\to\infty} H_1(x_k)$. Let

$$h_k = \operatorname*{argmin}_{h\in\partial B(0,1)} \sum_{i,j} f_{i,j}(x_k)h^i h^j$$

I.e.

$$H_1(x_k) = \sum_{i,j} f_{i,j}(x_k)h_k^i h_k^j.$$

(Note that $h_k \in \mathbb{R}^n$ might not be unique!.) Since the unit sphere is compact, there is a subsequence of h_k, $h_{k(m)}$ converging to some point $h^* \in B(0,1)$. Recall that $Q(x,h) \equiv \sum_{i,j} f_{i,j}(x)h^i h^j$ and because the $f_{i,j}$'s are continuous, we have that $Q(x,h)$ is continuous.

Consequently we have that the sequence $(x_k, h_k) \in \mathbb{R}^{2n}$ has the property that

$$Q(x_k, h_k) = H_1(x_{k(m)}) = Q(x_{k(m)}, h_{k(m)}) \to_{m\to\infty} Q(x^*, h^*) > H_1(x^*) = Q(x^*, \hat{h})$$

for some $\hat{h} \in \partial B(0,1)$.

But, since $Q(x,h)$ is continuous and

$$(x_{k(m)}, \hat{h}) \to_{m\to\infty} (x^*, \hat{h}),$$

this implies that $Q(x_{k(m)}, h_{k(m)}) > Q(x_{k(m)}, \hat{h})$ for some large enough m which is a contradiction because

$$Q(x_{k(m)}, h_{k(m)}) = H_1(x_{k(m)}) = \min_{h\in\partial B(0,1)} \sum_{i,j} f_{i,j}(x)h^i h^j.$$

□

Lemma 4.5.2. *If $Q(x_0, h) > 0$ for all $h \in \partial B(0,1)$, then $Q(x,h) > 0$ for all $h \in \partial B(0,1)$ and all $x \in B(x_0, \epsilon)$ and some $\epsilon > 0$.*

Proof. Because $H_1(x_0) > 0$ and H_1 is continuous, there is an $\epsilon > 0$ such that $H_1(x) > 0$ for $x \in B(x_0, \epsilon)$. Thus $Q(x,h) > 0$ for all $h \in \partial B(0,1)$ and $x \in B(x_0, \epsilon)$. □

Exercise 4.5.2. Show that $Q(x,h) > 0$ for all $|h| \neq 0$ if and only if $Q(x,h) > 0$ for all $|h| = 1$

Theorem 4.5.1. *Assume that $f \in C^2$. If $(\nabla f)(x_0) = 0$ and $Q(x_0,) > 0$ (equivalently, if $Q(x_0, h) > 0$ for all $|h| \neq 0$) then there is an $\epsilon > 0$ such that $f(x) > f(x_0)$ for $x \in B(x_0, \epsilon)$.*

Proof. The second order Taylor series tells us that

$$f(x) = f(x_0) + (\nabla f)(x_0) \cdot h + \frac{1}{2}\sum_{i,j} f_{i,j}(c)h^i h^j$$

with $h = x - x_0$, $c = x_0 + sh$, and $s \in (0,1)$. Because $(\nabla f)(x_0) = 0$ this reduces to:

$$\begin{aligned} f(x) &= f(x_0) + \frac{1}{2}\sum_{i,j} f_{i,j}(c)h^i h^j \\ &= f(x_0) + \frac{1}{2}Q(c,h) \end{aligned}$$

and if $Q(x_0, h) > 0$ for all $|h| \neq 0$, Lemmas 4.5.1 and 4.5.2 show that $Q(c,h) > 0$ for all $|h| \neq 0$ and $c \in B(x_0, \epsilon)$. This implies that

$$f(x) > f(x_0) \ \ \forall x \in B(x_0, \epsilon).$$

□

4.5.2 A Slightly More Geometric Approach

Here is a somewhat more geometric approach to Lemma 4.5.2:

1. Because $Q(x_0, h) > 0$ for $h \in \partial B(0,1)$ and $Q(x,h)$ is continuous, then for every $h^* \in \partial B(0,1)$ there is an $\epsilon(h^*) > 0$ and a ball $B((x_0, h^*), \epsilon(h^*)) \subset \mathbb{R}^{2n}$ such that for every $(x,h) \in B((x_0, h^*), \epsilon(h^*))$, $Q(x,h) > 0$.

2. The union of these balls is an open set U, containing the compact set $\mathcal{S} \equiv \{x_0\} \times \partial B(0,1) \subset \mathbb{R}^{2n}$.

3. U^c is closed and $\mathcal{S}$ is compact; this implies that the distance from U^c to $\mathcal{S}$ is bounded below by some $\delta > 0$.

   Proof. (a) Since every point $s \in \mathcal{S}$ is contained in an open ball centered on s that is disjoint from U^c, $d(s, U^c) \equiv \inf_{y \in U^c} |y - s| > 0$.

   (b) Use the the triangle inequality proves that $f(s) \equiv d(s, U^c)$ is a continuous function on $\mathcal{S}$.

   (c) Because $\mathcal{S}$ is compact and because $f(s) > 0$ for every $s \in \mathcal{S}$, f attains a minimum at some point $s_m \in \mathcal{S}$ and $\delta \equiv f(s_m) > 0$.

   □

4. Therefore $\bigcup_{s \in \mathcal{S}} B((s, \delta) \subset U$.

5. Note that for every $s = (x_0, h) \in \mathcal{S}$, we have that for $\{x \mid |x - x_0| < \delta\}$, $|(x,h) - (x_0,h)| < \delta$.

6. This implies that $Q(x,h) > 0$ for all $h \in \partial B(0,1) \subset \mathbb{R}^n$ and $x \in B(x_0, \delta) \subset \mathbb{R}^n$.

4.5.3 A $o(|h|^2)$ Approach

Another approach again assumes that $f \in C^2$ and uses the Taylor series for $q = 2$, but then expresses each of the partial derivatives using the fact that they are continuous in x. I.e

$$f_{i,j}(c) = f_{i,j}(x_0) + g_{i,j}(c)$$

where $g_{i,j}(c) \to 0$ as $c \to x_0$. This leads to

$$f(x) = f(x_0) + (\nabla f)(x_0) \cdot h + \frac{1}{2} \sum_{i,j} f_{i,j}(x_0) h^i h^j + \frac{1}{2} \sum_{i,j} g_{i,j}(c) h^i h^j$$

with $h = x - x_0$, $c = x_0 + sh$, and $s \in (0,1)$.

Because $(\nabla f)(x_0) = 0$ this reduces to:

$$f(x) = f(x_0) + \frac{1}{2} \sum_{i,j} f_{i,j}(x_0) h^i h^j + \frac{1}{2} \sum_{i,j} g_{i,j}(c) h^i h^j$$

Again, as above, because $Q(x_0,) > 0$ we have that

$$\frac{1}{2} \sum_{i,j} f_{i,j}(x_0) h^i h^j > \alpha |h|^2$$

for some $\alpha > 0$. Now choose $\epsilon > 0$ such that $|c - x_0| < \epsilon$ implies that $g_{i,j}(c) < \frac{\alpha}{2n^2}$ for all i and j. Because $x \in B(x_0, \epsilon)$ implies that $c \in B(x_0, \epsilon)$, we have that

$$x \in B(x_0, \epsilon)$$

implies that

$$\begin{aligned} f(x) &= f(x_0) + \frac{1}{2}\sum_{i,j} f_{i,j}(x_0)h^i h^j + \frac{1}{2}\sum_{i,j} g_{i,j}(c)h^i h^j \\ &\geq f(x_0) + \alpha|h|^2 - \frac{\alpha}{2n^2}n^2|h|^2 \\ &= f(x_0) + \frac{\alpha}{2}|h|^2 \\ &> f(x_0) \end{aligned}$$

Remark 4.5.1. *Note that because*

$$f_{i,j}(c) = f_{i,j}(x_0) + g_{i,j}(c)$$

and $g_{i,j}(c) \to 0$ *as* $c \to x_0$*, we can express*

$$f(x) = f(x_0) + (\nabla f)(x_0) \cdot h + \frac{1}{2}\sum_{i,j} f_{i,j}(x_0)h^i h^j + \frac{1}{2}\sum_{i,j} g_{i,j}(c)h^i h^j$$

as

$$f(x) = f(x_0) + (\nabla f)(x_0) \cdot h + \frac{1}{2}\sum_{i,j} f_{i,j}(x_0)h^i h^j + o(|h|^2),$$

explaining the title of this subsection.

4.5.4 When $Q(X_0,) = M(X_0)$ has non-zero determinant (and when it doesn't)

The results we have obtained above make it simple to conclude that when $(\nabla f)(x_0) = 0$ and $\det(M(x_0)) \neq 0$, either is either a local maximum, a local minimum or a saddle point. Exercise 4.5.1 showed that $Q(x_0,) < 0$ implies that $f(x) < f(x_0)$ for all $x \in B(x_0, \epsilon)$, though it assumed that $f \in C^3$.

Exercise 4.5.3. Assuming that $f \in C^2$, show that there is an $\epsilon > 0$ such that $Q(x_0,) < 0$ implies that $f(x) < f(x_0)$ for all $x \in B(x_0, \epsilon)$.

We know that by changing the basis of the domain at x_0 using the eigenvectors of $M(x_0)$ as basis vectors, we get that

$$Q(x_0, h) = \sum_{i=1}^{n} \lambda_i (h^i)^2$$

If $\det(M(x_0)) \neq 0$, then none of the eigenvalues are equal to 0 and if some are positive and some are negative, then we have the following:

Exercise 4.5.4. Assuming that $f \in C^2$, show that if (a) $\det(M(x_0)) \neq 0$, (b) some $\lambda_i > 0$ and (c) some $\lambda_i < 0$, then x_0 is a saddle point, i.e. in some directions f has a local minimum at x_0 and in other directions it has a local maximum at x_0.

Exercise 4.5.5. Find examples of functions $f : \mathbb{R}^2 \to \mathbb{R}$ such that $(\nabla f)(0) = 0$, all the eigenvalues of $M(0)$ are non-negative, and:

1. $f(x) < f(x_0)$ for some $h \in \partial B(0,1)$ and all $x = x_0 + sh$ where $s \in (-\epsilon, 0) \cup (0, \epsilon)$.
2. $f(x) < f(x_0)$ for all $x \in B(0, \epsilon)$ for some some $\epsilon > 0$.

4.6 Convex and Concave Functions

Convex and concave functions are the nicest that functions can get and still be non-linear. That is, if $f : \mathbb{R}^n \to \mathbb{R}$ is not linear, but it is convex, then many properties are still very nice.

Definition 4.6.1 (epigraph)**.** *A epigraph of a function* $f : \mathbb{R}^n \to \mathbb{R}$, $\text{epi}(f)$ *is all the region in the graph space above the graph of* f*:* $\text{epi}(f) \equiv \{(x, y) \in \mathbb{R}^{n+1} | y \geq f(x)\}$.

Definition 4.6.2 (epograph)**.** *A epigraph of a function* $f : \mathbb{R}^n \to \mathbb{R}$, $\text{epo}(f)$ *is all the region in the graph space below the graph of* f*:* $\text{epo}(f) \equiv \{(x, y) \in \mathbb{R}^{n+1} | y \leq f(x)\}$.

Definition 4.6.3 (Convex Function: Definition 1)**.** *A function* $f : \mathbb{R}^n \to \mathbb{R}$ *is convex if* $\text{epi}(f)$ *is convex.*

Definition 4.6.4 (Convex Function: Definition 2)**.** *A function* $f : \mathbb{R}^n \to \mathbb{R}$ *is convex if*

$$f(\alpha x + \beta y) \leq \alpha f(x) + \beta f(y)$$

for all $0 \leq \alpha, \beta \leq 1$ *such that* $\alpha + \beta = 1$.

Exercise 4.6.1. Prove that the two different definitions of *Convex Function* are equivalent.

Definition 4.6.5 (Concave Function: Definition 1)**.** *A function* $f : \mathbb{R}^n \to \mathbb{R}$ *is concave if* $\text{epo}(f)$ *is convex.*

Definition 4.6.6 (Concave Function: Definition 2)**.** *A function* $f : \mathbb{R}^n \to \mathbb{R}$ *is concave if*

$$f(\alpha x + \beta y) \geq \alpha f(x) + \beta f(y)$$

for all $0 \leq \alpha, \beta \leq 1$ *such that* $\alpha + \beta = 1$.

Exercise 4.6.2. Prove that the two different definitions of *Concave Function* are equivalent.

Exercise 4.6.3. Show that the only functions that are both convex and concave are linear functions.

Exercise 4.6.4. Suppose that $f : \mathbb{R}^n \to \mathbb{R}$ is a convex function. Show that if x and y are both local minima of f, then

1. $f(x) = f(y)$,
2. for all $w = \alpha x + (1 - \alpha)y$, where $0 \leq \alpha \leq 1$, $f(w) = f(x)$,
3. and therefore, every local minimum is a global minimum.

Exercise 4.6.5. Use Exercise 4.6.4 to show that any nonempty set of minimizers of a convex function f, is a convex set.

Exercise 4.6.6. Give an example of a convex function that is bounded from below, but has no minimum value.

Exercise 4.6.7. Give an example of a convex function that is not bounded from below or above.

Definition 4.6.7 (Supporting Hyperplane)**.** *Let $f : \mathbb{R}^n \to \mathbb{R}$ be convex. We say that $g = (g_1, g_2, ..., g_n) \in \mathbb{R}^{n*}$ ($\mathbb{R}^{n*}$ is the dual space of $\mathbb{R}^n$) defines a supporting hyperplane,*

$$h_g(x^*) \equiv \{(x, y)|y = \langle g, x - x^* \rangle\},$$

of f at x^* if

$$f(x) \geq f(x^*) + \langle g, x - x^* \rangle$$

for all $x \in \mathbb{R}^n$.

Exercise 4.6.8. Let $f : \mathbb{R}^n \to \mathbb{R}$ be convex. Show that the set of $g \in \mathbb{R}^{n*}$, such that $h_g(x^*)$ are supporting hyperplanes of f at x^*, is both closed and convex. We will denote this set of g by $\partial_{x^*} f$.

Definition 4.6.8 (Left and Right Derivatives)**.** *Let $f : \mathbb{R} \to \mathbb{R}$ be convex. We define the left and right derivatives at x^* to be the limits*

$$\frac{df}{dx}^l \equiv \lim_{x \uparrow x^*} \frac{f(x) - f(x^*)}{x - x^*}$$

and

$$\frac{df}{dx}^r \equiv \lim_{x \downarrow x^*} \frac{f(x) - f(x^*)}{x - x^*}.$$

In exercises 4.6.9 to 4.6.13 we will assume that $f : \mathbb{R} \to \mathbb{R}$ and f is convex.

Exercise 4.6.9. Show that for every $w \in \mathbb{R}$, $\partial_w f$ is a closed bounded interval. Now show that the endpoints of $\partial_w f$ correspond to the left and right derivatives of f at w.

Definition 4.6.9. *(Notation - Interior) We denote the interior of a set E by E^o.*

Exercise 4.6.10. Why is f differentiable at x if $\partial_x f = [s, s]$, i.e. the closed interval consisting of a single point? Show that f is not differentiable at x if and only if $(\partial_x f)^o \neq \emptyset$.

Exercise 4.6.11. Show that $(\partial_u f)^o \cap (\partial_v f)^o$ for all $u, w \in \mathbb{R}$, $u \neq w$.

Exercise 4.6.12. Show that $F_j \equiv \{x|$ f is not differentiable at x$\}$ is at most countably infinite. The j in F_j stands for "jump" – explain why I use the term jump for this set. Hint: plot the derivative of a function f for which F_j is nonempty.

Definition 4.6.10 (Measure Zero Sets)**.** *We will say that $E \in \mathbb{R}$ has 1-dimensional measure zero if, for any $\epsilon > 0$, we can find a collection of intervals I_i of lengths d_i such that $E \subset \bigcup_i I_i$ and $\sum_i d_i < \epsilon$.*

Exercise 4.6.13. Show that F_j defined in Exercise 4.6.12 has measure zero. This implies that outside a set of measure zero, any convex function from $\mathbb{R}$ to $\mathbb{R}$ is differentiable. This is commonly written as f is differentiable almost everywhere or f is differentiable a.e.

Exercise 4.6.14. Show that for a convex $f : \mathbb{R}^n \to \mathbb{R}$, f is differentiable at x if and only if the set $\partial_x f$ is a single point.

Definition 4.6.11 (Convex Envelope)**.** *Define $\mathcal{C}_f$ to be the set of all convex functions c such that $c(x) \leq f(x)\ \forall\ x \in \mathbb{R}^n$. We define the Convex Envelope* $\operatorname{cnv}(f)$*,*

$$\operatorname{cnv}(f)(x) \equiv \sup \{f_\eta(x)\ | f_\eta \in \mathcal{C}_f,\ f_\eta(x) \leq f(x) \forall x \in \mathbb{R}^n\}$$

Exercise 4.6.15. Assume that $C < f(x)$ for all x for some $C > -\infty$. Show that $\operatorname{cnv}(f)$ is convex.

Exercise 4.6.16. Challenge Problem: Suppose that $\epsilon > 0$. Let us call a function a *ϵ-approximate convex function* if for all $x, y \in \mathbb{R}$,

$$f(\alpha x + (1 - \alpha)y) \leq \alpha f(x) + (1 - \alpha)f(y) + \epsilon.$$

Suppose that f is ϵ-approximate convex. Prove: $|\operatorname{cnv}(x) - f(x)| \leq \epsilon$ for all $x \in \mathbb{R}$

Exercise 4.6.17. Challenge Problem: Suppose that $\epsilon > 0$. Define the *ϵ-subdifferential* $\partial^\epsilon_{x^*} f$, of f at x^* to be the set of g such that

$$f(x) + \epsilon \geq f(x^*) + \langle g, x - x^* \rangle$$

for all $x \in \mathbb{R}^n$. Prove: $\partial \operatorname{cnv}(f)_x \subset \partial^\epsilon_x f$

Chapter 5

Properties of Differentiable Functions

5.1 Linear Transformations

Only comment on this section is that you should have a book on linear algebra handy. You should also review the *singular value decomposition* (SVD) and learn to use this when reasoning about linear transformations. See also the Appendix Linear Algebra section in these notes.

5.2 Affine Transformations

Exercise 5.2.1. Show that if $T : \mathbb{R}^n \to \mathbb{R}^n$ is a transformation that is one-to-one and onto (i.e. bijective), so that there is an inverse of T, T^{-1}, **then** $T \circ T^{-1} = I_n$ **and** $T^{-1} \circ T = I_n$, where $I_n : \mathbb{R}^n \to \mathbb{R}^n$ is the identity transformation on $\mathbb{R}^n$. Applying this to matrices, show that an invertible matrix commutes with its inverse.

Exercise 5.2.2. Look up a proof, or create a proof of the following facts: if A and B are $n \times n$ matrices, then

1. $\det(A) = \det(A^t)$,
2. $\det(AB) = \det(A)\det(B)$.

5.3 Differentiable Transformations

Definition 5.3.1. *(little o of h, $o(h)$) Suppose that $f : \mathbb{R}^n \to \mathbb{R}^m$. We say that*

$$f(h) \in o(h)$$

if

$$\lim_{h \to 0} \frac{f(h)}{|h|} = 0$$

.

A mapping F is said to be *differentiable at x* if there is a good linear approximation to the mapping at x. More precisely, we say that:

Definition 5.3.2. *(Differentiability and Derivatives)* $F : \mathbb{R}^n \to \mathbb{R}^m$ *is* differentiable *at* x *if there is a linear map,* L_F^x *such that*

$$F(x+h) = F(x) + L_F^x(h) + o(h).$$

Often $L_F^x(h)$ *is denoted by* $DF(x)(h)$ *and once a basis is chosen,* $DF(x)$ *is represented by an* $m \times n$ *matrix. Another equivalent way to say this is that* F *is differentiable at* x *if there is a linear transformation* L_F^x *such that*

$$\lim_{|h|\to 0} \frac{f(x+h) - F(x) - L_F^x(h)}{|h|} \to 0$$

5.4 Compositions

Theorem 5.4.1 (Composite Function Theorem)**.** *If* g *is differentiable at* t_0 *and* f *is differentiable at* $x_0 \equiv g(t_0)$*, then* $F \equiv f \circ g$ *is differentiable at* t_0 *and* $DF(t_0) = Df(x_0) \circ Dg(t_0)$*.*

Proof. In this proof, I will drop the 0 subscript on t and x.

$$\begin{aligned}
F(t+h) &= f(g(t+h)) \\
&= f(g(t) + L_g^t(h) + k(h)) \\
&= f(g(t)) + L_f^{g(t)}(L_g^t(h) + k(h)) + l(L_g^t(h) + k(h)) \\
&= F(t) + L_f^{g(t)}(L_g^t(h)) + L_f^{g(t)}(k(h)) + l(L_g^t(h) + k(h))
\end{aligned}$$

where $k(h), l(h) \in o(h)$. If we assume that for $k(h), l(h), w_1(h), w_2(h) \in o(h)$ we have

1. $L_f^{g(t)}(k(h)) \in o(h)$
2. $l(L_g^t(h) + k(h)) \in o(h)$
3. $w_1(h) + w_2(h) \in o(h)$

we get that

$$\begin{aligned}
F(t+h) &= F(t) + L_f^{g(t)}(L_g^t(h)) + L_f^{g(t)}(k(h)) + l(L_g^t(h) + k(h)) \\
&= F(t) + L_f^{g(t)}(L_g^t(h)) + o(h)) \\
&= F(t) + \left(Df(g(t)) \circ Dg(t)\right)(h) + o(h)
\end{aligned}$$

and we are done. To finish, we need to show the (1-3) above. First some exercises.

Exercise 5.4.1. Show that if $k(h) \in o(h)$, then $C\, k(h) \in o(h)$.

Exercise 5.4.2. Show that $k(h), l(h) \in o(h)$ implies $k(h) + l(h) \in o(h)$.

Exercise 5.4.3. Show that if $A : \mathbb{R}^n \to \mathbb{R}^m$ is a linear transformation, then $|A(y)| \leq C|y|$ for some $C < \infty$.

Moving on with the proof:

1. Because $L_f^{g(t)}$ is a linear transformation, the above exercises imply that $|L_f^{g(t)}(y)| \leq C|y|$, so that $L_f^{g(t)}(k(h)) \leq Ck(h)$ which implies that

$$L_f^{g(t)}(k(h)) \in o(h).$$

2. Now we show that $l(L_g^t(h) + k(h)) \in o(h)$. Observe that

$$\left\{ \frac{l(h)}{|h|} \underset{|h|\to 0}{\to} 0 \right\} \Rightarrow \left\{ \frac{l(L_g^t(h) + k(h))}{|L_g^t(h) + k(h)|} \underset{|L_g^t(h)+k(h)|\to 0}{\to} 0 \right\}.$$

   We note also that $\frac{|L_g^t(h)+k(h)|}{|h|} < C < \infty$, implying $\{|h| \to 0\} \Rightarrow \{|L_g^t(h) + k(h)| \to 0\}$. As a result, we have that

$$\frac{l(L_g^t(h) + k(h))}{|h|} = \frac{l(L_g^t(h) + k(h))}{|L_g^t(h) + k(h)|} \cdot \frac{|L_g^t(h) + k(h)|}{|h|} \quad (5.1)$$

$$\leq C \frac{l(L_g^t(h) + k(h))}{|L_g^t(h) + k(h)|} \underset{|h|\to 0}{\to} 0 \quad (5.2)$$

   using $\{|h| \to 0\} \Rightarrow \{|L_g^t(h) + k(h)| \to 0\}$.

   Thus, $l(L_g^t(h) + k(h)) \in o(h)$.

□

5.5 Inverse Function Theorem

Theorem 5.5.1 (Inverse Function Theorem). *Suppose that $F : \mathbb{R}^n \to \mathbb{R}^n$, $F \in C^1(\mathbb{R}^n)$, and $D_{x_0}F$ is invertible. Then F^{-1} exists in a neighborhood of x_0, $D_{x_0}(F^{-1}) = (D_{x_0}F)^{-1}$ and F^{-1} is also C^1 in some neighborhood of $F(x_0)$. (In fact, if $F \in C^q(\mathbb{R}^n)$ then $F^{-1} \in C^q(\mathbb{R}^n)$.)*

Proof. The first thing to note is that intuitively, this is reasonable, since, if the linear approximation to F at x_0, $D_{x_0}F$, is invertible at x_0, it is reasonable to think the thing that $D_{x_0}F$ approximates, F, is also locally invertible at x_0.

1. Recall that for $a \in \mathbb{R}$, $|a| < 1$, $\sum_{i=0}^{\infty} a^i = 1 + a + a^2 + a^3 + \cdots = \frac{1}{1-a}$.

2. Without any loss of generality, we can assume that $x_0 = 0$ and $F(0) = 0$. If not, then translate in the domain and range: $\tilde{F} \equiv F(x + x_0) - F(x_0)$.

3. Define $G \equiv I - (D_0F)^{-1} \circ F$ and let $\epsilon = \frac{1}{2}$.

4. Calculating, we see that $D_0G = 0$ and since $G \in C^1$, we have that $||D_xG|| < \epsilon$ for $x \in B(0, \delta(\epsilon)) \subset \mathbb{R}^n$. To see this note that the operator norm of a matrix is less than the sum of the absolute values of its entries. Since at $x = 0$ all the entries of the derivative of G are 0, and they are all continuous, we get the existence of that open $\delta(\epsilon)$ ball in which $||DG||$ is less than ϵ.

5. The mean value theorem gives that G is a contraction mapping in $B(0, \delta(\epsilon))$, i.e. $|G(h) - G(k)| < \epsilon|h-k|$ for all $h, k \in B(0, \delta(\epsilon))$. Apply the mean value theorem to $g(s) \equiv G(k+s(h-k))$. We get that $g(1)-g(2) = g'(t)\cdot(1-0)$ for some $t \in (0,1)$. But $g'(t) = D_{k+t(h-k)}G\cdot(h-k)$ so we get that

$$\begin{aligned} |G(h) - G(k)| &= |g(1) - g(0)| \\ &= |g'(t) \cdot (1-0)| \\ &= |D_{k+t(h-k)}G \cdot (h-k)| \\ &\leq ||D_{k+t(h-k)}G||\ |h-k| \\ &< \epsilon|h-k| \end{aligned}$$

6. Define $H = I + G + G \circ G + g \circ G \circ G + \cdots = \sum_{i=0}^{\infty} G^i$.

7. Note that $|H(h)| < 2|h|$. We see this by observing that:

$$\begin{aligned} |H(h)| &= |h + G(h) + G^2(h) + \cdots| \\ &\leq |h| + |G(h)| + |G^2(h)| + \cdots \\ &< |h|(1 + \epsilon + \epsilon^2 + \cdots) \\ &= \frac{|h|}{1-\epsilon} \\ &= 2|h| \end{aligned}$$

8. $DH = I + DG + DG \circ DG + \cdots = \sum_{i=0}^{\infty} DG^i$ exists for all $x \in B(0, \delta(\epsilon))$. Note: if H were a finite sum, this would follow immediately from the fact that the derivative of a finite sum of functions is the sum of the derivatives of the functions. The claim that $DH = \sum_{i=0}^{\infty} DG^i$ is equivalent to the claim that for x and $x+h$ in $B(0, \delta(\epsilon))$

$$H(x+h) - H(x) - (\sum_{i=0}^{\infty} DG^i)(h) \in o(h).$$

We split this into 3 pieces:

$$\begin{aligned} H(x+h) - H(x) - (\sum_{i=0}^{\infty} DG^i)(h) &= (\text{Term } 1) \left\{ \sum_{i=0}^{N} G^i(x+h) - \sum_{i=0}^{N} G^i(x) - (\sum_{i=0}^{N} DG^i)(h) \right\} \\ &+ (\text{Term } 2) \left\{ \sum_{i=N+1}^{\infty} (G^i(x+h) - G^i(x)) \right\} \\ &- (\text{Term } 3) \left\{ (\sum_{i=N+1}^{\infty} DG^i)(h) \right\} \end{aligned}$$

To show that $\{\text{Term } 1 + \text{Term } 2 + \text{Term } 3\} \in o(h)$, we must show that for any $\eta > 0$ there is a δ_η such that

$$|h| < \delta_\eta \Rightarrow \{\text{Term } 1 + \text{Term } 2 + \text{Term } 3\} < \eta.$$

(a) Choose $0 < \eta$.

(b) Choose N big enough that $\frac{\epsilon^{N+1}}{1-\epsilon} < \frac{\eta}{3}$.

(c) Choose $\tilde{\delta}$ small enough that $\{\text{Term 1}\} < \frac{\eta}{3}|h|$ which we can do since a finite sum of differentiable functions is differentiable.

(d) Note also that because G is a contraction, we get that

$$|G^i(x+h) - G^i(x)| < \epsilon^i |h|$$

for x and $x+h$ in $B(0, \delta(\epsilon))$.

(e) Define

$$\delta_\eta = \min((\delta(\epsilon) - |x|), \tilde{\delta}).$$

(f) Using $\sum_{i=N+1}^{\infty} \epsilon^i = \frac{\epsilon^{N+1}}{1-\epsilon}$, we calculate:

$$\begin{aligned} \{\text{Term 1 + Term 2 + Term 3}\} &< \frac{\eta}{3}|h| + \frac{\epsilon^{N+1}}{1-\epsilon}|h| + \frac{\epsilon^{N+1}}{1-\epsilon}|h| \\ &< \eta|h| \end{aligned}$$

whenever $|h| < \delta_\eta$.

9. Define $\delta \equiv \frac{\delta(\epsilon)}{2}$.

10. $H \circ (I - G) = I$ and $(I - G) \circ H = I$ on $B(0, \delta)$. Why $\delta = \frac{\delta(\epsilon)}{2}$? We have to make sure that the output of H and the output of $I - G$ stays in the set of points where G is a contraction mapping, so, since H can double the size of whatever you stick in H (and $I - G$ can increase the size by a factor of at most $\frac{3}{2}$), we restrict ourselves to the ball $\frac{1}{2}$ the size of the ball on which G is a contraction mapping.

11. This implies that $H \circ (D_0F)^{-1} \circ F = I$ and $F \circ H \circ (D_0F)^{-1} = I$. In other words, on

$$F(B(0, \delta)),$$

we have that

$$F^{-1} = H \circ (D_0F)^{-1}.$$

Simply compute it.

12. Also, from Step 8 we have that H is differentiable so that

$$DF^{-1} = DH \circ (D_0F)^{-1}$$

and

$$D_{F(x)}(F^{-1}) = (D_xF)^{-1}.$$

We have that

(a) $H \circ (D_0F)^{-1} \circ F = I$ implies that $DH \circ (D_0F)^{-1} \circ DF = I$ on $B(0, \delta)$.

(b) Likewise we get $DF \circ DH \circ (D_0F)^{-1} = I$ on $F(B(0, \delta))$.

(c) We conclude that $DF^{-1} = DH \circ (D_0F)^{-1}$ and $D(F^{-1}) \circ DF = I$ and $DF \circ D(F^{-1}) = I$.

Note that H being differentiable at all $x \in B(0, \delta(\epsilon))$ is enough since

$$(D_0F)^{-1}(F(B(0,\delta))) \subset B(0, \delta(\epsilon)).$$

We have that

$$DF^{-1} = DH \circ (D_0F)^{-1}$$

for all $x \in F(B(0,\delta))$.

13. Recall that for any invertible n×n matrix A,

$$A^{-1} = \frac{1}{\det(A)} adj(A)$$

where

   (a) $adj(A) = C^T$
   (b) C, the cofactor matrix, has elements $C_{i,j} = (-1)^{i+j} M(A, i, j)$
   (c) $M(A, i, j)$ is the determinant of the (n-1)×(n-1) matrix obtained by removing the ith row and the jth column of A.

14. This implies that if, for $x \in E$, E open, $\det(DF) \neq 0$, $(D_xF)^{-1}$ has the same differentiability as D_xF does. Because $\det(D_xF)$ is a polynomial of elements of the matrix D_xF that is nonzero when $x \in E$, $(DF)^{-1} = \frac{1}{\det(DF)} adj(DF)$ implies that

$$(DF)^{-1} = \frac{\text{Matrix of polynomials of elements of DF}}{\text{non-zero polynomial of elements of DF}}.$$

This gives us the result we want.

15. We therefore have that all the partial derivatives of $D_{F(x)}(F^{-1}) = (D_xF)^{-1}$ with respect to x are continuous up to order $q-1$.

16. This implies that $D_y(F^{-1}) = (D_{F^{-1}(y)}F)^{-1}$ is continuous in y as long as $F^{-1}(y)$ is continuous. But Step 8 implies that F^{-1} is continuous. This gives us that F^{-1} is in C^1: all the partial derivatives of $F^{-1}(y)$ with respect to the y_i are continuous. Using the fact that $D_y(F^{-1}) = (D_{F^{-1}(y)}F)^{-1}$ and iterating, we get the result we want. Namely ...

17. We conclude that $F \in C^q \to F^{-1} \in C^q$.

□

5.6 Implicit Function Theorem

Theorem 5.6.1 (Implicit Function Theorem). *Suppose that $F : \mathbb{R}^n \to \mathbb{R}^m$, $m < n$ is C^q in a neighborhood of $a \in \mathbb{R}^n$. Suppose also that D_aF is full rank and that, without loss of generality, that the first m columns of D_aF are independent. Express $x \in \mathbb{R}^n$ as $x = (x_1, x_2)$ where $x_1 \in \mathbb{R}^m$ and $x_2 \in \mathbb{R}^{n-m}$, so that $a = (a_1, a_2)$. Then there exists a neighborhood of $a_2 \in \mathbb{R}^{n-m}$, U_{a_2}, and a function $g_a : \mathbb{R}^{n-m} \to \mathbb{R}^m$, such for all*

$$y \in U_{a_2} \Rightarrow F(g_a(y), y) = F(a).$$

Proof. We prove the theorem in stages:

We assume $F(0) = 0$*.* Note that if $F(a) = b$, then $\tilde{F}(x) \equiv F(a+x) - b$ satisfies $\tilde{F}(0) = 0$ and $D_0\tilde{F} = D_aF$. So we can assume that $F(0) = 0$. □

To illuminate the Implicit Function Theorem, we look at the case of a linear function in detail.

The Linear Case.

1. If $F : \mathbb{R}^n \to \mathbb{R}^m$ is linear, then $DF = F$.
2. From the assumptions in the theorem, the first m columns of F are linearly independent. Write $DF = F = [F_1\ F_2]$ in block form, where the first m columns of F form the m×m matrix F_1 and the last *n-m* columns form the m×(n-m) matrix F_2.
3. Note that $Fx = F_1x_1 + F_2x_2$ and that $F_1x_1 = -F_2x_2$ if and only if $Fx = 0$.
4. Since F_1 is invertible, we get
$$x_1 = -F_1^{-1}F_2x_2 \text{ if and only if } Fx = 0.$$
5. So
$$x = (x_1, x_2) = (-F_1^{-1}F_2x_2, x_2) \text{ for any } x_2 \in \mathbb{R}^{n-m}$$
if and only if
$$F(x) = 0.$$

In the language of the theorem $g_0(x_2) = -F_1^{-1}F_2x_2$ so that
$$F(g_0(x_2), x_2) = F(-F_1^{-1}F_2(x_2), x_2) = 0$$
for not just some neighborhood of $x_2 \in \mathbb{R}^{n-m}$ but for all $x_2 \in \mathbb{R}^{n-m}$. □

The Nonlinear Case. The basic idea will be to embed F in a mapping from $\mathbb{R}^n$ to $\mathbb{R}^n$ and then use the inverse function theorem.

1. In fact, the embedding is a simple one:
$$G : x = (x_1, x_2) \in \mathbb{R}^n \to y = (y_1, y_2) \in \mathbb{R}^n$$
is defined by
$$y = (y_1, y_2) = (F(x_1, x_2), x_2).$$
2. Denote the first m columns of DF by D_1F and the last *n-m* columns by D_1F. (We are suppressing the subscript that indicates where the derivative is being evaluated.)
3. With this notation, we get that the derivative of G, DG is given by:
$$DG = \begin{array}{c} \\ \text{m} \\ \text{n-m} \end{array} \begin{array}{c} \begin{array}{cc} \text{m} & \text{n-m} \end{array} \\ \left[\begin{array}{ll} DF_1 & DF_2 \\ 0 & I_{n-m} \end{array} \right] \end{array}$$
where I_{n-m} denotes the (n-m)×(n-m) identity matrix
4. Computing, we get that $\det(G) = \det(F_1) \neq 0$.

5. Thus, there is an inverse function

$$(x_1, x_2) = (g_1(y_1, y_2), g_2(y_1, y_2))$$

such that

$$F(g_1(y_1, y_2), g_2(y_1, y_2)), g_2(y_1, y_2)) = (y_1, y_2)$$

for all $(y_1, y_2) \in B(0, \delta) \subset \mathbb{R}^n$ for some sufficiently small $\delta > 0$.

6. But we already knew that $g_2(y_1, y_2) = y_2$ and we get that:

$$F(g_1(y_1, y_2), y_2), y_2) = (y_1, y_2).$$

7. Choosing $y_1 = 0$, we get that

$$F(g_1(0, y_2), y_2), y_2) = (0, y_2)$$

and defining $g(y_2) \equiv g_1(0, y_2)$ we get that

$$F(g(y_2), y_2) = 0$$

for all $y_2 \in B(0, \delta) \subset \mathbb{R}^{n-m}$.

□

The differentiability follows from the Inverse Function Theorem. □

Remark 5.6.1. *The idea behind the implicit function theorem is that, assuming that $D_{x_1}F$ (The first m columns of the derivative of F) is invertible at $a = (a_1, a_2)$, then any changes caused by jiggling a_2 a bit can be undone by changing a_1 a bit.*

In a little more detail:

1. *Since $D_{x_1}F(a_1, a_2)$ is invertible and this derivative is continuous, $D_{x_1}F(a_1 + h_1, a_2 + h_2)$ is invertible for $(h_1, h_2) \in B(0, \delta) \subset \mathbb{R}^n$ for small enough δ.*

2. *This in turn implies that for any fixed $h_2 \in B(0, \frac{\delta}{2}) \subset \mathbb{R}^{n-m}$, $F(a_1 + h_1, a_2 + h_2)$ is an invertible function of $h_1 \in B(0, \frac{\delta}{2}) \subset \mathbb{R}^m$, with $B(F(a), \epsilon) \subset F(a_1 + B(0, \frac{\delta}{2}), a_2 + h_2)$ for all $h_2 \in B(0, \frac{\delta}{2})$ – this is the thing that takes the most to prove carefully, if done directly as suggested here.*

3. *Using the fact that because F is continuous, for small enough $\delta_2 \le \frac{\delta}{2}$ we will get that when $h_2 \in B(0, \delta_2) \Rightarrow |F(a_1, a_2 + h_2)| < \epsilon$.*

4. *Putting all this together, we get that for $h_2 \in B(0, \delta_2)$, there is a function g such that $F(g(h_2), h_2) = F(a_1, a_2) = F(a)$*

Note: To get correspondance between remark and the theorem above, note that set $a_2 + g(x_2 - a_2)$ replaces $g(x_2)$ in the proof of the theorem which replaces $g_a(x_2)$ in the statement of the theorem.

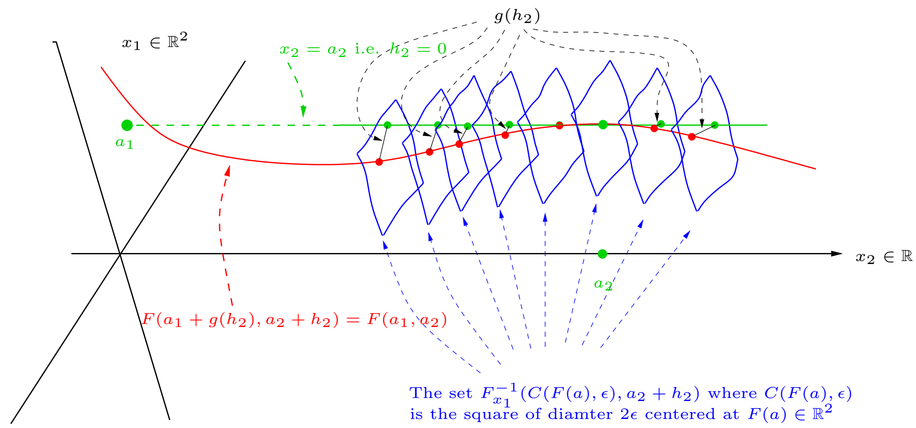

Figure 5.1: Illustration for remark 5.6.1: the F in this illustration maps $\mathbb{R}^3$ to $\mathbb{R}^2$. Continuity of the local inverse in the first factor, x_1, allows us to define an implicit function. Note: We are illustrating the implicit function whose graph lives in $\mathbb{R}^3$, the domain of F, not the graph of F, which lives in $\mathbb{R}^5$. Note: the deformation shown of the square $C(F(a), \epsilon)$ as we move along the x_2 direction is a deformation in the 2 dimensions associated with x_1, not the single dimension associated with x_2). That is those wiggly squares are still flat!

Remark 5.6.2. *Note that if, instead of Steps 1-2 in the remark 5.6.1, we assume that F is continuous,*

$$F(a_1 + h_1, a_2 + h_2) : h_1 \in B(0, \delta) \to U_{h_2} \textit{ is invertible}$$

and

$$B(F(a), \epsilon) \subset U_{h_2} \textit{ for some } \epsilon > 0 \textit{ and all } h_2 \in B(0, \delta),$$

we get a locally valid implicit function.

Thus differentiability is not necessary as long as we know something about the local invertibility of the F in the first factor of the domain (i.e. invertibility in h_1).

Remark 5.6.3. *Though I did notice, in the wikpedia article on the implicit function theorem, that someone had worked out the implicit function theorem for continuous, non-differentiable functions that are invertible in the first factor, I did not go back to see how closely what I wrote corresponds to what they quote there.*

Remark 5.6.4. *I recommend working through Problem 7 in Section 4.6 on page 152 of Fleming's book.*

5.7 Manifolds

I like Fleming's section on manifolds and the fact he takes a very concrete approach to the definition of a manifold. But I will add a few notes that, to a large degree, overlap what he does.

5.7.1 Embedded manifolds are enough

Assuming that the manifold is a subset of $\mathbb{R}^n$ for some n is, technically speaking, not very restrictive since any r-manifold can be embedded in some $\mathbb{R}^n$ for an $n \geq 2r$.

Of course, such an embedding might not be helpful, or might even make things less clear, so in many practical cases, you will work with a more intrinsic definition. Nevertheless, a great deal of intuition can be built and broadly useful results can be constructed considering only submanifolds of $\mathbb{R}^n$. Of course, all local properties and ideas can be understood using only submanifolds.

The basis for the statements that embedded manifolds are (almost) enough is based on the Whitney and Nash emebdding theorems.

5.7.2 Definitions

Note: I now adopt Fleming's choice of which variables to assume invertible. That is, I assume that our defining maps $F(x_1, x_2)$ are locally invertible in the *second factor* x_2 which is n-m dimensional, so that the level sets defined by F are m-dimensional.

Definition 5.7.1 (m-dimensional manifolds)**.** *Suppose that a set $E \subset \mathbb{R}^n$ has the property that for every point $\hat{x} \in E$, there is an open ball $B(\hat{x}, \delta)$ such that $E \cap B(\hat{x}, \delta) = \{x|F(x) = 0\} \cap B(\hat{x}, \delta)$ for some $F : \mathbb{R}^n \to \mathbb{R}^{n-m}$ where DF is full rank in $B(\hat{x}, \delta)$. Then we say that E is an m-***dimensional manifold**. *Because $E \subset \mathbb{R}^n$, we also say that E is an m-***submanifold of** $\mathbb{R}^n$

Definition 5.7.2 (m-Slices of $\mathbb{R}^n$)**.** *Define the m-slice of R^n, R_m, by $R_m \equiv \{x \in \mathbb{R}^n | x_{m+1} = x_{m+2} = ... = x_n = 0\}$.*

Definition 5.7.3 (Slice map for a manifold)**.** *We start with an m-manifold E.*

1. *Given a defining map $F : \mathbb{R}^n \to \mathbb{R}^{n-m}$ for an m-manifold E at a point $\hat{x} = (\hat{x}_1, \hat{x}_2) \in E$.*
2. *Inspired by the proof of the implicit function theorem (but assuming F_{x_2} invertible instead of F_{x_1} invertible), define the (local) C_1 diffeomorphism $G(x_1, x_2)$ to be $(y_1, y_2) = G(x_1, x_2) = (G_1(x_1, x_2), G_2(x_1, x_2)) = (x_1, F(x_1, x_2))$.*
3. *Suppose G is invertible on $B((\hat{x}_1, \hat{x}_2), \delta)$.*
4. *Choose $B((\hat{x}_1, 0), \eta) \subset G(B(\hat{x}, \delta))$.*
5. *Define $H = G^{-1}$. Note that H is a C^1 diffeommorphism.*
6. *Define $B \equiv B((\hat{x}_1, 0), \eta)$*

Then we have that

$$E \cap H(B) = H(R_m \cap B)$$

i.e., we can conclude that G localy straightens out E into a piece of R_m. We will call H the **slice map** *because it locally maps R_m difeomorphically onto E.*

Definition 5.7.4 (Tanget vectors I)**.** *At any point $x \in E \subset \mathbb{R}^n$, where E is an m-manifold, we define the set of* **tangent vectors** *to be vectors $v \in \mathbb{R}^n$ such that there is a differentiable map $\gamma : (-\epsilon, \epsilon) \subset \mathbb{R} \to E$ satisfying (1) $\gamma(0) = x$ and (2) $\dot{\gamma}(0) = v$. ($\dot{\gamma}$ denotes the derivative of γ with respect to t.)*

Since there is locally a one to one corredpondence between differentiable paths in E through x and paths in $\mathbb{R}_m$ through $(x_1, 0)$, we can modify the last definition:

Definition 5.7.5 (Tanget vectors II). *At any point $x \in E \subset \mathbb{R}^n$, where E is an m-manifold, we define the set of* **tangent vectors** *to be vectors $v \in \mathbb{R}^n$ such that there is a differentiable map $\gamma : (-\epsilon, \epsilon) \subset \mathbb{R} \to R_m$ satisfying (1) $H(\gamma(0)) = x$ and (2) $DH(\dot{\gamma}(0)) = v$. ($\dot{\gamma}$ denotes the derivative of γ with respect to t.)*

Fact 5.7.1. *Tangent vectors at $x \in E$ are in the null space of $D_x F$, the derviative of the defining function F at x.*

Proof. Since $G_2(H(\gamma(t))) = 0$ for all t in a neighborhood of 0, we get that $DG_2 \circ DH\dot{\gamma}(0) = 0$. In other words the tangent vectors of E at x are null vectors of $D_x G_2 = D_x F$. □

Definition 5.7.6 (Tangent Space). *The* **tangent Space** *of E at a point x, $T_x E$, is the set of all tangent vectors to E at x.*

Definition 5.7.7 (Normal Vectors). *Suppose that M is a k-dimensional submanifold of $\mathbb{R}^n$. Then at each point $x \in M$, we define the set of* **normal vectors,** *$N_x M$ to be all $w \in \mathbb{R}^n$ such that $\langle v, w \rangle = 0$ for all $v \in T_x M$.*

Remark 5.7.1. *Suppose $\phi : \mathbb{R}^n \to \mathbb{R}^{n-k}$ is a defining function for a k-submanifold M. Then the columns of $(D_x\phi)^t$ (the transpose of $D_x\phi$) span $N_x M$, the normal space of M at x.*

Exercise 5.7.1. Prove the statement in Remark 5.7.1

5.7.3 Intersections

Definition 5.7.8 (Transverse Intersections). *Suppose that K and M are submanifolds of $\mathbb{R}^n$ of dimension k and m and suppose that $x \in K \cap M$. Then we say that K and M* **intersect transversly** *if $\dim(T_x K \cap T_x M) = k + m - n$. When $k + m - n < 0$, we interpret this to mean that $K \cap M = \emptyset$. We denote the fact that K and M intersect transversly by $K \pitchfork M$.*

Remark 5.7.2 (Transverse Linear Subspaces). *Note that if K is a linear subspace of dimension k, then there are $n-k$ vectors $N_K \equiv \{n_1, ..., n_{n-k}\}$ which are mutually orthogonal to each other and to every vector in K. Likewise, if M is a linear subspace of dimension m, then there are $n - m$ vectors $N_M \equiv \{m_1, ..., m_{n-k}\}$ which are mutually orthogonal to each other and to every vector in M. It is clear that every vector in $K \cap M$ will be orthogonal to every vector in $N_K \cup N_M$. If the set*

$$N_K \cup N_M$$

is linearly independent, we get that

$$\dim(K \cap M) = n - (n-k) - (n-m) = k + m - n.$$

Assume $(n-k) + (n-m) \leq n$. Then $N_K \cup N_M$ will be an independent set of vectors whenever (a) the K and M are chosen randomly or equivalently, when the elements of N_K and N_M are chosen randomly and (in the stochastic sense) independently. Thus we see that transverse intersections are what is expected when Linear subspaces intersect – they are typical *in a sense that can be made precise using ideas from probability.*

Note that if we consider submanifolds K and M, with defining equations

$$\{\phi_i\}_1^{n-k} \text{ which define } \Phi : \mathbb{R}^n \to \mathbb{R}^{n-k}$$

and

$$\{\psi_i\}_1^{n-m} \text{ which define } \Phi : \mathbb{R}^n \to \mathbb{R}^{n-k},$$

then, together they Φ and Ψ define a map G, from $\mathbb{R}^n$ to $\mathbb{R}^{2n-k-m}$. Wherever DG is full rank and $2n - k - m = (n - k) + (n - m) \leq n$, we will get that G locally defines a manifold of dimension $n - (2n - k - m) = k + m - n$. Of course,

$$\{x|G(x) = c = (c_1, c_2)\}$$

for some

$$c = (c_1, c_2) \in \mathbb{R}^{n-k} \times \mathbb{R}^{n-m} = \mathbb{R}^{2n-k-m},$$

is the intersection of

$$\{x|\Phi(x) = c_1\} \text{ and } \{x|\Psi(x) = c_2\}.$$

Thus transverse intersection translates into well behaved intersection in these sense of the manifold character of the intersection. A little more precisely, if

1. Φ is full rank at x and therefore locally Φ defines a k-submanifold, and

2. Ψ is full rank at x and therefore locally Ψ defines a k-submanifold

then

$$\{K \pitchfork M\} \Rightarrow \{K \cap M \text{ is locally a } (k+m\text{-}n)\text{-submanifold}\}$$

.

5.7.4 Comments

For further study, I recommend a combination of John M. Lee's three books on manifolds [10, 11, 12], Boothby's book, "Introduction to Differentiable manfolds and Riemannian Geometry" [1] and Do Carmo's book "Riemannian Geometry" [6]. If you are going to pick only one, I would pick Boothby's book, but I would recommend having Lee's book on smooth manifolds on hand as well. For Riemannian Geometry, I like Do Carmo's book best, but even here, the others have their merits.

Chapter 6

Measure and Integration

6.1 Riemann vs Lebesgue

Since you have already been exposed to the idea of integration in the usual calculus course, you will already have an intutively correct idea of what integration is all about. What we change here is the kinds of things we integrate and the measures we integrate over. And we get much more deeply involved in the details, which turn out to be wonderfully rich.

Lebesgue integration is the typical choice of analysts when they want to think about integrating things. But it is not the only choice. Daniell integrals, Steltjes integrals, and a bunch of others are out there, all with their particular uses and enthusiasts. Our approach here is pragmatic: Lebesgue works for most things and for those things we will use it. When it doesn't quite fit the bill, we use what does work.

So, what is Lebesgue integration and how does it differ from Riemann integration? In Riemann integration, we partition the *domain* into regular subsets (intervals or rectangles) and take the the largest and smallest functional values attained in each subset, multiply these values by the measure of those subsets and sum these up, after which we take infimums and supremums:

$$\int^* f d\mu \equiv \inf_P \sum_i \sup_{x \in I_i} f(x)\mu(I_i)$$

$$\int_* f d\mu \equiv \sup_P \sum_i \inf_{x \in I_i} f(x)\mu(I_i)$$

where P is the partition of the domain into intervals I_i. If $\int^* f d\mu = \int_* f d\mu$ then we say f is Riemann integrable.

In Lebesgue integration, we parition the *range* into intervals I_i and pull them back to a partition of the domain: $E_i = f^{-1}(I_i)$. (This paritition can be very far from regular!) We get:

$$\int^* f d\mu \equiv \inf_P \sum_i \left(\sup_{y \in I_i} y \right) \mu(f^{-1}(I_i)) = \inf_P \sum_i b_i \mu(f^{-1}(I_i))$$

$$\int_* f d\mu \equiv \sup_P \sum_i \left(\inf_{y \in I_i} y \right) \mu(f^{-1}(I_i)) = \sup_P \sum_i a_i \mu(f^{-1}(I_i))$$

We are rewarded for our change in pespective by the result that now, *every* respectable function is integrable! (By integrable we will mean the upper and lower integrals are equal). As a result, we like the Lebesgue integral and are not so inclined to like the Riemann integral, even though for many practical purposes they are indistinguishable (because for really nice functions, they are the same.) Figure 6.1 illustrates both versions of integration.

Exercise 6.1.1. Recall that the function $f : \mathbb{R}^n \to \mathbb{R}$ is said to be *measurable* if the set $f^{-1}(I)$ is $\mathcal{L}^n$ measurable whenever I is a (possibly infinite) interval. Suppose that the support of f is bounded. Show that the Riemann integral $\int f(x)\, dx$ exists when f is continuous, but that it is even easier to show that the Lebesgue integral $\int f(x)\, d\mathcal{L}^1 x$ exists when f is merely measurable. Hint: A continuous function on a compact set is uniformly continuous.

6.2 Iterated Integrals

Integraing [x+y] over $[0,r] \times [0,s]$**:** The problem of integrating $\phi(x,y) = [x+y]$ over the region $A(r,s) \equiv \{(x,y) \mid 0 \le x \le r \text{ and } 0 \le y \le s\}$ can be done by brute force (which is what I started to do when working on it, but realized later I could do this by rearranging the terms a bit). We can use the symmetry of the region to simplify the calculation a great deal.

1. Note that ϕ is constant on any line with a slope of -1.
2. Let R_L and R_U are upper and lower regions with the same area, $V_2(R_L) = V_2(R_U)$,
3. Let $\phi(R_L)$ and $\phi(R_U)$ be the values of ϕ on those regions, noting that ϕ is constant on those regions.
4. Then:

$$\begin{aligned}
\phi(R_L)V_2(R_L) + \phi(R_U)V_2(R_U) &= (\phi(R_L) + \phi(R_U))\, A \\
&= \frac{\phi(R_L) + \phi(R_U)}{2} 2A \\
&= \left(\frac{\phi(R_L) + \phi(R_U)}{2}\right) (V_2(R_L) + V_2(R_U)) \\
&= \left(\frac{r+s-1}{2}\right) (V_2(R_L) + V_2(R_U))
\end{aligned}$$

5. If the number of regions is even, we are done.
6. If the number of regions (steps of ϕ in $[0,r] \times [0,s]$ with non-zero area) is odd the value of the middle step is $\frac{r+s-1}{2}$. **Proof:** since there are $r+s$ steps, the middle step, when $r+s$ is odd is $\frac{r+s-1}{2} + 1$ step. but the value on the kth step is always $k-1$, so the value of ϕ on the middle step is $\frac{r+s-1}{2} + 1 - 1 = \frac{r+s-1}{2}$.
7. See Figure 6.2

Exercise 6.2.1. Suppose that $\phi(x,y) = [x+y]$ and $A = [0,r] \times [0,s] \subset \mathbb{R}^2$. State why

$$\int_A \phi(x,y)\, d\mathcal{L}^2 = \int_A (r+s-1) - \phi(x,y)\, d\mathcal{L}^2.$$

Use this to compute $\int_A \phi(x,y)\, d\mathcal{L}^2$.

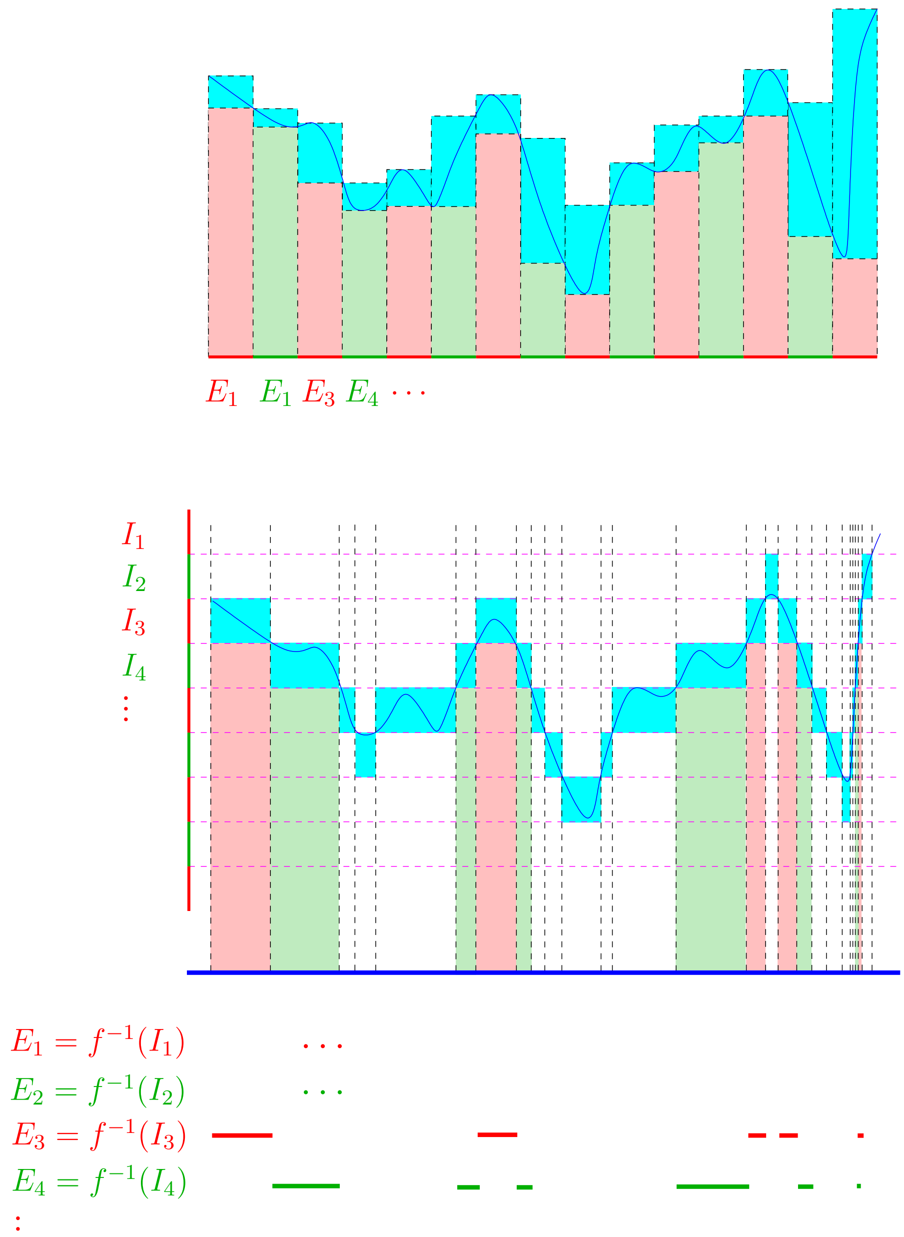

Figure 6.1: **Riemann versus Lebesgue Integration:** the upper figure illustates the partition of the domain dictated by the Riemannian approach. The green and red rectangles live completely below the graph of f. Call the area they sum to $A_{\text{lower}}(P)$ where P is the paritition. The red and green plus the cyan rectangles live completely above the graph. Call their area A_{upper}. If $\sup_P A_{\text{lower}}(P) = \inf_P A_{\text{upper}}$ then f is Riemann integrable. The lower figure illustrates that key difference for the Lebesgue case: we partition the range and pull that back by f^{-1} to a partition of the domain. It turns out that this is exactly what is needed to make all reasonable functions integrable. Now $A_{\text{lower}}(P) = \sum_i a_i \mu(E_i)$ and $A_{\text{upper}}(P) = \sum_i b_i \mu(E_i)$ where P is a parition of the range into the intervals $I_i = [a_i, b_i)$.

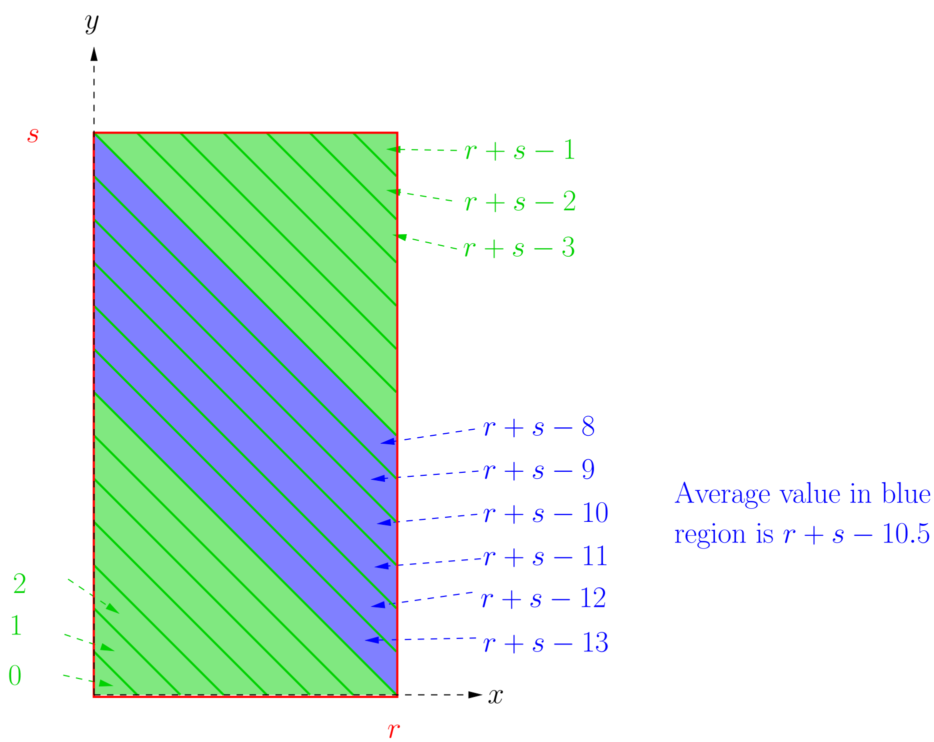

Figure 6.2: **Integrating** $[x+y]$ **over** $[0,r]\times[0,s]$**:** The function ϕ is a step function and in $[0,r]\times[0,s]$ there are $r+s$ steps with positive area, and one – $\{(r,s)\}$ – that has one point in it. In the case pictured here, $r+s=20$. If we combine regions of equal size, starting at regions touching the opposite corners $(0,0)$ and (r,s), the average of the ϕ on those regions is always $\frac{r+s-1}{2}$. Thus the integal is $\frac{rs(r+s-1)}{2}$. I.e since $\frac{0+r+s-1}{2} = \frac{1+r+s-2}{2} = \frac{2+r+s-3}{2} = \cdots$ the integral is equivalent to integrating $\frac{r+s-1}{2}$ over the rectangle. Note that in this case, $r=7$ and $s=13$ so that the $\frac{r+s-1}{2} = r+s-10.5$

Exercise 6.2.2. Suppose $f(x,y) \equiv y$ and we define $A \equiv \{(x,y) \mid y \geq x,\ 0 \leq x \leq 1,\ y \leq 1\}$. Evaluate $\int_A f(x,y)d\mathcal{L}^2$ in two different ways by using iterated integrals. Draw figures illustrating what you are doing.

Exercise 6.2.3. Suppose $f(x,y,z) \equiv 1$ and we define $A \equiv \{(x,y,z) \mid y \geq x,\ 0 \leq x \leq 1,\ y \leq 1, 0 \leq z \leq y-x\}$. Evaluate $\int_A f(x,y,z)d\mathcal{L}^3$ in three different ways by using iterated integrals. Draw figures illustrating what you are doing.

Exercise 6.2.4. Suppose $f(x,y,z) \equiv z$ and we define A to be the region in Exercise 6.2.3. Evaluate $\int_A f(x,y,z)d\mathcal{L}^3$ in three different ways by using iterated integrals. Draw figures illustrating what you are doing.

There are difficult problems that can be solved by switching order of integration or summation. Examples include the proof of the Cauchy Binet formula (a very significant generalization of the pythagorean theorem), the proof of the central result in Vapnik-Chernovenkis theory (a very important piece of statistical learning theory), and the proof of the deformation theorem in geometric measure theory (which is a very powerful approximation theorem for currents which are generalized surfaces). A reference for the first theorem is Evans and Gariepy's *Measure Theory and Fine properties of functions* [7], references for the second include the very nice notes on statistical learning theory by Rob Nowak nowak-2009-notes as well as *A probabilistic Theory of Pattern Recognition* by Devroye, Györfi and Lugosi [5], and a reference for the last theorem is Krantz and Parks', *Geometric Integration Theory* [9]. These are all more advanced than Fleming's book, but a little bit of coaching would be sufficient for the motivated student to dig into any one of those. (But the coaching would be important in most cases.)

6.3 When Integrals Diverge.

In Fleming's book, there is a problem that asks you to show that

$$\lim_{r\to\infty} \int_1^r \frac{\sin(x)}{x} dx$$

exists and is finite. This is the launching point for the exploration in this section. Recall that $f^+ \equiv \max(f,0)$, $f^- \equiv \max(-f,0)$, and that

$$\chi_E(y) \equiv \begin{cases} 1 \text{ when } y \in E \\ 0 \text{ when } y \notin E \end{cases} \tag{6.1}$$

I will define

$$\text{sgn}(y) \equiv \begin{cases} 1 \text{ when } y > 0 \\ 0 \text{ when } y \leq 0 \end{cases} \tag{6.2}$$

(Note that this definition of sgn is non-standard.) We define $\mathbb{R}^+ \equiv \{x \in \mathbb{R} \mid x > 0\}$ and $\mathbb{Z}^+ \equiv \{x \in \mathbb{Z} \mid x > 0\}$. Note also that we will often use dx to denote $d\mathcal{L}^k$, k-dimensional Lebesgue measure in $\mathbb{R}^k$.

Exercise 6.3.1. Note that $\lim_{r\to\infty} \int_1^r \frac{1}{x}dx \to \infty$. The following sequence is inspired by (and aimed at) the following problem: show that

$$\lim_{r\to\infty} \int_1^r \left(\frac{\sin^+(x)}{x}\right) dx = \infty.$$

1. Show that if $a_i \geq 0$ for all positive integers i and a_i are non-increasing: $a_1 \geq a_2 \geq \cdots$, then, for any $k \in \mathbb{Z}^+$, $\sum_{i=1}^{\infty} a_i$ is finite if and only if $\sum_{i=1}^{\infty} a_{1+k\cdot i}$ is finite.
2. Define $S^+(x) = \text{sgn}(\sin(x))$: the function that is 1, when $\sin(x) > 0$, and 0, when $\sin(x) \leq 0$.
3. Show that $\lim_{r\to\infty} \int_1^r \frac{1}{x} dx \to \infty$ implies that $\lim_{r\to\infty} \int_1^r \frac{S^+(x)}{x} dx \to \infty$.
4. For any positive integer n, define:
   (a)
   $$\begin{aligned} E_n &\equiv \cdots [-2n, -2n+1] \cup [-n, -n+1] \cup [0,1] \cup [n, n+1] \cup [2n, 2n+1] \cup \cdots \\ &= \bigcup_{k\in\mathbb{Z}} [kn, kn+1] \end{aligned}$$
   (b) $B_n(x) \equiv \chi_{E_n}$, and
   (c) $B_{n,\tau,s}(x) = B_n(\tau x + s)$ where τ and s are real numbers.
5. Show that $\lim_{r\to\infty} \int_1^r \frac{\beta B_{n,\tau,s}(x)}{x} dx \to \infty$ for any choice of $n \in \mathbb{Z}^+$ and $\tau, s, \beta \in \mathbb{R}^+$.
6. Find $n \in \mathbb{Z}^+$ and $\tau, s, \beta \in \mathbb{R}^+$ such that $\beta B_{n,\tau,s}(x) \leq \sin^+(x)$ for all x.
7. Now use the previous steps to get that $\lim_{r\to\infty} \int_1^r \left(\frac{\sin^+(x)}{x}\right) dx = \infty$.

Exercise 6.3.2. We depended on the fact that $f(x) = \frac{1}{x}$ was nonincreasing in Exercise 6.3.1, but that was more than we actually needed.

1. Suppose that $0 < f(1) < \infty$ and $f(x) \geq 0$ for all $x \in \mathbb{R}^+$.
2. Suppose $\frac{f(x)}{f(y)} \leq C < \infty$ for all $0 < y < x$.
3. Show that $\int_1^\infty f(x) \sin^+(x) dx$ is finite if and only if $\int_1^\infty f(x) dx$ is finite

Exercise 6.3.3. Show that if $k \in \mathbb{Z}^+$, $a_i \geq 0$ for all $i \in \mathbb{Z}^+$ and $\frac{a_i}{a_j} \leq C < \infty$ for all $j < i$, then $\sum_{i=1}^{\infty} a_i$ is finite if and only if $\sum_{i=1}^{\infty} a_{1+k\cdot i}$ is finite.

Exercise 6.3.4. Here is another variation:

1. Suppose that $0 < f(1) < \infty$ and $f(x) \geq 0$ for all $x \in \mathbb{R}^+$.
2. Suppose $\frac{f(x)}{f(y)} \leq C < \infty$ for all $0 < y < x$.
3. Suppose that $N, k \in \mathbb{Z}^+$.
4. Show that $\sum_{i=N}^{\infty} f(1 + k \cdot i)$ is finite if and only if $\int_1^\infty f(x) dx$ is finite

Exercise 6.3.5. And another:

1. Suppose again that $0 < f(1) < \infty$, $f(x) \geq 0$ for all $x \in \mathbb{R}^+$ and $\frac{f(x)}{f(y)} \leq C < \infty$ for all $0 < y < x$.

2. Suppose that $\{x_i\}_{i=1}^{\infty}$ has the property that $0 < \alpha < \frac{i}{x_i} < \beta < \infty$.

3. Show that $\sum_{i=1}^{\infty} f(x_i)$ is finite if and only if $\int_1^{\infty} f(x)dx$ is finite.

4. Show that if we only know that $0 < \alpha < \frac{i}{x_i}$, then, if $\sum_{i=1}^{\infty} f(x_i)$ is finite, we know that $\int_1^{\infty} f(x)dx$ is finite. Find an example of a sequence such that $0 < \alpha < \frac{i}{x_i}$, $\int_1^{\infty} f(x)dx$ is finite, but $\sum_{i=1}^{\infty} f(x_i)$.

Exercise 6.3.6. **[Challenge Problem]** Suppose again that $0 < f(1) < \infty$, $f(x) \geq 0$ for all $x \in \mathbb{R}^+$ and $\frac{f(x)}{f(y)} \leq C < \infty$ for all $0 < y < x$.

1. Assume the measure μ has the property that $0 < \alpha \leq \frac{\mu([n,n+m])}{m} \leq \beta < \infty$ for $n, m \in \mathbb{Z}^+$ such that $n, m \geq N$, where $N \in \mathbb{Z}^+$.

2. Show that $\int_1^{\infty} f(x)\, d\mu$ is finite if and only if $\int_1^{\infty} f(x)\, dx$ is finite.

3. Show that Exercises (6.3.1-6.3.5) are special cases of this result.

The key idea in Exercises (6.3.1-6.3.6) is that to know whether or not $\int_1^{\infty} f dx < \infty$, you need only have a very rough approximation to the measure in the integral (i.e. the Lebesgue measure $\mathcal{L}^1$) and the property that $f(y)$ bounds $f(x)$ for $y < x$, modulo the constant C.

6.4 Transforming Volumes

If we map $E \subset \mathbb{R}^n$, $\mathcal{L}^n(E) < \infty$, by an affine map $F : x \to Ax + b$, where $F : \mathbb{R}^n \to \mathbb{R}^n$, then $\mathcal{L}^n(F(E)) = |\det(A)|\mathcal{L}^n(E)$. How can we see this?

First we observe that if

$$A = \begin{bmatrix} | & | & \dots & | \\ a_1 & a_2 & \dots & a_n \\ | & | & \dots & | \end{bmatrix}$$

and

$$e_i = \left.\begin{bmatrix} 0 \\ \vdots \\ 0 \\ 1 \\ 0 \\ \vdots \\ 0 \end{bmatrix}\right\} \text{all entries except the } i\text{th is zero, the } i\text{th} = 1,$$

then $a_i = Ae_i$ and we have that the unit cube, defined by the e_i's gets mapped to the parallelpiped defined by the a_i.

Lebesgue measure is rotationally invariant and so we can compute the volume of the parallelpiped defined by the columns of A or we can compute the volume of OA for any orthogonal matrix O. The QR decomposition gives us $A = QR$ where Q is an n by n orthogonal matrix and

R is an upper triangular matrix. (You might know this as the gram-schmidt orthogonalization of the columns of A.) Note that

$$A = \begin{bmatrix} | & | & \dots & | \\ a_1 & a_2 & \dots & a_n \\ | & | & \dots & | \end{bmatrix} = \begin{bmatrix} | & | & \dots & | \\ q_1 & q_2 & \dots & q_n \\ | & | & \dots & | \end{bmatrix} \begin{bmatrix} r_{1,1} & r_{1,2} & r_{1,3} & \dots & r_{1,n} \\ 0 & r_{2,2} & r_{2,3} & \dots & r_{2,n} \\ 0 & 0 & r_{3,3} & \dots & r_{3,n} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \dots & r_{n,n} \end{bmatrix}$$

Since the volume of the parallelpiped defined by the columns of A is the same as the volume of the parallelpiped defined by the columns of $Q^t A = R$.

Now we change the signs of the columns so that all the $r_{i,i}$'s are positive. This only changes the determinant by a sign, if at all. We ignore the sign since a negative sign just indicates a reflection, i.e. a change in orientation and we are here only focusing on dilations and contractions here. We denote this changed R by $\hat{R}$ and the elements of are $\hat{r}_{i,j}$, instead of $r_{i,j}$.

Now we note that the n-parallelpiped defined by the columns of $\hat{R}$, $P_n^{\hat{R}}$, equals the object obtained by taking $P_{n-1}^{\hat{R}}$, the $n-1$-parallelpiped defined by the first $n-1$ columns of A, and stacking them up with shifts $(\hat{r}_{1,n}, \hat{r}_{2,n}, \dots, \hat{r}_{n-1,n}, 0)$ direction. More precisely, the slice of $P_n^{\hat{R}}$ at a height h above the plane of the first $n-1$ coordinates equals the set

$$P_{n-1} + \frac{h}{\hat{r}_{n,n}}(\hat{r}_{1,n}, \hat{r}_{2,n}, \dots, \hat{r}_{n-1,n}, \hat{r}_{n,n}).$$

Fubini's theorem give us that the volume of $P_n^{\hat{R}}$ must therefore be the volume of $P_{n-1}^{\hat{R}}$ times the height of the stack, $\hat{r}_{n,n}$. That is,

$$\mathcal{L}^n(P_n^{\hat{R}}) = \mathcal{L}^{n-1}(P_{n-1}^{\hat{R}})\hat{r}_{n,n}.$$

Continuing in this way, we get that

$$\begin{aligned} \mathcal{L}^n(P_n^{\hat{R}}) &= \Pi_{i=1}^n \hat{r}_{i,i} && (6.3) \\ &= \det(\hat{R}) && (6.4) \\ &= |\det(R)| && (6.5) \\ &= |\det(Q^t)\det(A)| && (6.6) \\ &= |\det(A)|. && (6.7) \end{aligned}$$

Thus we get that A changes volumes of sets in $\mathbb{R}^n$ by a factor of $|\det(A)|$

6.5 Regions Bounded by Simple Closed Curves

If $\omega(t)$ is a smooth, arc-length parameterization of $\partial\Omega \subset \mathbb{R}^2$ and is a simple closed curve with length L, then

$$\mathcal{L}^2(\Omega) = \int_0^L \omega \times \dot{\omega}\, dt$$

We are assuming the standard counterclockwise orientation of the boundary.

There are two ways in which we will prove this.

Proof 1. This is a direct proof, using an approximation argument:

1. We can approximate ω as closely as we like, both in terms of length and the area of the enclosed approximation to Ω, with a polygon whose vertices lie on ω. We will denote the region bounded by P, Ω_P.

2. We can create this polygon by choosing a set of rays from the origin that are spaced finely enough that the P and Ω_P appriximate ω and Ω as closely as we like. See Figure 6.3.

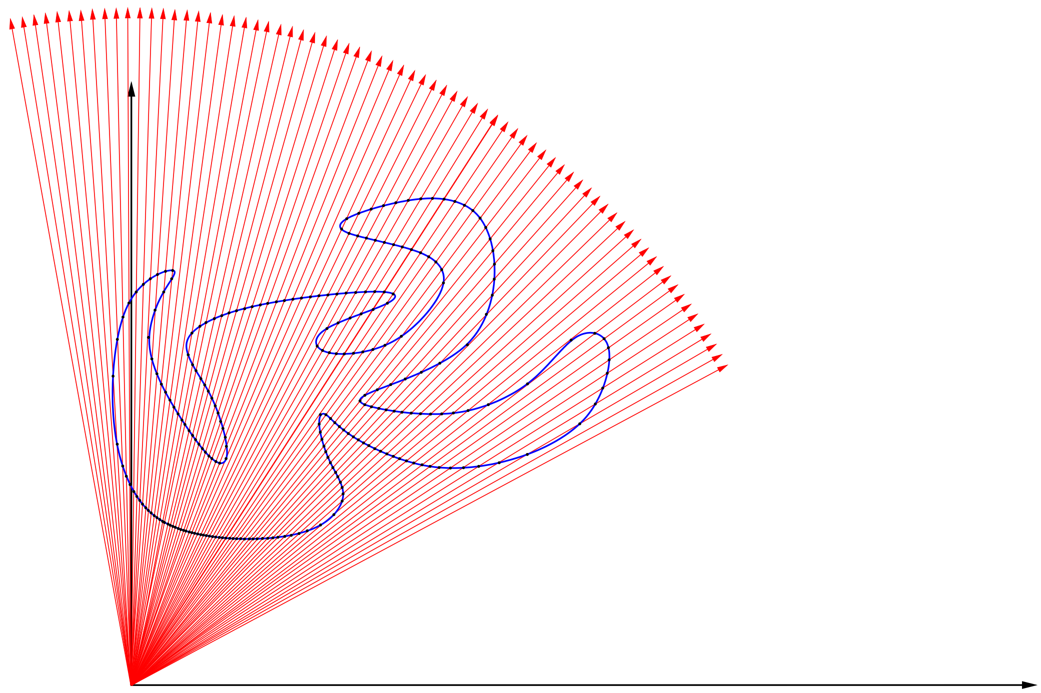

Figure 6.3: Rays that generate a polygon P, approximating Ω.

3. We will denote the vertices by $\{v_i\}_{i=1}^N$. The side that starts at v_i and ends at v_{i+1} will be denoted by S_i. The orientation of the polygon P is also counterclockwise.

4. The oriented area of the triangle formed by any side S_i of P and the ray from the origin to v_i is given by $\frac{v_i \times S_i}{2}$.

5. You can convince yourself that the area of Ω_P is given the sum of the oriented areas, $\sum_{i=1}^{N-1} \frac{v_i \times S_i}{2}$.

6. Now, using Taylor series approximations for $\omega(t_i) = v_i$, we can get the fact that S_i is equal to $\dot{\omega}(t_i)\Delta t_i + \ddot{\omega}(\alpha_i t_i)(\Delta t_i)^2$ for some $|\alpha_i| < 1$.

7. Suppose that ω is contained in $B(0, R)$, that $\Delta t_i \leq \Delta$ for all i and $|\ddot{\omega}(t)| < C$ for all t. We

obtain that:

$$\begin{aligned}
\sum_{i=1}^{N-1} \frac{v_i \times S_i}{2} &= \sum_{i=1}^{N-1} \frac{v_i \times \left(\dot\omega(t_i)\Delta t_i + \ddot\omega(\alpha_i t_i)(\Delta t_i)^2\right)}{2} \\
&= \sum_{i=1}^{N-1} \frac{v_i \times (\dot\omega(t_i)\Delta t_i)}{2} + \underbrace{\sum_{i=1}^{N-1} \frac{v_i \times \left(\ddot\omega(\alpha_i t_i)(\Delta t_i)^2\right)}{2}}_{\leq C(N-1)\Delta} \\
&\underset{\Delta\to 0}{\to} \int_0^L \frac{\omega \times \dot\omega}{2}\, dt
\end{aligned}$$

□

Another, simpler proof uses the divergence theorem and it applies to non-smooth $\partial\Omega$ as well.

Proof 2. If Ω is an open, connected set of finite perimeter, and we parameterize the boundary to get ω, we also have a divergence theorem (see Evans and Gariepy [7], chapter 5). The reduced boundary $\partial^*\Omega$ of sets of finite perimeter, which is the boundary that one can "see" when you integrate, is the boundary we pay attention to, not the topological boundary.

1. Define the radial vectorfield $v(x) = x$ and notice that $\nabla \cdot v(x) = 2$.
2. We will have that the oriented unit normal exists $\mathcal{H}^1$ almost every $t \in \partial^*\Omega$ the unit normal $\vec{n}(t)$. Then the orienting unit tanget at those t will be $\vec{n}(t)$ rotated $\frac{\pi}{2}$ counterclockwise and we will denote it by $\dot\omega(t)$.
3. We get that:

$$\begin{aligned}
\int_\Omega \nabla \cdot v(x) dx &= \int_\omega v(\omega(t)) \cdot \vec{n}(t)\, dt \\
&= \int_\omega v(\omega(t)) \times \dot\omega(t)\, dt \\
&= \int_\omega \omega(t) \times \dot\omega(t)\, dt \quad (\text{since } v(\omega(t)) = \omega(t))
\end{aligned}$$

4. using the fact that $\nabla \cdot v(x) = 2$ everywhere, we get

$$2\mathcal{H}^2(\Omega) = \int_\omega \omega(t) \times \dot\omega(t)\, dt$$

□

Remark 6.5.1. *Actually this volume calculation works even in higher dimensions, but we need to use k-vectors and wedge products to state it. If $v(x)$ is the radial vectorfield and η is the orienting, simple unit $n-1$-vectorfield for $\partial^*\Omega$, the reduced boundary of any set of finite perimeter $\Omega \subset \mathbb{R}^n$, we get:*

$$n\mathcal{H}^n(\Omega) = \int_{\partial^*\Omega} v(s) \wedge \eta\, d\mathcal{H}^{n-1}s.$$

Chapter 7 of Fleming's book introduces the ideas needed here. See also Frank Morgan's book [13] for more about k-vectors and and wedge products.

Remark 6.5.2. *A good start at understanding what a simple k-vector is can be obtained by thinking of it as the k-dimensional parallelpiped defined by k independent vectors. So, in our case, we take* $n-1$ *indepepdent vectors in the tangent space of* $\partial^*\Omega$ *and "wedge" them together to get something that is sort of like a parallelpiped. The wedge product* $v \times \eta$ *is then the n-dimensional parallelpiped defined by* v *and* η*. (Switching the order of the product, switches the orientation of the product.) What I have said is not completely correct, since* η *is unchanged by rotations that keep the plane that* η *defines invariant. So, if you think of* η *as the equivalence class of all parallelpipeds that span the same* $n-1$ *dimensional subspace and have the same volume, you are not very close to one accurate picture of the* $n-1$*-vector* η*.*

6.6 The Area Formula

Recall the definition of Lipschitz:

Definition 6.6.1 (Lipschitz)**.** *A function* $F : \mathbb{R}^n \to \mathbb{R}^m$ *is said to be Lipschitz if there is a constant* $L < \infty$ *such that* $|F(x) - F(y)| \le L|x - y|$ *for all* $x, y \in \mathbb{R}^n$*. The optimal (smallest) constant L for which the inequality holds is called the Lipschitz constant.*

Understanding how measures of sets and integrals of functions over those sets transform under mappings is the point of the **area formula** (and also of the coarea formula, which we will touch on at the end of he section). The formula is a very general formula that works even for maps which are merely Lipschitz.

We also need a generalization of the $|\det(DF)|$:

Definition 6.6.2 (Jacobian)**.** *The Jacobian of a mapping* $F : \mathbb{R}^n \to \mathbb{R}^m$*,* JF *is defined to be*

$$JF = \begin{cases} \sqrt{\det(DF^t \circ DF)} & \text{when } m \ge n \\ \sqrt{\det(DF \circ DF^t)} & \text{when } m < n \end{cases} .$$

Notice that the definition agrees with $|\det(DF)|$ *when* $n = m$*.*

We will need the notion of **Hausdorff measure**, $\mathcal{H}^k$, as well. Referring the reader to [7, 13] for details, it will suffice to note that Hausdorff measure is an outer measure that does exactly what you think it should do for k-dimensional sets in $\mathbb{R}^n$. Zero dimensional Hausdorff Measure is the counting measure: $\mathcal{H}^0 E =$ the number of point in E. Here is the formal definition that takes a bit of time to sink in:

Definition 6.6.3 (Hausdorff Measure)**.** *Suppose that* $\nu \in [0, \infty)$*. Letting* $\mathcal{F}$ *denote any countable collection of subsets of* $\mathbb{R}^n$ *such that* $E \subset \bigcup_{F_i \in \mathcal{F}} F_i$ *and* $D(\mathcal{F})$ *be the supremum of the diameters of the sets* $F_i \in \mathcal{F}$*, we define:*

$$\mathcal{H}^\nu(E) \equiv \lim_{\delta \to 0} \left(\inf_{\{\mathcal{F} | D(\mathcal{F}) \le \delta\}} \sum_{F_i \in \mathcal{F}} \alpha(\nu) \left(\frac{\text{diam}(F_i)}{2} \right)^\nu \right),$$

where $\alpha(\nu)$ *is the volume of the unit ball in* ν *dimensions for* $\nu \in \mathbb{Z}^+$ *and is extended to* $\nu \notin \mathbb{Z}^+$ *using the* Γ *function.*

Remark 6.6.1. *Again, I want to emphasize that this yields a measure that does what you think it should do on smooth, k-dimensional subsets of* $\mathbb{R}^n$*. The measure is defined for any* $\nu \in [0, \infty)$ *and is indispensible in geometric analysis and geometric measure theory.*

If F is Lipschitz, the Rademacher's theorem (see Evans and Gariepy [7], chapter 6) tells us that F is differentiable almost everywhere and the following theorem is valid for any mapping from $\mathbb{R}^n$ to $\mathbb{R}^m$, $n \leq m$ that is at least Lipschitz.

Theorem 6.6.1 (Area formula I)**.** *Suppose that F mapping $\mathbb{R}^n$ to $\mathbb{R}^m$ is Lipschitz and $n \leq m$. Then*

$$\int_E JF \, d\mathcal{L}^n = \int_{F(E)} \mathcal{H}^0(F^{-1}(y)) \, d\mathcal{H}^n y$$

The key idea is that F is linearly approximable at almost every point in E, so we can integate the dilation factor JF over E to get how much the set has contracted or dilated after mapping. BUt we have to take into account multiplicity. That is it can be that several pieces map on top of each other, so the image appears smaller than it should. That is the reason we count the number of preimages of any point in the image when integrating over $F(E)$. See Figure 6.4.

There is a similar formula that transforms the integral of a function over E to an integral over F(E) in the range, again accounting for multiplicities.

Theorem 6.6.2 (Area formula II)**.** *Suppose that F mapping $\mathbb{R}^n$ to $\mathbb{R}^m$ is Lipschitz, $n \leq m$ and g is integrable function from $\mathbb{R}^n$ to $\mathbb{R}$. Then*

$$\int_E g \; JF \, d\mathcal{L}^n = \int_{F(E)} \sum_{x \in F^{-1}(y)} g(x) \, d\mathcal{H}^n y$$

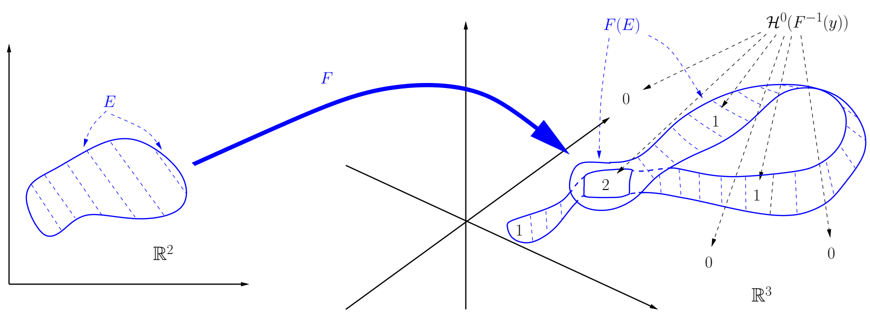

Figure 6.4: Area formula example

The coarea formula, introduced by Herbert Federer in 1959, in a paper titled "Curvature Measures" [8] is a powerful generalization of Fubini's theorem. In Fubini's Theorem, we are allowed to slice a domain up by x_i = constant sets and turn a multidimensional integral into a sequence of one dimensional integrals. Using the Coarea formula we can slice the domain up using level sets of Lipschitz functions. Here is the formula:

Appendix

Linear Algebra

Here is the brief outline of topics whose mastery will enable you to have a solid, working grasp of linear spaces and subspaces, linear transformations, and the properties of matrices that represent transformations and subspaces:

Vector spaces

1. **Linear Independence and Vector Space bases.** Linear Independence of $\{v_i\}_{i=1}^n$ means that linear combinations of the v_i's, $\sum_{i=1}^n \alpha_i v_i$ are not zero unless $\alpha_i = 0$'s for all i. A basis $H = \{h_i\}_{i=1}^n$ is a set of linear independent vectors such that every vector in $v \in V$ can be written as a linear combination of elements of H: $v = \sum_{i=1}^n \beta_i h_i$.
2. **Subspaces and Affine subspaces.** Subspaces include 0, affine subspaces need not include 0. (Therefore a subspace is an affine subspace, but not vice versa.)
3. **Examples.** A rich diversity: $\mathbb{R}^n$, spaces of polynomials, sequence spaces, other function spaces.

Norms

1. **In $\mathbb{R}^n$.** Norms map vector to the non-negative real numbers: $||\cdot|| : V \to \mathbb{R}^+ \cup \{0\}$, satisfying $||\alpha v|| = |\alpha|||v||$, $||v+w|| \le ||v|| + ||w||$. Important examples: 1-norm, 2-norm, ∞-norm, p-norm
2. **In functions spaces.** We have the same important examples in function spaces: 1-norm, 2-norm, ∞-norm, p-norm

Linear operators; affine operators

1. **Operator norm.** The definition: $\sup_{x\neq 0} \frac{|L(x)|}{|x|}$
2. **Reduced echelon form and what is tells you about a matrix.** You can directly read off *what the null space is* and therefore the *dimension of the null space*, which also gives you the *dimension of the range*. The reduced echelon form also tells you what columns can be used to span the range of A, so we can read off a *parameterization of the range of A*.
3. **Null space and level sets.** Suppose that N_A is the null space of A. If $A : \mathbb{R}^n \to \mathbb{R}^m$ and $y \in \mathbb{R}^m$, then if $x \in \mathbb{R}^n$ satisfies $A(x) = y$, then $L_y = x + N_A$ is the set of all points in $\mathbb{R}^n$ that map to y.

4. **Span of a set of vectors.** We denote the set of all linear combinations of the columns of A, i.e. the span of the columns of A, by $\text{span}(A)$.

5. **Determinants.** Let $E \subset \mathbb{R}^n$ and $F_A : \mathbb{R}^n \to \mathbb{R}^n$ be a linear map represented by a matrix A. The determinant of a matrix is the volume dilation factor: $vol(F_A(E)) = |det(A)|vol(E)$. The sign of $det(A)$ tells you if the orientation of $F_A(E)$ has switched or has stayed the same as the orientation of E.

Inner products and Orthogonality

1. **Orthogonality.** When we have an inner product $\langle x, y\rangle$, we say x is orthogonal to y when $\langle x, y\rangle = 0$. This is sometimes denoted $x \perp y$. If all the columns of O are orthogonal to each other and they each have norm equal to 1, we say that O is an orthogonal matrix and the columns are orthonormal. Then $OO^t = O^tO = I$, the $n \times n$ identity matrix with 1's down the diagonal and zeros everywhere else.

2. **Projections.** let P be a matrix with m orthonormal n-dimensional columns. Let $P^\perp$ be the matrix of $n-m$ orthonormal columns each of which is orthogonal to the columns of P. Then $M_P = PP^t$ is the operator which projects $\mathbb{R}^n$ onto the span of the columns of P and if x is in $\text{span}(P)$ then $M_P(x) = x$, otherwise, we can decompose $x = x_P + x_{P^\perp}$ where $x_P = M_P(x)$ and $x_{P^\perp} = M_{P^\perp}$. This decomposition into an element in P and an element in $P^\perp$ is unique.

3. **Nilpotent operators.** N is nilpotent if $N^p = 0$ for some $p > 1$.

4. **QR decomposition.** Relation to Gram Schmidt orthogonalization: they are basically the same thing. Suppose A is a matrix of m, n-dimensional vectors. Then $A = QR$, where R is upper triangular and Q has orthonormal columns. Thus, $\text{span}(A) = \text{span}(R)$.

5. **Convex functions and supporting hyperplanes.** Convex functions are to optimization what linear systems of equations are to differential equations: the "easy" case (which is not so easy all the time). A closed convex subset $E \subset \mathbb{R}^n$ equals the intersection of closed half spaces containing E. If $x \in E^c$ and E is convex, then there exists a $v \in \mathbb{R}^n$ such that $\langle y - x, v\rangle < 0$ for all $y \in E$. If E is closed and $x \in E^c$, there there is a closest point in E x^*, such that $operatornamedis(x, E) = ||x^* - x|| > 0$. We have that if we define $v = x - x^*$, $\langle y - x^*, v\rangle \leq 0$ for all $y \in E$.

Symmetric operators; normal operators

1. **Eigenvectors and eigenvalues.** $(A - \lambda I)v = 0 \to v$ is an eigenvector corresponding to the eigenvalue λ. Eigenvectors can be complex numbers.

2. **Diagonalization.** A is diagonalizable if there is an invertible matrix Q such that $A = QDQ^{-1}$ where D is a diagonal matrix. Some matrices are diagonalizable if and only if we allow Q and D to be complex. If $A = A^t$, where A^t is the transpose of A, then Q can be taken to be an orthogonal matrix: $Q^{-1} = Q^t$

3. **Jordan Normal Form** Generalization of diagonalization that works or all square matrices: allowing complex values, we are able to get that any square matrix A can be decomposed – $A = Q^{-1}JQ$ where the J is an upper triangular matrix with the eigenvalues of A appearing on the diagonal of J.

4. **Relation to operator norm.** The Jordan normal form tells us that the determinant of A equals the product of the eigenvalues of A.

Singular value decomposition (SVD)

1. **All matrices have an SVD.** It is not necessarily unique, but non-uniqueness harmless
2. **Relation to operator norm.** If $||A||$ denotes the operator norm of A, then $||A|| = \sup_i \sigma_i$.
3. **Relation to the determinant.** $\Pi_{i=1}^{n} \sigma_i = |\det(A)|$ – when determinant is defined. When A not square, $\Pi_{i=1}^{n}$ is the correct Extension of the determinant since it measure the expansion or contraction of the subspace normal to the null space of A.
4. **How it illuminates the geometry. Either:**
{rotation/reflection → orthogonal projection → dilation along coordinate axes → rotation/reflection}
Or:
{rotation/reflection → dilation along coordinate axes → embedding in higher dimensional space → rotation/reflection}

What is different about infinite dimensions?

1. **Hamel Bases versus Schauder Bases.** finite combinations get everything versus infinite sums of a countable basis gets everything. (These exist if and only if the space is separable.)
2. **Subspaces need not be closed.** For example, take any Schauder basis $S \equiv \{s_i\}_{i=1}^{\infty}$ and consider all finite linear combinations of elements of S. The result is a subspace but it is not closed.
3. **All norms are not equivalent.** For example: the 1-norm of the function $\frac{1}{\sqrt{x}}$ on the unit interval is finite but the 2-norm is infinite.
4. **The unit ball is not compact.** Using the topology induced by the norm, the unit ball is not compact.
5. **Not all linear operators are continuous.** bounded = continuous.
6. **The spectrum is complicated.** There are multiple ways that $A - \lambda I$ can fail to be non-invertible. Each way generates different types of elements of the spectrum.
7. **Proving the spectral theorem.** Proving the spectral theorem for normal operators in Banach spaces is *very* involved. Proof of this statement: See Conway's book [4] on functional analysis and his proof of the spectral theorem for normal operators. For example, it involves measures on the complex plane which take values in the space of projection operators.
8. **Hilbert spaces are easier than Banach spaces.** Having an inner product and a notion of orthogonality makes many things easier/possible.
9. **Continuous, self-adjoint operators on Hilbert spaces are nice.** ... things are fairly similar to finite dimensions
10. **Recommendation.** Read through Chapters 1 and 2 of Cheney's book [3], mentioned below in the "Further Reading" Chapter, to get a sense for the main results in infinite dimensional linear theory.

Bibliography

[1] William M. Boothby. *An Introduction to Differentiable Manifolds and Riemannian Geometry*. Academic Press, 1975.

[2] Robert Pemberton Burn. *Numbers and functions: steps into analysis*. Cambridge University Press, 2015.

[3] E. Ward Cheney. *Analysis for Applied Mathematics*. Springer, 2001.

[4] John B. Conway. *A Course in Functional Analysis*, volume 96 of *Graduate Texts in Mathematics*. Springer Verlag, 2nd edition, 1990.

[5] Luc Devroye, László Györfi, and Gabor Lugosi. *A probabilistic Theory of Pattern Recognition*, volume 31 of *Applications of Mathematics*. Springer, 1996.

[6] Manfredo Perdigao Do Carmo. *Riemannian geometry*. Birkhauser, 1992.

[7] Lawrence C Evans and Ronald F Gariepy. *Measure theory and fine properties of functions*, volume 5. CRC press, 1992.

[8] Herbert Federer. Curvature measures. *Transactions of the American Mathematical Society*, 93(3):418–491, 1959.

[9] Steven G. Krantz and Harold R. Parks. *Geometric Integration Theory*. Birkhauser, 2008. ISBN-13: 978-0-8176-4676-9.

[10] John M. Lee. *Riemannian Manifolds: An Introduction to Curvature*, volume 176 of *Graduate Texts in Mathematics*. Springer-Verlag, 1997. ISBN 0-387-98322-8.

[11] John M. Lee. *Introduction to Smooth Manifolds*, volume 218 of *Graduate Texts in Mathematics*. Springer, 2000. http://www.math.washington.edu/~lee/Courses/545-2000/notes.html.

[12] John M. Lee. *Introduction to Topological Manifolds*, volume 202 of *Graduate Texts in Mathematics*. Springer-Verlag, 2000. ISBN 0-387-95026-5.

[13] Frank Morgan. *Geometric measure theory: a beginner's guide*. Academic Press, fourth edition, 2008.