

MATH 547, STATISTICAL LEARNING THEORY, FALL 2019

STEVEN HEILMAN

CONTENTS

1. Introduction	1
2. A General Supervised Learning Problem	13
2.1. The Perceptron Algorithm	15
2.2. Embeddings and the “Kernel Trick”	18
2.3. Optional: Proof of Mercer’s Theorem	21
3. Probably Approximately Correct (PAC) Learning	23
3.1. Learning Boolean Conjunctions	24
3.2. Learning DNF Formulas	25
3.3. Boosting	27
3.4. Occam’s Razor	31
3.5. Additional Comments	32
4. Vapnis-Chervonenkis (VC) Theory	33
4.1. Applications of the Fundamental Theorem	38
5. Some Concentration of Measure	39
5.1. Concentration for Independent Sums	39
5.2. Concentration for Lipschitz Functions	40
5.3. Additional Comments	43
6. Empirical Risk Minimization (ERM) and Concentration	46
6.1. Gaussian Processes	47
6.2. Sub-Gaussian Processes	51
6.3. General Empirical Processes	51
6.4. Additional Comments	51
7. Concentration of Empirical Processes	51
8. Deep Learning	51
9. Appendix: Basics of Complexity Theory	51
10. Appendix: Some Functional Analysis	54
11. Appendix: Isoperimetric Inequalities	84
12. Appendix: Notation	99
References	101

1. INTRODUCTION

With four parameters I can fit an elephant, and with five I can make him wiggle his trunk.

Date: August 6, 2019 © 2019 Steven Heilman, All Rights Reserved.

John von Neumann

In Statistical Learning Theory, i.e. the statistical theory of machine learning, we will focus on the following questions:

- How do we choose a model that fits the data? (This is the one basic question in statistics.)
- How do we find the best model parameter? (This is the another basic question in statistics.)
- What algorithm for solving a given problem is most efficient? (Here “most efficient” can have several meanings, such as run time or memory usage, in the worst case or average case.)
- Even if we can efficiently find the best model parameter in a statistical problem, is that parameter meaningful? (As the quote of von Neumann points out, it may not be meaningful to fit a model with too many parameters to data. Moreover, finding correlations in data is not often meaningful. For example, the number of Nobel prizes awarded to a country is highly correlated with that country’s chocolate consumption, but this correlation is not at all meaningful. See also Exercise 1.1.)
- Can we design algorithms that function when data arrives in a stream? (That is, suppose we only have access to a small amount of a larger group of data at any time, e.g. due to memory constraints on a computer. Can we still find good algorithms in this case?)

When answering these questions, we should also consider the following dichotomies for algorithms:

- Deterministic algorithms vs. randomized algorithms (i.e. those that use randomness)
- Exact vs. approximation algorithms. (Sometimes an exact efficient algorithm does not exist, while an efficient algorithm that is approximately correct does exist.)
- Theoretical vs. practical algorithms. (Some algorithms work well in theory, but implied constants in their run times can be so large that such algorithms are impractical. Also, some algorithms work well according to practitioners, but there is no known theoretical guarantee that they work well for an arbitrary data set.)
- Supervised vs. unsupervised learning. (Supervised learning uses labelled data, unsupervised learning does not.)

In practice it is also important to consider when an algorithm can run in parallel on different computers, but we will not focus on this topic in this course.

Landau’s Asymptotic Notation. Let $f, g: \mathbb{R} \rightarrow \mathbb{C}$. We use the notation $f(t) = o(g(t))$, $\forall t \in \mathbb{R}$ to denote $\lim_{t \rightarrow \infty} \left| \frac{f(t)}{g(t)} \right| = 0$. We use the notation $f(t) = O(g(t))$ to denote that $\exists c > 0$ such that $|f(t)| \leq c|g(t)|$ for all $t \in \mathbb{R}$. We write $f(t) = \Omega(g(t))$ when $\exists c > 0$ such that $|f(t)| \geq c|g(t)|$ for all $t \in \mathbb{R}$. We write $f(t) = \Theta(g(t))$ when $f(t) = O(g(t))$ and $g(t) = O(f(t))$.

Standard Norm/Inner Product Notation. For any $x = (x_1, \dots, x_n), y = (y_1, \dots, y_n) \in \mathbb{R}^n$, define the standard inner product $\langle x, y \rangle := \sum_{i=1}^n x_i y_i$. We also denote $\|x\| := (\sum_{i=1}^n x_i^2)^{1/2}$ as the standard norm on \mathbb{R}^n .

For more notation see Section 12.

Exercise 1.1. Let $x^{(1)}, \dots, x^{(m)}$ be m vectors in \mathbb{R}^n with $\|x^{(i)}\| = 1$ for all $1 \leq i \leq m$. Let $\varepsilon > 0$. Assume that $m > (1 + 2/\varepsilon)^n$. Show that there exists $i, j \in \{1, \dots, m\}$ such that $\|x^{(i)} - x^{(j)}\| < \varepsilon$.

Consequently, the vectors $x^{(i)}$ and $x^{(j)}$ are highly correlated, so that $\langle x^{(i)}, x^{(j)} \rangle > 1 - \varepsilon^2/2$. That is, if you have enough vectors on a unit sphere, at least two of them will be correlated with each other.

(If you want a big hint, look ahead to Proposition 4.3.)

To better understand our basic questions and dichotomies, we consider them for several specific examples.

Example 1.2 (Computing Determinants). Let $n > 0$ be an integer. Suppose we want to compute the determinant of a real $n \times n$ matrix A with entries A_{ij} , $i, j \in \{1, \dots, n\}$. An inefficient but straightforward way to do this is to directly use a definition of the determinant. Let S_n denote the set of all permutations on n elements. For any $\sigma \in S_n$, let $\text{sign}(\sigma) := (-1)^j$, where σ can be written as a composition of j transpositions (Exercise: this quantity is well-defined). (A transposition $\sigma \in S_n$ satisfies $\sigma(i) = i$ for at least $n - 2$ elements of $\{1, \dots, n\}$.) Then

$$\det(A) = \sum_{\sigma \in S_n} \text{sign}(\sigma) \prod_{i=1}^n A_{i\sigma(i)}.$$

This sum has $|S_n| = n!$ terms. So, if we use this formula to directly compute the determinant of A , in the worst case we will need to perform at least $(n + 1) \cdot n!$ arithmetic operations. This is quite inefficient. We know a better algorithm from linear algebra class. We first perform row operations on A to make it upper triangular. Suppose B is an $n \times n$ real matrix such that BA represents one single row operation on A (i.e. adding a multiple of one row to another row). Then there are real $n \times n$ matrices B_1, \dots, B_m such that

$$B_1 \cdots B_m A \quad (*)$$

is an upper triangular matrix. The matrices B_1, \dots, B_m can be chosen to first eliminate the left-most column of A under the diagonal, then the second left-most column entries under the diagonal, and so on. That is, we can choose $m \leq n(n - 1)/2$, and each row operation involves at most $3n$ arithmetic operations. So, the multiplication of $(*)$ uses at most

$$3mn \leq 2n^3$$

arithmetic operations. The determinant of the upper diagonal matrix $(*)$ is then the product of its diagonal elements, and

$$\det(B_1 \cdots B_m A) = \det(B_1) \cdots \det(B_m) \det(A).$$

That is,

$$\det(A) = \frac{\det(B_1 \cdots B_m A)}{\det(B_1) \cdots \det(B_m)}.$$

So, $\det(A)$ can be computed with at most $2n^3 + m + n \leq 4n^3 = O(n^3)$ arithmetic operations. Can we do any better?

It turns out that this is possible. Indeed, if it is possible to multiply two $n \times n$ real matrices with $O(n^a)$ arithmetic operations for some $a > 0$, then it is possible to compute the determinant of an $n \times n$ matrix with $O(n^a)$ arithmetic operations. The naïve way to

multiply two real $n \times n$ matrices requires $O(n^3)$ arithmetic operations, so that $a = 3$ is achievable. However $a < 2.3728639$ is the best known upper bound [Gal14] (building upon Coppersmith-Winograd, Stothers, and Williams.) I do not think the algorithm with such a value of a has been implemented in practice, since the implied constants in its analysis are quite large, and apparently the algorithm does not parallelize. On the other hand, Strassen’s algorithm has been implemented, and it has $a = \log 7 / \log 2 \approx 2.807$.

What if we only have access to the matrix in the streaming setting? That is, suppose n is so large or the memory of the computer is so limited that we can only store a few of the rows of the matrix at one time. In this case, it is impossible to know the determinant of the whole matrix. For example, suppose the first $n - 1$ rows of the matrix A are known linearly independent vectors, and the last row of the matrix is known, except its last entry. If the last row is identical to the first row (except for their last entries), then changing the last entry could make the determinant of A zero or nonzero. So, it seems like we need to know essentially the whole matrix to even approximately know the determinant of the matrix, or even more fundamentally, the rank of the matrix. Indeed this is the case [CW09].

All of the above algorithms are deterministic, and they compute the determinant exactly (up to machine precision).

Remark 1.3. Interestingly, computing the permanent of a matrix

$$\text{per}(A) = \sum_{\sigma \in S_n} \prod_{i=1}^n A_{i\sigma(i)}$$

is $\#\mathbf{P}$ -complete. However, for any $\varepsilon > 0$, there is a $(1 + \varepsilon)$ polynomial time approximation algorithm for computing the permanent.

Example 1.4 (Least Squares). Suppose we want to solve a least squares minimization problem. Suppose $x \in \mathbb{R}^m$ is an unknown vector, and A is a known $m \times n$ real matrix. Let $Z \in \mathbb{R}^n$ be a vector of i.i.d. standard Gaussian random variables. Our observation is $y := Aw + Z$, and the goal is to recover the unknown vector w . In linear least squares regression, we try to determine the best linear relationship w between the rows of A and the observation y . Assume that $n \leq m$ and the matrix A has full rank (so that $A^T A$ is invertible). The vector $x \in \mathbb{R}^m$ that minimizes the quantity

$$\|y - Ax\|^2 := \sum_{i=1}^n (y_i - (Ax)_i)^2$$

is then

$$x := (A^T A)^{-1} A^T y. \quad (*)$$

Equivalently, x minimizes

$$\mathbf{E} \|y - x\|^2$$

over all choices of vectors $x \in \mathbb{R}^m$ such that $x = By$ for some $n \times m$ real matrix B , and such that $\mathbf{E}x = w$. (Since Z is the only random variable here, \mathbf{E} denotes expected value with respect to Z .)

So, solving the linear least squares minimization problem could “just” use equation (*). However, inverting a matrix directly is inefficient and could introduce numerical error. Below are some alternative ways of computing x . From (*), we first write

$$A^T Ax = A^T y,$$

and we then do any one of the following.

- Compute the Cholesky decomposition of $A^T A$. That is, we write $A^T A = R^T R$ where R is an upper triangular $n \times n$ matrix with positive diagonal elements. Then, $R^T R x = A^T y$, and solve the following simpler problems: (1) solve $R^T z = A^T y$ for the unknown $z \in \mathbb{R}^n$, then (2) solve $R x = z$ for $x \in \mathbb{R}^n$. This is our desired x , since

$$R^T R x = R^T z = A^T y.$$

- Compute the QR decomposition of the matrix A . That is, we write $A = QR$ where Q is an $m \times n$ rectangular matrix with $Q^T Q = I_n$, where I_n denotes the $n \times n$ identity matrix, and R is an upper triangular $n \times n$ matrix with positive diagonal elements. Then $(A^T A)^{-1} = (R^T Q^T Q R)^{-1} = (R^T R)^{-1}$, so $(A^T A)^{-1} A^T = (R^T R)^{-1} R^T Q^T$, and $(R^T R)^{-1} R^T = R^{-1}$, so we have

$$x \stackrel{(*)}{=} R^{-1} Q^T y.$$

- Compute the singular value decomposition of the matrix A . (That is, we write $A = USV$, where U is an $m \times m$ orthogonal matrix, V is an $n \times n$ orthogonal matrix, and S is an $n \times n$ diagonal matrix with nonnegative entries.) Then $A^T A x = V^T S^2 V x = A^T y$, and solve the following simpler problems: (1) solve $V^T z = A^T y$ for the unknown $z \in \mathbb{R}^n$, then (2) solve $S^2 V x = z$ for $x \in \mathbb{R}^n$. This is our desired x , since

$$A^T A x = V^T S^2 V x = V^T z = A^T y.$$

These algorithms all use $O(n^3)$ arithmetic operations, and they assume we have access to the whole matrix A .

What if we only have streaming access to A ? Put another way, is there an algorithm that only needs to access a small amount of A at any single time? The answer is yes. Recall that A is an $m \times n$ matrix with $m \geq n$. In some cases, m will be much larger than n . The recursive algorithm described below only needs to store an $n \times n$ matrix of A at any given time, and it is given access to one row of A at a time.

For a more modern view of this problem, see e.g. [KKP17].

Exercise 1.5. Let A be an $m \times n$ real matrix with $m \geq n$. Then A has rank n if and only if $A^T A$ is positive definite.

(Hint: $A^T A$ is always positive semidefinite.)

Algorithm 1.6 (Recursive Least Squares/ Online Learning). Let $m \geq n$, let A be an $m \times n$ real matrix. Let $a^{(1)}, \dots, a^{(m)} \in \mathbb{R}^n$ be row vectors which are the rows of A (data), and let $b \in \mathbb{R}^m$. For any $j \geq n$, let

$$A_j := \begin{pmatrix} a^{(1)} \\ \vdots \\ a^{(j)} \end{pmatrix}, \quad b^{(j)} := \begin{pmatrix} b_1 \\ \vdots \\ b_j \end{pmatrix}.$$

Assume that A_n has rank n (so that $A_n^T A_n$ is invertible). Define

$$x^{(n)} := (A_n^T A_n)^{-1} A_n^T b^{(n)} \in \mathbb{R}^n, \quad P_n := (A_n^T A_n)^{-1}.$$

For any $j \geq n$, define

$$P_{j+1} = P_j - \frac{P_j(a^{(j+1)})^T a^{(j+1)} P_j^T}{1 + a^{(j+1)} P_j (a^{(j+1)})^T}.$$

$$x^{(j+1)} = x^{(j)} + P_{j+1} (a^{(j+1)})^T (b_{j+1} - a^{(j+1)} x^{(j)}).$$

The vectors $x^{(n)}, \dots, x^{(m)}$ recursively minimize the quantity $\|Ax - b\|^2$ in the following sense.

Proposition 1.7. *Let $\lambda > 0$. Let $x^{(n)}, \dots, x^{(m)}$ be the output of Algorithm 1.6. Let $n \leq j \leq m$. Define $f_j: \mathbb{R}^n \rightarrow \mathbb{R}$ by*

$$f_j(x) := \frac{1}{2} \sum_{i=1}^j (\langle x, a^{(i)} \rangle - b_i)^2, \quad x \in \mathbb{R}^n.$$

Then $x^{(j)}$ minimizes f_j on \mathbb{R}^n . In particular, when $j = m$, $x^{(m)}$ minimizes $\|Ax - b\|^2$.

Proof. We induct on j . The case $j = n$ follows by definition of $x^{(n)}$ and by (*). We now complete the inductive step. Assume the Proposition holds for j , and consider the case $j + 1$. Define $G_j := (A_j^T A_j)$.

First, note that

$$G_{j+1} = \begin{pmatrix} A_j^T & (a^{(j+1)})^T \end{pmatrix} \begin{pmatrix} A_j \\ a^{(j+1)} \end{pmatrix} = A_j^T A_j + (a^{(j+1)})^T a^{(j+1)} = G_j + (a^{(j+1)})^T a^{(j+1)}. \quad (*)$$

By the inductive hypothesis and (*), we have $x^{(j)} = G_j^{-1} A_j^T b^{(j)}$. So,

$$A_j^T b^{(j)} = G_j G_j^{-1} A_j^T b^{(j)} = G_j x^{(j)} \stackrel{(*)}{=} (G_{j+1} - (a^{(j+1)})^T a^{(j+1)}) x^{(j)}. \quad (**)$$

From (*), the minimum of f_{j+1} on \mathbb{R}^n occurs when

$$\begin{aligned} x &= G_{j+1}^{-1} A_{j+1}^T b^{(j+1)} = G_{j+1}^{-1} \begin{pmatrix} A_j \\ a^{(j+1)} \end{pmatrix}^T \begin{pmatrix} b^{(j)} \\ b_{j+1} \end{pmatrix} = G_{j+1}^{-1} (A_j^T b^{(j)} + b_{j+1} (a^{(j+1)})^T) \\ &\stackrel{(**)}{=} G_{j+1}^{-1} (G_{j+1} x^{(j)} - (a^{(j+1)})^T a^{(j+1)} x^{(j)} + b_{j+1} (a^{(j+1)})^T) \\ &= x^{(j)} + G_{j+1}^{-1} (a^{(j+1)})^T (b_{j+1} - a^{(j+1)} x^{(j)}). \end{aligned}$$

Comparing this formula to the definition of $x^{(j+1)}$ in Algorithm 1.6, it remains to manipulate the G_{j+1}^{-1} term. Applying Exercise 1.8 to (*),

$$G_{j+1}^{-1} = (G_j + (a^{(j+1)})^T a^{(j+1)})^{-1} = G_j^{-1} - \frac{G_j^{-1} (a^{(j+1)})^T a^{(j+1)} G_j^{-1}}{1 + a^{(j+1)} G_j^{-1} (a^{(j+1)})^T}.$$

Finally, note that $P_n = G_n^{-1}$, and since the matrices P_j and G_j^{-1} satisfy the same recursion, we get $P_j = G_j^{-1}$, completing the proof. □

Exercise 1.8. Show the following identity. Let A be an $r \times r$ real matrix, let U be an $r \times s$ real matrix, and let V be an $s \times r$ real matrix. Assume that A is invertible and that $I + VA^{-1}U$ is invertible, where I is the $s \times s$ identity matrix. Then $A + UV$ is invertible and

$$(A + UV)^{-1} = A^{-1} - (A^{-1}U)(I + VA^{-1}U)^{-1}(VA^{-1}).$$

In particular, if $s = 1$, we get the Sherman-Morrison formula:

$$(A + UV)^{-1} = A^{-1} - \frac{A^{-1}UV A^{-1}}{1 + VA^{-1}U}.$$

Example 1.9 (Minimum Vertex Cover). Suppose we have a set of vertices $V := \{1, \dots, n\}$ and a set of undirected edges $E \subseteq \{\{i, j\} : i, j \in V\}$. The goal of the minimum vertex cover problem is to find the smallest vertex cover of the graph $G = (V, E)$. A vertex cover is a subset $S \subseteq V$ such that every $\{i, j\} \in E$ satisfies $i \in S$ or $j \in S$. More generally, for any $i \in V$, let $c_i \in \mathbb{R}$, $c_i \geq 0$. We are asked to minimize the weighted sum

$$\sum_{i \in S} c_i$$

over all $S \subseteq V$ such that every $\{i, j\} \in E$ satisfies $i \in S$ or $j \in S$. (To recover the unweighted minimum vertex cover problem, let $c_i := 1$ for all $1 \leq i \leq n$.) For a somewhat contrived example, we could think of the vertices as cities, and the set S as a subset of cities where cell phone towers are placed. And each cell phone tower is designed to cover the city in which it resides, and any adjacent cities.

There is a simple polynomial time algorithm that, given $G = (V, E)$ whose minimum vertex cover has size $a > 0$, finds a vertex cover $S \subseteq V$ such that

$$a \leq |S| \leq 2a.$$

This algorithm is therefore called a 2-approximation algorithm for the minimum vertex cover problem.

It is known that there exists an $\alpha > 1$ such that finding an α -approximation to the minimum vertex cover problem is NP-complete. So, it seems impossible to efficiently solve the minimum vertex cover problem. More specifically, if we were given $G = (V, E)$ whose minimum vertex cover has size $c > 0$, and we could find a vertex cover $S \subseteq V$ in polynomial time such that

$$c \leq |S| \leq \alpha c,$$

then $P = NP$, and we would solve one of the Millennium Prize Problems. Since it is widely believed that $P \neq NP$, it is doubtful that the Minimum Vertex Cover problem can be solved in time polynomial in n . In fact, the constant $\alpha = 2$ is believed to be the best possible approximation. That is, it is conjectured that it is NP-hard to find a 1.9999-approximation to the minimum vertex cover problem [KR08].

Let us describe the deterministic algorithm.

Algorithm 1.10 (Greedy Algorithm for Unweighted Min-Vertex Cover). Given a graph $G = (V, E)$, begin with $S := \emptyset$. Iterate the following procedure until E is empty.

- Choose some edge $\{i, j\} \in E$.
- Redefine $S := S \cup \{i\} \cup \{j\}$, and remove from E any edge whose endpoint is either i or j .

When E is empty, output S .

Proposition 1.11. *Algorithm 1.10 is a 2-approximation algorithm for the unweighted minimum vertex cover problem.*

Proof. Let a be the size of the minimum vertex cover of a graph $G = (V, E)$. Let E' be the set of edges that is found in the first step of all iterations of Algorithm 1.10. By step two of the iteration, each edge in E' is encountered exactly once in step one, and the vertices of each edge in E' are all disjoint from each other. Moreover, steps one and two of the iteration imply that S is a vertex cover. Also, step two implies that the minimum vertex cover must contain at least one endpoint of each edge in E' , otherwise some edge would be not covered by the minimum vertex cover. So, $|E'| \leq a$. In summary,

$$|S| = 2|E'| \leq 2a.$$

Lastly, $|S| \geq a$ by definition of c . □

There is also a randomized 2-approximation algorithm for the general (weighted) minimum vertex problem using linear programming. The Minimum Vertex Cover problem is:

$$\begin{aligned} &\text{minimize} && \langle c, x \rangle && \text{subject to the constraints} \\ & && x_i \in \{0, 1\}, && \forall i \in V \\ & && x_i + x_j \geq 1, && \forall \{i, j\} \in E \end{aligned}$$

We then solve the following linear program and try to relate its solution to the above problem.

$$\begin{aligned} &\text{minimize} && \langle c, x \rangle && \text{subject to the constraints} \\ & && x_i \in [0, 1], && \forall i \in V \\ & && x_i + x_j \geq 1, && \forall \{i, j\} \in E \end{aligned}$$

Recall that an $n \times n$ matrix A is **symmetric** if $A = A^T$. We denote the $n \times n$ identity matrix as I or I_n . Recall that an $n \times n$ matrix A is **orthogonal** if A is a real matrix such that $A^T A = A A^T = I$.

Theorem 1.12 (Spectral Theorem for Real Symmetric Matrices). *Let A be a real $n \times n$ symmetric matrix. Then there exists an $n \times n$ orthogonal matrix Q and a real diagonal matrix D such that*

$$A = Q^T D Q.$$

Equivalently, there exists an orthonormal basis of \mathbb{R}^n consisting of eigenvectors of A with real eigenvalues.

Example 1.13 (Singular Value Decomposition (SVD) and Principal Component Analysis (PCA)). Suppose A is an $m \times n$ real matrix. We think of each row of A as a vector of data. The matrix $A^T A$ is then an $n \times n$ real symmetric matrix. By the Spectral Theorem 1.12, there exists an $n \times n$ orthogonal matrix Q and a real diagonal matrix D such that

$$A^T A = Q^T D Q.$$

Similarly, the matrix $A A^T$ is an $m \times m$ real symmetric matrix. The Spectral Theorem 1.12 implies that there exists an $m \times m$ orthogonal matrix R and a real diagonal matrix G such that

$$A A^T = R^T G R.$$

Note that all entries of D are nonnegative, since if $v \in \mathbb{R}^n$ is an eigenvector of $A^T A$ with eigenvalue $\lambda \in \mathbb{R}$, then

$$\lambda \|v\|^2 = \lambda \langle v, v \rangle = \langle v, A^T A v \rangle = \langle A v, A v \rangle = \|A v\|^2.$$

That is, we must have $\lambda \geq 0$. So, the square root of D (i.e. the diagonal matrix all of whose entries are the square roots of the entries of D) is well-defined. We denote the square root of D as \sqrt{D} , so that $(\sqrt{D})^2 = D$. Note that we can then write

$$A = R^T \sqrt{D} Q.$$

This factorization of A is called a **singular value decomposition** of A . (In general, the matrices R^T, \sqrt{D}, Q are not uniquely determined by A , though the entries of \sqrt{D} are uniquely determined by A .) The entries of \sqrt{D} are called the **singular values** of A .

So, at least theoretically, a singular value decomposition exists. How can we find it on a computer? If the matrix A is relatively small, then we can compute $A^T A$ and find its eigenvalues and eigenvectors by the Power Method (see Exercise 1.18 below). The matrix D then consists of the eigenvalues and Q contains the eigenvectors of $A^T A$. The power method is especially efficient when A has very few nonzero entries.

If the matrix A is fairly large, we can instead randomly sample a small number of the rows and columns of A , perform a singular value decomposition on this smaller matrix, and use it to approximate the SVD of A itself [KV09, HMT11, Mah11].

SVD is used widely in practice. Note that since each data vector is a vector in \mathbb{R}^n , and each eigenvector of $A^T A$ is also a vector in \mathbb{R}^n , any vector $x \in \mathbb{R}^n$ can be written uniquely as a linear combination of eigenvectors of $A^T A$ (since these eigenvectors form an orthonormal basis of \mathbb{R}^n .) In many applications, it is generally expected that $A^T A$ has low rank or “approximately low rank.” In such cases, the matrix $A^T A$ is well understood by examining its eigenvectors with largest eigenvalues. For example, if $A^T A$ has rank k , then only k of its eigenvectors are needed to understand $A^T A$. Similarly, if $A^T A$ has k large eigenvalues with the rest of them being close to zero, then only the first k eigenvectors are needed to approximately understand $A^T A$, and hence A .

PCA is designed with this observation in mind. In PCA, one uses the SVD of A (or of A with each row of A subtracted by the average of all rows of A), and we examine the eigenvectors of $A^T A$ (and AA^T) with the largest eigenvalues. In particular, the eigenvectors of $A^T A$ with largest eigenvalues are judged to be the “principal components” of any data vector $x \in \mathbb{R}^n$. PCA is used e.g. when Netflix tries to find recommendations for movies or shows you might enjoy. We consider each row of A to contain the preferences of one user, e.g. containing their ratings for various movies (say entry $A_{ij} \in \{1, 2, 3, 4, 5\}$ is the rating of user i for movie j). Suppose the eigenvectors of $A^T A$ with the k largest eigenvalues are $y^{(1)}, \dots, y^{(k)}$. We can then e.g. examine the k^{th} **spectral embedding** $f_k: \{1, \dots, n\} \rightarrow \mathbb{R}^k$ that maps each movie to its vector of values on these eigenvectors:

$$f_k(i) := \left(y_i^{(1)}, \dots, y_i^{(k)} \right), \quad \forall i \in \{1, \dots, n\}.$$

In practice, one observes that “similar” movies appear in clusters in the spectral embedding set $\{f_k(i)\}_{i=1}^n \subseteq \mathbb{R}^k$. We can then infer that someone who enjoys one movie in one cluster will enjoy other movies in the same cluster. The actual application is a bit more complicated since users do not rank all movies, so extra steps are made to fill or infer the missing entries of A . Since the initial data vectors are in \mathbb{R}^n but the set $\{f_k(i)\}_{i=1}^n$ is in \mathbb{R}^k with $k < n$, PCA is considered a **dimension reduction** method.

PCA is also used in facial recognition with the so-called “eigenfaces” method. In this case, each row of A is a vector (say of bitmap image values) of the front view of someone’s

face. PCA is then performed on a standard large data set. Given a new face image $x \in \mathbb{R}^n$, we associate x to its “feature vector” $g_k(x) := (\langle x, y^{(1)} \rangle, \dots, \langle x, y^{(k)} \rangle)$. We can then guess if two different faces $x, y \in \mathbb{R}^n$ are the same or not the same by computing the distance $\|g_k(x) - g_k(y)\|$. Alternatively, if we want to compare x to the set of images in the initial data set (on which the original PCA was performed), we can try to find $\min_{i=1, \dots, n} \|g_k(x) - f_k(i)\|$, and then associate x with the value of $i \in \{1, \dots, n\}$ achieving this minimum. The actual application is more complicated. Images in practice do not often show the front of a face, so extra steps can be made to infer the front view of a face from the image of a non-frontal view of a face.

Since PCA and SVD use unlabelled data, they are considered methods in **unsupervised learning**.

Example 1.14 (k -means Clustering). Let k, m, n be positive integers. Given m vectors $x^{(1)}, \dots, x^{(m)} \in \mathbb{R}^n$, the k -means clustering problem asks for the partition S_1, \dots, S_k of $\{1, \dots, m\}$ into k sets which minimizes

$$\sum_{i=1}^k \sum_{j \in S_i} \left\| x^{(j)} - \frac{1}{|S_i|} \sum_{p \in S_i} x^{(p)} \right\|^2. \quad (*)$$

For each $1 \leq i \leq k$, the term $\frac{1}{|S_i|} \sum_{p \in S_i} x^{(p)}$ is the center of mass (or barycenter) of the points in S_i , so each term in the sum is the squared distance of some point in S_i from the barycenter of S_i . So, k -means clustering can be seen as a kind of geometric version of least-squares regression. We emphasize that k is fixed.

How can we solve this problem? The most basic algorithm is a “gradient-descent” procedure known as Lloyd’s Algorithm.

Algorithm 1.15 (Lloyd’s Algorithm). Let $x^{(1)}, \dots, x^{(m)} \in \mathbb{R}^n$. Begin by choosing $y^{(1)}, \dots, y^{(k)} \in \mathbb{R}^n$ (randomly or deterministically), and define $T_i := \emptyset$ for all $1 \leq i \leq k$. Repeat the following procedure:

- For each $1 \leq i \leq k$, re-define

$$T_i := \left\{ j \in \{1, \dots, m\} : \|x^{(j)} - y^{(i)}\| = \min_{p=1, \dots, k} \|x^{(j)} - y^{(p)}\| \right\}.$$

(If more than one p achieves this minimum, assign j to an arbitrary such minimal p .) (The sets T_1, \dots, T_m are called Voronoi regions.)

- For each $1 \leq i \leq k$, re-define $y^{(i)} := \frac{1}{|T_i|} \sum_{p \in T_i} x^{(p)}$.

Once this procedure is iterated a specified number of times, output $S_i := T_i$.

Algorithm 1.15 can be considered a “gradient-descent” procedure since the first step of the iteration always decreases the quantity $(*)$ by the definition of $(*)$, and the second step of the iteration always decreases $(*)$ by Exercise 1.16.

Exercise 1.16. Let $x^{(1)}, \dots, x^{(m)} \in \mathbb{R}^n$. Let $y \in \mathbb{R}^n$. Show that

$$\sum_{j=1}^m \left\| x^{(j)} - \frac{1}{m} \sum_{p=1}^m x^{(p)} \right\|^2 \leq \sum_{j=1}^m \|x^{(j)} - y\|^2.$$

That is, the barycenter is the point in \mathbb{R}^n that minimizes the sum of squared distances.

While Lloyd’s Algorithm 1.15 decreases the value of the quantity $(*)$, iterating this algorithm many times does not guarantee that a global minimum of $(*)$ is found. To see why, recall that the local minimum of a function $f: \mathbb{R} \rightarrow \mathbb{R}$ may not be the same as a global minimum. So, while Lloyd’s Algorithm 1.15 is simple and it might work well in certain situations, it has no general theoretical guarantees. Some work has been done to make a “wise” choice of the initial points $y^{(1)}, \dots, y^{(k)}$.

So, are there any efficient algorithms with theoretical guarantees? For any $\varepsilon > 0$, there is a $9 + \varepsilon$ factor approximation algorithm for the k -means clustering problem [KMN⁺04] with a polynomial running time (that does not depend on k). This algorithm is based upon [Mat00]. It was shown [ACKS15] that there exists some $\varepsilon > 0$ such that approximating the k -means clustering problem with a multiplicative factor of $1 + \varepsilon$ for all k is NP-hard. So, there is a rather large gap between the best general purpose algorithm, and the hardness result. Many algorithms can approximately solve the k -means clustering problem to a multiplicative factor of $1 + \varepsilon$, but these algorithms always have an exponential dependence on k for their run times [HPM04]. So, if we try to use $k = 100$, which occurs in many applications, these algorithms seem to be impractical.

It is possible to combine dimension-reduction techniques (such as PCA or the Johnson-Lindenstrass Lemma, Theorem 5.6) with the above algorithms [CEM⁺15, MMR18], thereby saving time by working in lower dimensions. However, these techniques do not seem to improve the exponential run times in k .

Some streaming algorithms are known for k -means clustering [Che09, FMS07]. For example, the algorithm of [Che09] uses memory of size $O(d^2 k^2 \varepsilon^{-2} (\log n)^8)$ to approximately solve the k -means problem within a multiplicative factor of $1 + \varepsilon$. Note that the points themselves are not actually stored in this algorithm, otherwise the memory requirement would be at least $\Omega(n)$. In fact, only the barycenters of the clusters are typically stored in these streaming algorithms, which drastically reduces the memory requirement.

Since k -means clustering uses unlabelled data (despite the fact that k needs to be specified), it is considered a method in **unsupervised learning**.

Exercise 1.17. Let $n \geq 2$ be a positive integer. Let $x = (x_1, \dots, x_n) \in \mathbb{R}^n$. For any $x, y \in \mathbb{R}^n$, define $\langle x, y \rangle := \sum_{i=1}^n x_i y_i$ and $\|x\| := \langle x, x \rangle^{1/2}$. Let $S^{n-1} := \{x \in \mathbb{R}^n : \|x\| = 1\}$ be the sphere of radius 1 centered at the origin. Let $x \in S^{n-1}$ be fixed. Let v be a random vector that is uniformly distributed in S^{n-1} . Prove:

$$\mathbf{E} |\langle x, v \rangle| \geq \frac{1}{10\sqrt{n}}.$$

Exercise 1.18 (The Power Method). This exercise gives an algorithm for finding the eigenvectors and eigenvalues of a symmetric matrix. In modern statistics, this is often a useful thing to do. The Power Method described below is not the best algorithm for this task, but it is perhaps the easiest to describe and analyze.

Let A be an $n \times n$ real symmetric matrix. Let $\lambda_1 \geq \dots \geq \lambda_n$ be the (unknown) eigenvalues of A , and let $v_1, \dots, v_n \in \mathbb{R}^n$ be the corresponding (unknown) eigenvectors of A such that $\|v_i\| = 1$ and such that $Av_i = \lambda_i v_i$ for all $1 \leq i \leq n$.

Given A , our first goal is to find v_1 and λ_1 . For simplicity, assume that $1/2 < \lambda_1 < 1$, and $0 \leq \lambda_n \leq \dots \leq \lambda_2 < 1/4$. Suppose we have found a vector $v \in \mathbb{R}^n$ such that $\|v\| = 1$ and $|\langle v, v_1 \rangle| > 1/n$. (From Exercise 1.17, a randomly chosen v satisfies this property.) Let k be

a positive integer. Show that

$$A^k v$$

approximates v_1 well as k becomes large. More specifically, show that for all $k \geq 1$,

$$\|A^k v - \langle v, v_1 \rangle \lambda_1^k v_1\|^2 \leq \frac{n-1}{16^k}.$$

(Hint: use the spectral theorem for symmetric matrices.)

Since $|\langle v, v_1 \rangle| \lambda_1^k > 2^{-k}/n$, this inequality implies that $A^k v$ is approximately an eigenvector of A with eigenvalue λ_1 . That is, by the triangle inequality,

$$\|A(A^k v) - \lambda_1(A^k v)\| \leq \|A^{k+1} v - \langle v, v_1 \rangle \lambda_1^{k+1} v_1\| + \lambda_1 \|\langle v, v_1 \rangle \lambda_1^k v_1 - A^k v\| \leq 2 \frac{\sqrt{n-1}}{4^k}.$$

Moreover, by the reverse triangle inequality,

$$\|A^k v\| = \|A^k v - \langle v, v_1 \rangle \lambda_1^k v_1 + \langle v, v_1 \rangle \lambda_1^k v_1\| \geq \frac{1}{n} 2^{-k} - \frac{\sqrt{n-1}}{4^k}.$$

In conclusion, if we take k to be large (say $k > 10 \log n$), and if we define $z := A^k v$, then z is approximately an eigenvector of A , that is

$$\left\| A \frac{A^k v}{\|A^k v\|} - \lambda_1 \frac{A^k v}{\|A^k v\|} \right\| \leq 4n^{3/2} 2^{-k} \leq 4n^{-4}.$$

And to approximately find the first eigenvalue λ_1 , we simply compute

$$\frac{z^T A z}{z^T z}.$$

That is, we have approximately found the first eigenvector and eigenvalue of A .

Remarks. To find the second eigenvector and eigenvalue, we can repeat the above procedure, where we start by choosing v such that $\langle v, v_1 \rangle = 0$, $\|v\| = 1$ and $|\langle v, v_2 \rangle| > 1/(10\sqrt{n})$. To find the third eigenvector and eigenvalue, we can repeat the above procedure, where we start by choosing v such that $\langle v, v_1 \rangle = \langle v, v_2 \rangle = 0$, $\|v\| = 1$ and $|\langle v, v_3 \rangle| > 1/(10\sqrt{n})$. And so on.

Google's PageRank algorithm uses the power method to rank websites very rapidly. In particular, they let n be the number of websites on the internet (so that n is roughly 10^9). They then define an $n \times n$ matrix C where $C_{ij} = 1$ if there is a hyperlink between websites i and j , and $C_{ij} = 0$ otherwise. Then, they let B be an $n \times n$ matrix such that B_{ij} is 1 divided by the number of 1's in the i^{th} row of C , if $C_{ij} = 1$, and $B_{ij} = 0$ otherwise. Finally, they define

$$A = (.85)B + (.15)D/n$$

where D is an $n \times n$ matrix all of whose entries are 1.

The power method finds the eigenvector v_1 of A , and the size of the i^{th} entry of v_1 is proportional to the "rank" of website i .

Exercise 1.19. Run PCA on a "planted" data set on a computer, consisting of 100 samples in \mathbb{R}^{10} of the random variable $(X, Y, Z_3, \dots, Z_{10}) \in \mathbb{R}^{10}$ where X, Y are standard Gaussian random variables, Z_i is a mean i Gaussian random variable with variance 10^{-2} , for all $3 \leq i \leq 10$, and X, Y, Z_3, \dots, Z_{10} are all independent. (You can use your favorite computer program to simulate the random variables.)

Then, run PCA on Airline Safety Information, and try to find out something interesting (this part of the question is intentionally open ended). The data is [here](#), with accompanying article [here](#). (See also [here](#).)

Exercise 1.20. Run a k -means clustering algorithm (e.g. Lloyd’s algorithm) on a “planted” data set in \mathbb{R}^2 consisting of 50 samples from (X, Y) and another 50 samples from (Z, W) where X, Y, Z, W are all independent Gaussians with variance 1, X, W have mean zero, Y has mean 1 and Z has mean 2. Try at least the values $k = 2, 3, 4, 5$.

Then, run a k -means clustering algorithm on Airline Safety Information, and try to find out something interesting (this part of the question is intentionally open ended).

2. A GENERAL SUPERVISED LEARNING PROBLEM

In this course, one main focus is the following.

Problem 2.1 (Supervised Learning Problem). Let A, B be sets. Let $f: A \rightarrow B$ be an unknown function. The goal of the learning problem is to determine the function f on all of A using a small number of known values of f on A . Let $x^{(1)}, \dots, x^{(k)} \in A$ and let $y^{(1)}, \dots, y^{(k)} \in B$. It is known that

$$f(x^{(i)}) = y^{(i)}, \quad \forall 1 \leq i \leq k.$$

We then want to exactly or approximately determine f on all of A .

The set of ordered pairs $\{(x^{(i)}, y^{(i)})\}_{i=1}^k$ is sometimes called the **training set**. The function f is sometimes called a **predictor**, **hypothesis** or **classifier**.

Without some assumptions on f , this problem is impossible to solve, since we could just arbitrarily define f on inputs other than $x^{(1)}, \dots, x^{(k)}$. So, one must add some assumptions on a class of functions f under consideration. One of the most basic and well studied examples is the following subset of Boolean functions $f: \{-1, 1\}^n \rightarrow \{-1, 1\}$.

Definition 2.2 (Linear Threshold Functions). A function $f: \{-1, 1\}^n \rightarrow \{-1, 1\}$ is called a **linear threshold function** if there exists $w \in \mathbb{R}^n$ and $t \in \mathbb{R}$ such that

$$f(x) = \text{sign}(\langle w, x \rangle - t), \quad \forall x \in \{-1, 1\}^n.$$

Here

$$\text{sign}(s) := \frac{s}{|s|} = \begin{cases} 1 & , \text{ if } s > 0 \\ 0 & , \text{ if } s = 0 \\ -1 & , \text{ if } s < 0. \end{cases}$$

A linear threshold function could also be called a **single-layer neural network**.

Here we are assuming that f defined in this way never takes the value 0, and we will always make this assumption unless otherwise stated.

Linear threshold functions can be understood geometrically as two sets of points on either side of a hyperplane.

Definition 2.3 (Hyperplane). Let $w \in \mathbb{R}^n, t \in \mathbb{R}$. A **hyperplane** in \mathbb{R}^n is a set of the form

$$\{x \in \mathbb{R}^n : \langle w, x \rangle = t\}.$$

That is, if $f: \{-1, 1\}^n \rightarrow \{-1, 1\}$ is a linear threshold function, then $\{x \in \{-1, 1\}^n: f(x) = 1\}$ lies on one side of the hyperplane, and $\{x \in \{-1, 1\}^n: f(x) = -1\}$ lies on the other side of the hyperplane. For this reason, linear threshold functions are sometimes called **half spaces**.

If we add the constraint that the unknown function f in Problem 2.1 is a linear threshold function, we arrive at the following problem.

Problem 2.4 (Supervised Learning Problem for Linear Threshold Functions).

Let $f: \{-1, 1\}^n \rightarrow \{-1, 1\}$ be an unknown linear threshold function. Let $x^{(1)}, \dots, x^{(k)} \in \{-1, 1\}^n$ and let $y_1, \dots, y_k \in \{-1, 1\}$. It is known that

$$f(x^{(i)}) = y_i, \quad \forall 1 \leq i \leq k.$$

The goal of the problem is to find $w \in \mathbb{R}^n$ and $t \in \mathbb{R}$ such that

$$\text{sign}(\langle w, x^{(i)} \rangle - t) = y_i, \quad \forall 1 \leq i \leq k.$$

Remark 2.5 (Reduction to Homogeneous $t = 0$ Case). For all $1 \leq i \leq n$, let $z^{(i)} := (x^{(i)}, 1) \in \{-1, 1\}^{n+1}$ and let $w' := (w, t) \in \mathbb{R}^{n+1}$ in Problem 2.4. The goal of Problem 2.4 can then be restated as: find $w' \in \mathbb{R}^{n+1}$ such that

$$\text{sign}(\langle w', z^{(i)} \rangle) = y_i, \quad \forall 1 \leq i \leq k.$$

For this reason, we will often assume $t = 0$ below when discussing Problem 2.4.

Remark 2.6. Problem 2.4 can be understood geometrically as trying to find a hyperplane that separates the points $\{x \in \{-1, 1\}^n: f(x) = 1\}$ from $\{x \in \{-1, 1\}^n: f(x) = -1\}$.

This problem is ill-posed as stated already when $n = 2$ and $k = 4$, since it is impossible to separate pairs of opposite corners of $\{-1, 1\}^2$ using a hyperplane.

Example 2.7. If $f(1, 1) = f(-1, -1) = 1$ and $f(1, -1) = f(-1, 1) = -1$, then there does not exist $w \in \mathbb{R}^2$ and $t \in \mathbb{R}$ such that $\text{sign}(\langle w, x^{(i)} \rangle - t) = y_i \forall 1 \leq i \leq 4$. To see this, we can look at the “partial derivatives” of f as follows. Let $w = (w_1, w_2) \in \mathbb{R}^2$ and suppose $\langle w, (1, 1) \rangle > t$, $\langle w, (-1, -1) \rangle > t$ and $\langle w, (1, -1) \rangle < t$, $\langle w, (-1, 1) \rangle < t$. Then

$$2w_2 = \langle w, (1, 1) - (-1, -1) \rangle > 0, \quad 2w_1 = \langle w, (1, 1) - (-1, 1) \rangle > 0,$$

So,

$$\langle w, (1, 1) \rangle = w_1 + w_2 > -w_1 - w_2 = \langle w, (-1, -1) \rangle.$$

Therefore, $f(1, 1) > f(-1, -1)$, a contradiction.

So, in order for the learning problem 2.4 to be exactly solvable, we must assume that there exists some $w \in \mathbb{R}^n, t \in \mathbb{R}$ such that

$$\text{sign}(\langle w, x^{(i)} \rangle - t) = y_i, \quad \forall 1 \leq i \leq k.$$

If we only want to find a $w \in \mathbb{R}^n, t \in \mathbb{R}$ such that this equality holds for 99% of $1 \leq i \leq k$, then that becomes a much different problem. We will discuss that problem in Section 3.

For a real world example of Problem 2.4, we can think of x as some sequence of letters (encoded in binary) from an email, so that f takes value 1 on a spam email, and f takes value -1 on non-spam emails. The goal is then to find the classifier f that can classify any given email as spam or non-spam, using the “training data” $x^{(1)}, \dots, x^{(k)} \in \{-1, 1\}^n$ and $y_1, \dots, y_k \in \{-1, 1\}$.

2.1. The Perceptron Algorithm. The Perceptron Algorithm is a basic algorithm for solving Problem 2.4 from the 1960's. In this section, we will consider the homogeneous version of Problem 2.4. Let $x^{(1)}, \dots, x^{(k)} \in \{-1, 1\}^n$ and let $y_1, \dots, y_k \in \{-1, 1\}$ be given. It is known that

$$f(x^{(i)}) = y_i, \quad \forall 1 \leq i \leq k.$$

The goal of the problem is to find $w \in \mathbb{R}^n$ such that

$$\text{sign}(\langle w, x^{(i)} \rangle) = y_i, \quad \forall 1 \leq i \leq k.$$

It is assumed that such a w exists.

Algorithm 2.8 (Perceptron Algorithm).

- Define $w^{(1)} := 0 \in \mathbb{R}^n$ and let $s := 1$
- If there exists some $1 \leq i \leq k$ such that $y_i \neq \text{sign}(\langle w^{(s)}, x^{(i)} \rangle)$, i.e. a mis-classification occurs, define

$$w^{(s+1)} := w^{(s)} + y_i x^{(i)}.$$

- Increase the value of s by one. Repeat the previous step until no such i exists.
- Output $w := w^{(s)}$.

After the first step of the algorithm, we have $w^{(2)} = y_i x^{(i)}$, so we know that $\langle w^{(2)}, x^{(i)} \rangle = y_i$. That is, after one step of the algorithm, there is some $1 \leq i \leq k$ that is classified correctly. Unfortunately, after obtaining $w^{(3)}$ in the next step of the algorithm, if $x^{(i)}$ was the vector obtained from the first step of the algorithm, then $\langle w^{(3)}, x^{(i)} \rangle$ and y_i may not have the same sign. That is, the second step of the algorithm could mis-classify the point we obtained in the first step. So, a priori, Algorithm 2.8 may never terminate. Or at very least, the Algorithm may take a very large number of steps until it terminates. Either situation would be quite bad.

Note that the algorithm can be updated as it receives new data. For this reason, Algorithm 2.8 is called an **online learning** algorithm. However, the algorithm must keep all vectors $x^{(1)}, \dots, x^{(k)}$ in memory, so it is not a streaming algorithm.

Theorem 2.9 (Perceptron Algorithm Run Time). Let $x^{(1)}, \dots, x^{(k)} \in \mathbb{R}^n$ and let $y_1, \dots, y_k \in \{-1, 1\}$ be given. Assume that there exists $w \in \mathbb{R}^n$ such that

$$\text{sign}(\langle w, x^{(i)} \rangle) = y_i, \quad \forall 1 \leq i \leq k.$$

Define

$$\beta := \max_{i=1, \dots, k} \|x^{(i)}\|, \quad \theta := \min \left\{ \|w\| : \forall 1 \leq i \leq k, y_i \langle w, x^{(i)} \rangle \geq 1 \right\}.$$

Let $w \in \mathbb{R}^n$ achieve the minimum value in the definition of θ . Then the Perceptron Algorithm 2.8 terminates with a value of s satisfying

$$s \leq (\beta\theta)^2.$$

Proof. We first show that the weight vector $w^{(s)}$ “moves towards” w at each iteration of the algorithm:

$$\langle w^{(s+1)}, w \rangle \geq \langle w^{(s)}, w \rangle + 1, \quad \forall s \geq 1. \quad (*)$$

To see this, let $1 \leq i \leq k$ such that $x^{(i)}$ is selected during step two in iteration s of the algorithm. By definition of w , note that $y_i \langle w, x^{(i)} \rangle \geq 1$. And by definition of $w^{(s+1)}$ and w ,

$$\langle w^{(s+1)}, w \rangle = \langle w^{(s)}, w \rangle + y_i \langle x^{(i)}, w \rangle \geq \langle w^{(s)}, w \rangle + 1.$$

We now show that the weight vector $w^{(s)}$ does not increase too much in length at each iteration:

$$\|w^{(s+1)}\|^2 \leq \|w^{(s)}\|^2 + \beta^2, \quad \forall s \geq 1. \quad (**)$$

Since $1 \leq i \leq k$ was selected in step two of the algorithm, we have $y_i \neq \text{sign}(\langle w^{(s)}, x^{(i)} \rangle)$, i.e. $y_i \langle w^{(s)}, x^{(i)} \rangle < 0$. Using this and the definition of $w^{(s+1)}$,

$$\begin{aligned} \|w^{(s+1)}\|^2 &= \|w^{(s)} + y_i x^{(i)}\|^2 = \|w^{(s)}\|^2 + 2\langle w^{(s)}, y_i x^{(i)} \rangle + \|x^{(i)}\|^2 \\ &\leq \|w^{(s)}\|^2 + \|x^{(i)}\|^2 \leq \|w^{(s)}\|^2 + \beta^2. \end{aligned}$$

We now conclude the proof. Let $s \geq 1$. Induction and (*) imply that $\langle w^{(s)}, w \rangle \geq s$. Induction and (**) imply that $\|w^{(s)}\|^2 \leq s\beta^2$. By definition of w , $\|w\| = \theta$. So, the cosine of the angle between w and $w^{(s)}$ satisfies

$$1 \geq \frac{\langle w^{(s)}, w \rangle}{\|w^{(s)}\| \|w\|} \geq \frac{s}{\theta \sqrt{s} \beta} = \frac{\sqrt{s}}{\theta \beta}.$$

That is, $\sqrt{s} \leq \beta \theta$. □

Remark 2.10. Instead of going through all points in Algorithm 2.8 until the separating hyperplane is found, one could simply run through all points $x^{(1)}, \dots, x^{(k)}$ in order. Doing so would speed up the algorithm, but the resulting weight vector w may not separate the points correctly. One might call this effect **underfitting**.

When θ is large, there are some data vectors $x^{(1)}, \dots, x^{(k)}$ close to the separating hyperplane $\{x \in \mathbb{R}^n : \langle w, x \rangle = 0\}$. More specifically,

Lemma 2.11. *Let $x^{(1)}, \dots, x^{(k)} \in \mathbb{R}^n$ and let $y_1, \dots, y_k \in \{-1, 1\}$ be given. Assume that there exists $w \in \mathbb{R}^n$ such that*

$$\text{sign}(\langle w, x^{(i)} \rangle) = y_i, \quad \forall 1 \leq i \leq k.$$

Let $w \in \mathbb{R}^n$ achieve the minimum in $\min \{ \|w\| : \forall 1 \leq i \leq k, y_i \langle w, x^{(i)} \rangle \geq 1 \}$. Then $w / \|w\|$ achieves the maximum

$$\max_{v \in \mathbb{R}^n, \|v\|=1} \min_{i=1, \dots, k} y_i \langle v, x^{(i)} \rangle. \quad (*)$$

Proof. Let v achieve the maximum value $a > 0$ in (*). Then $y_i \langle (v/a), x^{(i)} \rangle \geq 1$ by (*). So, by definition of w , $\|w\| \leq \|v/a\| = 1/a$. If $1 \leq i \leq k$, by definition of w ,

$$y_i \langle (w / \|w\|), x^{(i)} \rangle = \frac{1}{\|w\|} y_i \langle w, x^{(i)} \rangle \geq \frac{1}{\|w\|} \geq a.$$

So, w also achieves the maximum value in (*). (Consequently, when we take $\min_{i=1, \dots, k}$, all inequalities on the previous line must be equalities, so that $a = 1 / \|w\|$.) □

When $\|v\| = 1$, $y_i \langle v, x^{(i)} \rangle$ is the distance from $x^{(i)}$ to the hyperplane $\{x \in \mathbb{R}^n : \langle x, v \rangle = 0\}$. So, in Theorem 2.9, the quantity

$$\theta := \min \left\{ \|w\| : \forall 1 \leq i \leq k \ y_i \langle w, x^{(i)} \rangle \geq 1 \right\}.$$

produces a vector w whose perpendicular hyperplane has the largest uniform distance a to the vectors $x^{(1)}, \dots, x^{(k)}$. Since $a = 1/\|w\| = 1/\theta$, the quantity $1/\theta$ is sometimes called the **margin** of the vectors $x^{(1)}, \dots, x^{(k)}$, as it represents the “widest” symmetric slab through the origin that can fit between all of the vectors.

2.1.1. *Variants of Learning Linear Threshold Functions.* In the case that the points *cannot* be separated by a hyperplane, suppose we try to find a hyperplane that correctly classifies the largest number of points. Unfortunately, this problem is known to be NP-hard [JP78] [ABSS97, Theorem 4]. In the literature, this problem is called the Open Hemisphere Problem, or Densest Hemisphere Problem.

Also, given that there is some half space that correctly classifies 99% of data points, it is NP-hard to find a half space that correctly classifies 51% of data points [FGKP06]. Note that correctly classifying 50% of data points is easy by just choosing any half space and deciding which side to label +1. In fact, if $\mathbf{NP} \subsetneq \mathbf{TIME}(2^{(\log n)^{o(1)}})$, no polynomial time algorithm can distinguish between the following cases for learning half spaces with k input points: (i) $1 - 2^{-\Omega(\sqrt{\log k})}$ fraction of points can be correctly classified by some halfspace, or (ii) no more than $(1/2) + 2^{-\Omega(\sqrt{\log k})}$ fraction of points can be correctly classified by any halfspace [FGKP06, Theorem 4].

On the other hand, variants of the Perceptron algorithm do still learn linear threshold functions efficiently when we start with a separating hyperplane and then change each label independently with fixed probability less than 1/2 [BFKV98]. Under these assumptions, efficient learning can then occur in the PAC model (discussed in Section 3) and in the Statistical Query Model. Also, when a small fraction of labels can be arbitrarily corrupted (the so-called agnostic learning model) and some assumptions are made on allowable random samples of inputs, linear threshold functions can be learned efficiently [KKMS08]. With access to Gaussian samples, where a fraction of both samples and labels can be arbitrarily corrupted (the so-called nasty noise model), linear threshold functions can be learned efficiently [DKS18].

Instead of returning the final vector w that is output by the Perceptron algorithm, practitioners use an average of all weight vectors w over all time steps. Supposedly this works better than the Perceptron itself, and it is called the averaged perceptron.

Logistic regression is also said to work well in practice for classification tasks.

Remark 2.12 (Linear Programming). Let $x^{(1)}, \dots, x^{(k)} \in \{-1, 1\}^n$ and let $y_1, \dots, y_k \in \{-1, 1\}$ be given. Let A be the $k \times n$ matrix whose i^{th} row is $x^{(i)}$. We can rewrite the goal of Problem 2.4 (with $t = 0$) as finding $w \in \mathbb{R}^n$

$$\langle w, x^{(i)} \rangle y_i > 0, \quad \forall 1 \leq i \leq k.$$

Since the $y_1, \dots, y_k \in \{-1, 1\}$ are fixed, this is a set of linear inequalities in w , i.e. finding the existence of such a w is a linear program in n dimensions with k constraints. It is conceivable to use a linear programming algorithm (such as the ellipsoid method or interior point method) to solve Problem 2.4, however it would most likely be slower than the Perceptron Algorithm.

On the other hand, the run time of the linear program would not depend on the geometry of the points $x^{(1)}, \dots, x^{(k)}$, while our run time guarantee in Theorem 2.9 does depend on the geometry of the points (and this worst-case guarantee can increase exponentially with n).

Exercise 2.13. Let n be a positive integer. Let c_n be the number of boolean functions $f: \{-1, 1\}^n \rightarrow \{-1, 1\}$ that are linear threshold functions. This quantity is of interest since it roughly quantifies the “expressive power” of linear threshold functions for the supervised learning problem. It is known that

$$c_n = 2^{n^2(1+o(1))}$$

So, Problem 2.4 asks for the linear threshold function that fits the given data among a family of functions of super-exponential size. For another perspective on the “expressive power” of linear threshold functions, we will look into the VC-dimension in e.g. Proposition 4.11.

Using an inductive argument prove the weaker lower bound

$$c_n \geq 2^{n(n-1)/2}.$$

(Hint: induct on n . If $f: \{-1, 1\}^n \rightarrow \{-1, 1\}$, consider $\bar{f}: \{-1, 1\}^{n+1} \rightarrow \{-1, 1\}$ defined (partially for now) so that $\bar{f}(x_1, \dots, x_n, -1) := f(x_1, \dots, x_n)$ for all $(x_1, \dots, x_n) \in \{-1, 1\}^n$. How many ways can we define \bar{f} on the remaining “half” of the hypercube $\{-1, 1\}^{n+1}$ such that \bar{f} is a linear threshold function?)

As we will discuss in Section 8, it is of interest to state the general learning problem for compositions of linear threshold functions (i.e. neural networks). In this case, asymptotics for the number of such functions were recently found in <https://arxiv.org/pdf/1901.00434.pdf>.

Exercise 2.14. Let $a > 0$. Let $X^{(1)}, \dots, X^{(k)} \in \mathbb{R}^n$ be independent identically distributed samples from a Gaussian random vector with mean $(a, 0, \dots, 0)$ and identity covariance matrix). Let $X^{(k+1)}, X^{(k+2)}, \dots, X^{(2k)} \in \mathbb{R}^n$ be independent identically distributed samples from a Gaussian random vector with mean $(-a, 0, \dots, 0)$, where $a > 0$ is known. As in our analysis of the perceptron algorithm, define

$$\mathcal{B} := \max_{i=1, \dots, 2k} \|X^{(i)}\|$$

$$\Theta := \max \left\{ \|w\| : \forall 1 \leq i \leq 2k \ y_i \langle w, X^{(i)} \rangle \geq 1 \right\}.$$

Give some reasonable estimates for $\mathbf{E}\mathcal{B}$ and $\mathbf{E}\Theta$ as a function of a .

2.2. Embeddings and the “Kernel Trick”. In Remark 2.6, we noticed that exactly learning a linear threshold function amounts to separating data points into two sets by a separating hyperplane. However, some natural data sets may be sortable into two categories that cannot be separated by a hyperplane. For this reason, practitioners using the Perceptron Algorithm often “preprocess” their data so that it is sensibly separable by a hyperplane. The preprocessing can be described as a function $\phi: A \rightarrow C$, where A, C are sets. (Recall that we phrased Problem 2.1 as determining an unknown function $f: A \rightarrow B$.) The function $\phi: A \rightarrow C$ is sometimes called a **feature map**.

Problem 2.15 (Supervised Learning Problem for Linear Threshold Functions, with Kernel Trick). Let $f: \{-1, 1\}^n \rightarrow \{-1, 1\}$ be an unknown linear threshold function. Let $x^{(1)}, \dots, x^{(k)} \in \{-1, 1\}^n$ and let $y_1, \dots, y_k \in \{-1, 1\}$. It is known that

$$f(x^{(i)}) = y_i, \quad \forall 1 \leq i \leq k.$$

The goal of the problem is to find an inner product space C with inner product $\langle \cdot, \cdot \rangle_C$, an embedding $\phi: \{-1, 1\}^n \rightarrow C$, $t \in \mathbb{R}$ and $w \in C$ such that

$$\text{sign}(\langle w, \phi(x^{(i)}) \rangle_C - t) = y_i, \quad \forall 1 \leq i \leq k.$$

As above, we assume that Problem 2.15 can be solved when we apply the following modified Perceptron Algorithm.

Algorithm 2.16 (Kernel Perceptron Algorithm, Version 1).

- Define $w^{(1)} := 0 \in C$ and let $s := 1$
- If there exists some $1 \leq i \leq k$ such that $y_i \neq \text{sign}(\langle w^{(s)}, \phi(x^{(i)}) \rangle)$, i.e. a misclassification occurs, define

$$w^{(s+1)} := w^{(s)} + y_i \phi(x^{(i)}).$$

- Increase the value of s by one. Repeat the previous step until no such i exists.
- Output $w := w^{(s)}$.

Sometimes the embedding ϕ may be difficult to compute or write explicitly. In such cases, it might be desirable to only define ϕ implicitly. In such a case, for all $1 \leq i \leq k$ we only need to define $\phi(x^{(i)})$ to be an element of some inner product space. And if we can rewrite the algorithm so that it only uses the values of the inner products, $\langle \phi(x^{(i)}), \phi(x^{(j)}) \rangle_C \forall 1 \leq i, j \leq k$, then we need only specify the values of these inner products. We then use the following equivalence.

Exercise 2.17. Let M be a $k \times k$ real symmetric matrix. Then M is positive semidefinite if and only if there exists a real $k \times k$ matrix R such that

$$M = RR^T.$$

In either case, if $r^{(i)}$ denotes the i^{th} row of R , we have

$$m_{ij} = \langle r^{(i)}, r^{(j)} \rangle, \quad \forall 1 \leq i, j \leq k.$$

That is, the values of the inner products $\langle \phi(x^{(i)}), \phi(x^{(j)}) \rangle_C$ can be equivalently specified as the values m_{ij} of a real symmetric positive semidefinite matrix. More generally, we have the following infinite-dimensional result from 1909.

Theorem 2.18 (Mercer's Theorem). *Let μ be a Borel measure on \mathbb{R}^n such that the measure of any open set in \mathbb{R}^n is positive. We denote $L_2(\mu) := \{f: \mathbb{R}^n \rightarrow \mathbb{R} : \int_{\mathbb{R}^n} |f(x)|^2 d\mu(x) < \infty\}$, and we equip $L_2(\mu)$ with the standard inner product $\int_{\mathbb{R}^n} f(x)g(x)d\mu(x)$ defined for any $f, g \in L_2(\mu)$. Let $m: \mathbb{R}^n \times \mathbb{R}^n \rightarrow [0, \infty)$ be a continuous symmetric function ($m(x, y) = m(y, x)$ for all $x, y \in \mathbb{R}^n$) such that, for all $p \geq 1$, for all $z^{(1)}, \dots, z^{(p)} \in \mathbb{R}^n$, for all $\beta_1, \dots, \beta_p \in \mathbb{R}$ we have the **positive semi-definiteness condition***

$$\sum_{i,j=1}^p \beta_i \beta_j m(z^{(i)}, z^{(j)}) \geq 0, \quad \int_{\mathbb{R}^n} \int_{\mathbb{R}^n} |m(x, y)|^2 d\mu(x)d\mu(y) < \infty.$$

Then there exists an orthonormal basis $\{\psi_i\}_{i=1}^{\infty}$ of $L_2(\mu)$ and a sequence of nonnegative real numbers $\{\lambda_i\}_{i=1}^{\infty}$ such that m is equal to the following series, which converges absolutely pointwise:

$$m(x, y) = \sum_{i=1}^{\infty} \lambda_i \psi_i(x) \psi_i(y), \quad \forall x, y \in \mathbb{R}^n.$$

Exercise 2.19. Let μ be a Borel measure on \mathbb{R}^n such that the measure of any open set in \mathbb{R}^n is positive. Let $m: \mathbb{R}^n \times \mathbb{R}^n \rightarrow \mathbb{R}$ be continuous with $\int_{\mathbb{R}^n} \int_{\mathbb{R}^n} |m(x, y)|^2 d\mu(x) d\mu(y) < \infty$. Show that the following two positive semidefinite conditions on m are equivalent:

- $\forall p \geq 1$, for all $z^{(1)}, \dots, z^{(p)} \in \mathbb{R}^n$, for all $\beta_1, \dots, \beta_p \in \mathbb{R}$ we have

$$\sum_{i,j=1}^p \beta_i \beta_j m(z^{(i)}, z^{(j)}) \geq 0.$$

- $\forall f \in L_2(\mu)$, we have

$$\int_{\mathbb{R}^n} \int_{\mathbb{R}^n} f(x) f(y) m(x, y) d\mu(x) d\mu(y) \geq 0.$$

From either condition, we should see that the converse of Mercer's Theorem holds. We should also be able to deduce various properties of positive semidefinite (PSD) kernels. For example, a nonnegative linear combination of PSD kernels is PSD.

So, define $\ell_2 := \{(\beta_i)_{i=1}^\infty : \sum_{i=1}^\infty \beta_i^2 < \infty\}$ with the standard inner product $\langle (\beta_i)_{i=1}^\infty, (\gamma_i)_{i=1}^\infty \rangle := \sum_{i=1}^\infty \beta_i \gamma_i$ and $\phi: \mathbb{R}^n \rightarrow \ell_2$ by

$$\phi(x) := \{\sqrt{\lambda_i} \psi_i(x)\}_{i=1}^\infty, \quad \forall x \in \mathbb{R}^n,$$

then

$$\langle \phi(x), \phi(y) \rangle = \sum_{i=1}^\infty \lambda_i \psi_i(x) \psi_i(y) = m(x, y), \quad \forall x, y \in \mathbb{R}^n.$$

That is, Theorem 2.18 is an infinite-dimensional generalization of Exercise 2.17. Note that each ψ_i can be nonlinear, and in general the ϕ maps we use will be nonlinear maps.

Recall also that the Spectral Theorem 1.12 for real symmetric matrices $k \times k$ matrices M says there exists a real diagonal D and an orthogonal Q such that

$$M = Q^T D Q$$

We can write this in vector form to more closely match Theorem 2.18. For any $1 \leq p \leq k$, let λ_p denote the p^{th} diagonal entry of D and let $\psi^{(p)} \in \mathbb{R}^k$ denote the k^{th} row of Q . Then

$$m_{ij} = \sum_{p=1}^k \lambda_p \psi_i^{(p)} \psi_j^{(p)}, \quad \forall 1 \leq i, j \leq k.$$

The function $m(x, y)$ (or the matrix m_{ij}) is called a **kernel**.

Also, since all countable Hilbert spaces are isometric, we can and will assume that $C = \ell_2$ with the standard inner product, which we denote as $\langle \cdot, \cdot \rangle$. In order to rewrite Algorithm 2.16, we need to express the algorithm using only $\langle \phi(x^{(i)}), \phi(x^{(j)}) \rangle$ terms. To this end, for all $1 \leq i \leq k$, let α_i be the number of times that $i \in \{1, \dots, k\}$ is selected in step two of the algorithm. That is, α_i is the number of times that $x^{(i)}$ is “mis-classified” within the algorithm.

Let $s \geq 1$. Suppose at step s of the algorithm, $\tilde{x}^{(s)} \in \{x^{(1)}, \dots, x^{(k)}\}$ is selected. For any $s \geq 1$, induction implies that $w^{(s+1)}$ is a linear combination of $\phi(\tilde{x}^{(1)}), \dots, \phi(\tilde{x}^{(s)})$ in Algorithm 2.16:

$$w^{(s+1)} = \sum_{j=1}^s y_j \phi(\tilde{x}^{(j)}). \quad (*)$$

So, at the next step of the algorithm, we can compute

$$\text{sign}(\langle w^{(s+1)}, \phi(\tilde{x}^{(s+1)}) \rangle) = \text{sign}\left(\sum_{j=1}^s y_j \langle \phi(\tilde{x}^{(j)}), \phi(\tilde{x}^{(s)}) \rangle\right) =: \text{sign}\left(\sum_{j=1}^k y_j m(\tilde{x}^{(j)}, \tilde{x}^{(i)})\right).$$

Suppose the algorithm terminates after s steps. As mentioned above, for all $1 \leq j \leq k$, let α_j be the number of times that $\phi(x^{(j)})$ appears in the sum (*). Then we can rewrite (*) as

$$w := w^{(s+1)} = \sum_{j=1}^k \alpha_j y_j \phi(x^{(j)}).$$

We then plug in any x in the span of $x^{(1)}, \dots, x^{(k)}$ so that

$$\text{sign}(\langle w^{(k+1)}, \phi(x) \rangle) := \text{sign}\left(\sum_{j=1}^k \alpha_j y_j m(x^{(j)}, x)\right).$$

In summary, we can rewrite Algorithm 2.16 in the following way.

Algorithm 2.20 (Kernel Perceptron Algorithm, Version 2).

Let $m: \mathbb{R}^n \times \mathbb{R}^n \rightarrow [0, \infty)$ be a positive-definite function (kernel) satisfying the assumptions of Mercer's Theorem 2.18.

- Define $\alpha_1 := 0, \dots, \alpha_k := 0$ and let $s := 1$.
- If there exists some $1 \leq i \leq k$ such that $y_i \neq \text{sign}\left(\sum_{j=1}^k \alpha_j y_j m(\tilde{x}^{(j)}, \tilde{x}^{(i)})\right)$, i.e. a mis-classification occurs, then increase the value of α_i by one.
- Increase the value of s by one. Repeat the previous step until no such i exists.
- Output the function $f(x) := \text{sign}\left(\sum_{j=1}^k \alpha_j y_j m(x^{(j)}, x)\right)$, valid for any $x \in \mathbb{R}^n$.

For some rigorous guarantees of a “kernel trick” in the context of a clustering-type problem, see [LOGT12].

Exercise 2.21. For each kernel function $m: \mathbb{R}^n \times \mathbb{R}^n \rightarrow [0, \infty)$ below, find an inner product space C and a map $\phi: \mathbb{R}^n \rightarrow C$ such that

$$m(x, y) = \langle \phi(x), \phi(y) \rangle_C, \quad \forall x, y \in \mathbb{R}^n.$$

Conclude that each such m is a positive semidefinite function, in the sense stated in Mercer's Theorem.

- $m(x, y) := 1 + \langle x, y \rangle \forall x, y \in \mathbb{R}^n$.
- $m(x, y) := (1 + \langle x, y \rangle)^d \forall x, y \in \mathbb{R}^n$, where d is a fixed positive integer.
- $m(x, y) := \exp(-\|x - y\|^2)$.

Hint: it might be helpful to consider d -fold iterated tensor products of the form $x^{\otimes d} = x \otimes x \otimes \dots \otimes x$, along with their corresponding inner products.

2.3. Optional: Proof of Mercer's Theorem.

Proof of Mercer's Theorem 2.18. Consider the operator on $L_2(\mu)$ defined by

$$T(f)(x) := \int_{\mathbb{R}^n} f(y) m(x, y) d\mu(y), \quad \forall x \in \mathbb{R}^n, \forall f \in L_2(\mu)$$

By the Cauchy-Schwarz inequality and by assumption on m , if $f \in L_2(\mu)$, then $T(f) \in L_2(\mu)$, so that $T: L_2(\mu) \rightarrow L_2(\mu)$, with $\|T\|^2 \leq \int_{\mathbb{R}^n} |m(x, y)|^2 d\mu(x) d\mu(y)$, where $\|T\|$ denotes the

operator norm. We will show that T is a compact operator, i.e. $T(B(0,1))$ has compact closure in $L_2(\mu)$ with respect to the norm topology, where $B(0,1) := \{f \in L_2(\mu) : \|f\| \leq 1\}$, and $\|f\| := (\int_{\mathbb{R}^n} |f(x)|^2 d\mu(x))^{1/2}$. When $f, g \in L_2(\mu)$, we denote $f \otimes g: \mathbb{R}^n \times \mathbb{R}^n \rightarrow \mathbb{R}$ as the function $(f \otimes g)(x, y) := f(x)g(y)$. Let $\{\phi_i\}_{i=1}^\infty$ be an orthonormal basis of $L_2(\mu)$. Then $\{\phi_i \otimes \phi_j\}_{i,j=1}^\infty$ is an orthonormal basis of $L_2(\mu \times \mu)$. Since $m \in L_2(\mu \times \mu)$, m can be expressed as a linear combination of this orthonormal basis, in the sense that

$$\lim_{p \rightarrow \infty} \left\| m - \sum_{i+j \leq p} \langle m, \phi_i \otimes \phi_j \rangle \phi_i \otimes \phi_j \right\| = 0. \quad (*)$$

So, for any $p \geq 1$, define the following operator on $L_2(\mu)$

$$T_p(f)(x) := \int_{\mathbb{R}^n} f(y) \left(\sum_{i+j \leq p} \langle m, \phi_i \otimes \phi_j \rangle (\phi_i \otimes \phi_j)(x, y) \right) d\mu(y), \quad \forall x \in \mathbb{R}^n, \forall f \in L_2(\mu).$$

By (*), T is the limit of T_p in the sense that

$$\lim_{p \rightarrow \infty} \sup_{f \in L_2(\mu) : \|f\| \leq 1} \|(T(f) - T_p(f))\| = 0. \quad (**)$$

Also, by its definition, each T_p has finite-dimensional range, hence it is a compact operator. The operator T is compact if: given any norm-bounded sequence $\{f_i\}_{i=1}^\infty \in L_2(\mu)$, the sequence $\{Tf_i\}_{i=1}^\infty$ has a convergent subsequence. So, (**) implies that T itself is compact by the following diagonalization argument. Since T_1 is compact, there is a subsequence $\{f_{i,1}\}_{i=1}^\infty$ of $\{f_i\}_{i=1}^\infty$ such that $\{T_1 f_{i,1}\}_{i=1}^\infty$ converges. Since T_2 is compact, there is a subsequence $\{f_{i,2}\}_{i=1}^\infty$ of $\{f_{i,1}\}_{i=1}^\infty$ such that $\{T_2 f_{i,2}\}_{i=1}^\infty$ converges. And so on. Then let $g_i := f_{i,i}$ for all $i \geq 1$. Fix $i, j \geq 1$ and then let $p \geq 1$.

$$\|T(g_i) - T(g_j)\| \leq \|T(g_i) - T_p(g_i)\| + \|T_p(g_i) - T_p(g_j)\| + \|T_p(g_j) - T(g_j)\|.$$

The first and third terms become small when p is large by (**). The middle term can be made small by choosing p to be the minimum of i and j . In conclusion for all $\varepsilon > 0$, there exists $q \geq 1$ such that, for all $i, j \geq q$, we have $\|T(g_i) - T(g_j)\| < \varepsilon$. That is, $\{T(g_i)\}_{i=1}^\infty$ is a Cauchy sequence, hence convergent.

So, T itself is a compact operator, and it is positive semidefinite by Exercise 2.19. The Riesz Theory of Compact operators then applies (Theorem 10.38(4),(6).) That is, there exists a sequence of nonnegative eigenvalues $\{\lambda_i\}_{i=1}^\infty$ and an orthonormal system of eigenfunctions $\{\psi_i\}_{i=1}^\infty$ such that

$$Tf = \sum_{i=1}^{\infty} \lambda_i \langle f, \psi_i \rangle \psi_i, \quad \forall f \in L_2(\mu).$$

We now note that m is nonnegative on its diagonal. We show this by contradiction. Suppose $m(x, x) < 0$ for some $x \in \mathbb{R}^n$. Since m is continuous, there exists an open neighborhood U of x such that m is negative on $U \times U$. By assumption, $\mu(U) > 0$. So, if $f := 1_U$, we would have

$$0 \leq \langle Tf, f \rangle = \int_U \int_U m(x, y) d\mu(x) d\mu(y) < 0,$$

a contradiction. For any $p \geq 1$, define

$$S_p(f) := \sum_{i=1}^p \lambda_i \langle f, \psi_i \rangle \psi_i \quad \forall f \in L_2(\mu).$$

And define the corresponding kernel

$$m_p(x, y) := \sum_{i=1}^p \lambda_i \psi_i(x) \psi_i(y), \quad \forall x, y \in \mathbb{R}^n.$$

Then $T - S_p$ is also a positive semidefinite operator, so the corresponding kernel $m - m_p$ also has a nonnegative diagonal. That is, for all $p \geq 1$,

$$\sum_{i=1}^p \lambda_i (\psi_i(x))^2 \leq m(x, x).$$

The series on the left is therefore absolutely convergent as $p \rightarrow \infty$. By the Cauchy-Schwarz inequality, the series defining m_p is also absolutely convergent as $p \rightarrow \infty$, for all $x, y \in \mathbb{R}^n$. Denote

$$\bar{m}(x, y) := \lim_{p \rightarrow \infty} m_p(x, y) = \sum_{i=1}^{\infty} \lambda_i \psi_i(x) \psi_i(y), \quad \forall x, y \in \mathbb{R}^n.$$

It remains to show that $\bar{m} = m$. To see this, note that the operator with kernel \bar{m} has the same eigenvalues and eigenfunctions as T . Therefore T has kernel \bar{m} , as desired. \square

3. PROBABLY APPROXIMATELY CORRECT (PAC) LEARNING

Learning involves data from the real world, and real world data is noisy. To account for this, Valiant introduced the PAC learning model in 1984.

Problem 3.1 (Supervised PAC Learning Problem). Let A, B be sets. Let $f: A \rightarrow B$ be an unknown function. The goal of the learning problem is to determine the function f on all of A using a small number of *randomly drawn* values of f on A . Let $0 < \varepsilon, \delta < 1/2$ and let \mathbf{P} be a probability law on A . We are given access to ε, δ and to random samples from \mathbf{P} . (The distribution \mathbf{P} may or may not be known, depending on the problem at hand.) If $x \in A$ is a random sample from \mathbf{P} , the pair $(x, f(x))$ is also known.

The goal of the problem is the following. With probability at least $1 - \delta$ (with respect to randomly drawn samples, drawn according to \mathbf{P}), output a function $g: A \rightarrow B$ such that

$$\mathbf{P}(f(x) \neq g(x)) < \varepsilon.$$

That is, the hypothesis g is probably (with probability at least $1 - \delta$) approximately (up to error ε) correct.

We say a subset of functions $\mathcal{F} \subseteq \{f: A \rightarrow B\}$ is **PAC Learnable** if there exists an algorithm such that, for any probability distribution \mathbf{P} on A , for any $0 < \varepsilon, \delta < 1/2$, the algorithm achieves the above goal.

In the case $A = \{-1, 1\}^n$ and $B = \{-1, 1\}$, we say that \mathcal{F} is **efficiently PAC Learnable** if \mathcal{F} is PAC Learnable, and the algorithm has a run time that is polynomial in $n, 1/\varepsilon, 1/\delta$. When every $f \in \mathcal{F}$ has a positive integer “size” associated to it, we require the algorithm to have run time polynomial in $n, 1/\varepsilon, 1/\delta$ and in the size of the f , where $f \in \mathcal{F}$ is the unknown function being learned by the algorithm.

As discussed in Section 2, Problem 3.1 is too general to be tractable for arbitrary \mathcal{F} . In the literature, the family of functions \mathcal{F} is often called a **concept class**.

We include the number of samples of \mathbf{P} in the run time of a PAC learning algorithm. So, being efficiently PAC Learnable requires not taking too many samples from \mathbf{P} .

One criticism of PAC learning is that it cannot account for adversarial noise. That is, samples are assumed to independent and identically distributed, so any noise in the data is modeled as i.i.d. noise. On the other hand, one could imagine an adversary corrupts some fraction of the samples arbitrarily. Such noise be outside the scope of the PAC learning model. For this reason, some contemporary investigations in machine learning have investigated more general adversarial noise models.

3.1. Learning Boolean Conjunctions. Our first example of a class of PAC learnable functions is boolean conjunctions. For convenience, we use $\{0, 1\}$ instead of $\{-1, 1\}$ in this subsection.

Definition 3.2 (Boolean Conjunctions). Let $I, J \subseteq \{1, \dots, n\}$. A **boolean conjunction** is a function $f: \{0, 1\}^n \rightarrow \{0, 1\}$ of the form

$$f(x_1, \dots, x_n) = \prod_{i \in I} x_i \prod_{j \in J} (1 - x_j), \quad \forall (x_1, \dots, x_n) \in \{0, 1\}^n.$$

Example 3.3. The function f where $f = 1$ only when $x_1 = 1, x_3 = 0$ and $x_4 = 1$ can be written as

$$f(x_1, \dots, x_n) := x_1(1 - x_3)x_4, \quad \forall (x_1, \dots, x_n) \in \{-1, 1\}^n.$$

When f is a boolean conjunction, the set $\{x \in \{0, 1\}^n: f(x) = 1\}$ can be understood geometrically as an intersection of coordinate halfspaces. Alternatively, if we think of 0 as logical “false” and 1 as “true,” a conjunction f is then a combination of logical AND operations on a set of variables together with their negations.

Theorem 3.4. *The class of boolean conjunctions is efficiently PAC learnable.*

Proof. Let f be an unknown boolean conjunction of the form $f(x_1, \dots, x_n) = \prod_{i \in I} x_i \prod_{j \in J} (1 - x_j)$, for some $I, J \subseteq \{1, \dots, n\}$. We first “guess” a hypothesis g where $I = J = \{1, \dots, n\}$, so that

$$g(x_1, \dots, x_n) = \prod_{i=1}^n x_i(1 - x_i).$$

Initially, $g = 0$ on $\{-1, 1\}^n$. Suppose we sample $y = (y_1, \dots, y_n) \in \{-1, 1\}^n$ from \mathbf{P} , and we find that $f(y) = 1$. In such a case, we update g to agree with f . That is, for each $1 \leq i \leq n$, if $y_i = 0$, delete the x_i term from g , and if $y_i = 1$, delete the $(1 - x_i)$ term from g . We repeat this procedure. For any $k \geq 1$, let g_k denote the hypothesis g_k after k iterations of the algorithm, and let $g_0 := g$. Note that each deleted term does not appear in the formula for f (e.g. if $y_i = 0$ and $f(y) = 1$, then y_i cannot appear in the formula of f). So, for any $k \geq 1$, any term in f will be contained in g_k . Let us then try to find the probability that there is a term in g_k not appearing in f .

Let $1 \leq i \leq n$, and consider a term of the form $z := x_i$ or $z := 1 - x_i$. Suppose z appears in g_k but not in f . The term z will be deleted from g_k only when we have a sample where

$z = 0$ and $f = 1$. So, the probability that term z is deleted from g_k in one iteration of the algorithm is

$$p(z) := \mathbf{P}(f(y) = 1 \text{ and } z = 0).$$

For any $k \geq 1$, define

$$\varepsilon_k := \mathbf{P}(f(y) \neq g_k(y)).$$

Since $f \neq g_k$ only when $f = 1$, the union bound gives

$$\begin{aligned} \varepsilon_k &= \mathbf{P}(f(y) = 1 \text{ and } g_k(y) = 0) \\ &= \mathbf{P}(f(y) = 1 \text{ and some term in } g_k \text{ is zero}) \leq \sum_{\text{terms } z \text{ in } g_k} p(z). \quad (*) \end{aligned}$$

Let $\varepsilon > 0$. Suppose $p(z) < \varepsilon/(2n)$ for all terms z in g_k . Since g_k has at most $2n$ terms, we would then have $\varepsilon_k \leq 2n\varepsilon/(2n) = \varepsilon$. Let C_k be the event that \exists a term z in g_k with $p(z) > \varepsilon/(2n)$ and such that z does not appear in f . It remains to bound $\mathbf{P}(C_k)$.

Let z be a term in g_k that is not in f . The probability that z appears in g_k is at most $(1 - p(z))^k$, since each iteration of the algorithm gives us an independent sample from \mathbf{P} . Then by the union bound (since at most $2n$ terms appear in g_k)

$$\mathbf{P}(C_k) \leq 2n \left(1 - \frac{\varepsilon}{2n}\right)^k. \quad (**)$$

So, if $\delta > 0$, we choose k such that

$$2n \left(1 - \frac{\varepsilon}{2n}\right)^k \leq \delta.$$

Using $1 - t \leq e^{-t}$ for all $t \in [0, 1]$, we choose k so that $2ne^{-\varepsilon k/(2n)} \leq \delta$, i.e. $-\varepsilon k \leq 2n \log(\delta/(2n))$, i.e.

$$k \geq \frac{2n}{\varepsilon} \log(2n/\delta).$$

In summary, if the algorithm uses more than $\frac{2n}{\varepsilon} \log(2n/\delta)$ samples, then with probability at least $1 - \delta$ by (**), the hypothesis g_k of the algorithm satisfies $\mathbf{P}(f \neq g_k) \leq \varepsilon$ by (*). \square

Exercise 3.5. Show that the set of conjunctions is contained in the set of linear threshold functions. That is, given a boolean conjunction $f: \{0, 1\}^n \rightarrow \{0, 1\}$, find $w \in \mathbb{R}^n, t \in \mathbb{R}$ such that

$$f(x) = \text{sign}(\langle w, x \rangle - t), \quad \forall x = (x_1, \dots, x_n) \in \{0, 1\}^n.$$

3.2. Learning DNF Formulas. Surprisingly, the following slightly larger class of natural functions that contains boolean conjunctions is not efficiently PAC learnable, in its own function class.

Definition 3.6 (3-Term DNF Formula). Let $f_1, f_2, f_3: \{0, 1\}^n \rightarrow \{0, 1\}$ be boolean conjunctions. A **3-Term DNF Formula** is a function $f: \{0, 1\}^n \rightarrow \{0, 1\}$ of the form

$$f(x) = \max(f_1(x), f_2(x), f_3(x)), \quad \forall x \in \{-1, 1\}^n.$$

When f is a 3-term DNF formula, the set $\{x \in \{0, 1\}^n: f(x) = 1\}$ can be understood geometrically as a union of 3 conjunctions (which are themselves intersections of coordinate half spaces). Alternatively, if we think of 0 as logical “false” and 1 as “true,” f is then a 3-term logical OR operation applied to AND operations, on a set of variables together with their negations.

Theorem 3.7. *If $\mathbf{RP} \neq \mathbf{NP}$, then the class of 3-term DNF formulae is not efficiently PAC learnable, in its own function class.*

Remark 3.8. $\mathbf{P} \subseteq \mathbf{RP} \subseteq \mathbf{NP}$, so it could technically occur that $\mathbf{P} \neq \mathbf{NP}$ while $\mathbf{RP} = \mathbf{NP}$. However it is widely believed that $\mathbf{P} = \mathbf{RP}$.

The above Theorem in [KV94, Section 1.4] says that learning 3-term DNF formulae is not possible to do efficiently, in the class of 3-term DNF formula. However, it is still possible to learn 3-term DNF formulae, when we consider them in a larger class of functions. In fact, this is possible, as we now outline.

Let $a, b, c, d \in \{0, 1\}$. Using the distributive property

$$\max(ab, cd) = \max(a, c) \max(a, d) \max(b, c) \max(b, d),$$

we can rewrite any 3-term DNF formula as a 3-term CNF formula:

$$\max(f_1(x), f_2(x), f_3(x)) = \prod_{\substack{\text{terms } z_1 \text{ in } f_1 \\ \text{terms } z_2 \text{ in } f_2 \\ \text{terms } z_3 \text{ in } f_3}} \max(z_1, z_2, z_3).$$

(A 3-term CNF formula is a function on n variables of the form on the right.) Each term in this large product is treated as its own variable (among $(2n)^3$ possible variables), so that the formula on the right is then interpreted as a conjunction on $8n^3$ variables. We then apply the algorithm of Theorem 3.4 to learn the conjunction, and then rewrite the conjunction as a 3-term CNF formula. For more details see [KV94, Section 1.5]. This argument implies the following two things:

Theorem 3.9. *The class of 3-term CNF formulae is efficiently PAC learnable.*

Corollary 3.10. *The class of 3-term DNF formulae is efficiently PAC learnable when represented as 3-term CNF formulae. That is, the DNF function class is considered as a subset of the larger class of 3-term CNF formulae.*

This corollary is loosely analogous to the discussion of the Minimum Vertex Cover Problem after 1.11. Even though it might be hard to optimize a function on a small domain, enlarging the domain can make it easier to optimize the function (while changing our interpretation of the meaning of the optimum).

Remark 3.11. In the PAC learning framework, we always assume that the hypothesis class \mathcal{F} is **polynomially evaluable**. That is, there is an algorithm such that, for any $f \in \mathcal{F}$ and for any $x \in \{0, 1\}^n$, the value $f(x)$ is output in time polynomial in n and in the size of f .

It turns out the above argument can be generalized, so that, for any $k \geq 2$, k -term DNF formulae are efficiently PAC learnable within the class of k -CNF formulae. Though, as we can imagine, the number of variables involved in passing between DNF and CNF formula seems to be $(2n)^k$. So, if k itself depends on n (e.g. k is a polynomial in n), then the above argument no longer gives an efficient algorithm.

The following is a well-known open problem since the introduction of PAC learning in 1984: are DNF formula of polynomial size (i.e. with $n^{O(1)}$ terms) PAC learnable (within some larger function class)?

It is possible to PAC learn linear threshold functions, but we will postpone this discussion to Section 4. The story for more general learning models of halfspaces is more complicated. In the **agnostic learning model**, the learner has access to a random sample from $A \times B$, i.e. there is no a priori functional relationship that is assumed between A and B . The goal is then to find a $g \in \mathcal{F}$ minimizing $\mathbf{P}(g(X) \neq Y)$, where (X, Y) is the random sample from $A \times B$. (We can think intuitively think of this model as having a target function f with a fraction of its labels randomly corrupted.) In this setting, which is slightly more general than PAC learning, the agnostic learning of linear threshold functions is possible under some assumptions on the distribution \mathbf{P} , when the functions are considered in the larger class of polynomial threshold functions [KKMS08]. However, agnostic learning of linear threshold functions in the class of linear threshold functions is hard [FGRW12]. Moreover, (recalling Exercise 3.5), agnostic learning of conjunctions in the class of linear threshold functions is hard [FGRW12]. With access to Gaussian samples, where a fraction of both samples and labels can be arbitrarily corrupted (the so-called nasty noise model), linear threshold functions can be learned efficiently [DKS18].

Remark 3.12. When a PAC learning algorithm for a class of functions \mathcal{F} outputs a function in the class \mathcal{F} , the algorithm is called **proper**. Otherwise, the algorithm is called **improper**.

3.3. Boosting. In this section, we investigate a modification of the PAC learning model, where with high probability the algorithm is only guaranteed to classify slightly more than 50% of examples correctly. The main question is: can we somehow “boost” the performance of a “weak” learning algorithm like this to get “stronger” PAC learning algorithm.

Problem 3.13 (Supervised Weak Learning Problem). Let A, B be sets. Let $f: A \rightarrow B$ be an unknown function. The goal of the learning problem is to determine the function f on all of A using a small number of *randomly drawn* values of f on A . Let $0 < \varepsilon, \delta < 1/2$ and let \mathbf{P} be a probability law on A . We are given access to ε, δ and to random samples from \mathbf{P} . (The distribution \mathbf{P} may or may not be known, depending on the problem at hand.) If $x \in A$ is a random sample from \mathbf{P} , the pair $(x, f(x))$ is also known.

The goal of the problem is the following. With probability at least $1 - \delta$ (with respect to randomly drawn samples, drawn according to \mathbf{P}), output a function $g: A \rightarrow B$ such that

$$\mathbf{P}(x \in A: f(x) \neq g(x)) < \frac{1}{2} - \varepsilon.$$

That is, the hypothesis g is probably (with probability at least $1 - \delta$) slightly better than random guessing (since a random assignment g would have probability of mis-classification at most $1/2$). We emphasize that the function g can depend on the random sample.

We say a subset of functions $\mathcal{F} \subseteq \{f: A \rightarrow B\}$ is **ε -weak learnable** if there exists $\varepsilon > 0$ and an algorithm such that, for any probability distribution \mathbf{P} on A , there exists $\delta > 0$ such that, for any $0 < \delta < 1/2$, the algorithm achieves the above goal.

In the case $A = \{-1, 1\}^n$ and $B = \{-1, 1\}$, we say that \mathcal{F} is **efficiently weak learnable** if \mathcal{F} is weak learnable, and the algorithm has a run time that is polynomial in $n, 1/\varepsilon$ and $1/\delta$.

In the literature, the family of functions \mathcal{F} is often called a **concept class**.

As noted in Section 2.1.1, given that there is some half space that correctly classifies 99% of data points with $B = \{-1, 1\}$, it is hard to find a half space that correctly classifies

51% of data points [FGKP06]. Note that correctly classifying 50% of data points is easy by just choosing any half space and deciding which side to label +1. However, this result is not directly relevant in the PAC model, since if we are finding an approximation to a target function (such as a linear threshold function), then it is assumed that all points can be correctly classified. In the case that B has $k > 2$ values, random guessing only correctly classifies $1/k$ sample points, so an ε -weak learning algorithm with ε small is a rather strong assumption when compared to the case $k = 2$.

Suppose that we can weakly learn some concept class. Using our intuition from e.g. the Law of Large Numbers, we know that if we can double our money by betting on a game that has a 51% chance of success, then we can earn money with high probability by betting on independent repetitions of the game. Similarly, we can try to sample the weak learner many times and then take a majority vote of the output, thereby “boosting” the weak learner to a much “stronger” (PAC) learning [Sch90]. For an exposition of this argument, see [KV94, Chapter 4], where a nested sequence of majority votes is used, rather than a single majority function. Here we instead focus on a simpler and more practical boosting procedure, albeit without an explicit guarantee of PAC learning.

In this proof, we denote

$$\Delta_k := \{v \in \mathbb{R}^k : \sum_{i=1}^k v_i = 1, v_i \geq 0, \forall 1 \leq i \leq k\}.$$

Given the samples $x^{(1)}, \dots, x^{(k)} \in A$, labels $y_1, \dots, y_k \in \{-1, 1\}$, and given $v \in \Delta_k$, consider the probability law \mathbf{P}_v on A defined so that $\mathbf{P}_v(x^{(i)}) = v_i$ for all $1 \leq i \leq k$.

Algorithm 3.14 (Adaptive Boosting (AdaBoost) [FS97]). The input is $0 < \varepsilon < 1/2$, an ε -weak learning algorithm, a number of iterations t , samples $x^{(1)}, \dots, x^{(k)} \in A$, and labels $y_1, \dots, y_k \in \{-1, 1\}$. Initialize $v^{(1)} \in \Delta_k$, with $v^{(1)} := \frac{1}{t}(1, \dots, 1)$.

For each $1 \leq s \leq t$, do the following.

- Using the ε -weak learning algorithm on input $\mathbf{P}_{v^{(s)}}$, get hypothesis $g_s: A \rightarrow \{-1, 1\}$. (We assume the random samples given to the weak learning algorithm are independent of each other for all $1 \leq s \leq t$.)
- Let $\gamma_s := \sum_{i=1}^k v_i^{(s)} 1_{\{g_s(x^{(i)}) \neq y_i\}}$. (This is the \mathbf{P}_w probability that the weak learner misclassifies its input, so it is at most $1/2 - \varepsilon$ with \mathbf{P}_w probability close to 1. Similarly, the following quantity is typically less than one.)
- Let $\beta_s := \frac{\gamma_s}{1 - \gamma_s}$. (We may assume $\gamma_s > 0$ since if $\gamma_s = 0$, the learner makes no mistakes, so that $\gamma_1 = 0$, so we do not need to apply the boosting algorithm at all.) (So, β_s is the ratio of the \mathbf{P}_w probability of correct classification and the probability of mis-classification.)
- Define $v_i^{(s+1)} := \frac{v_i^{(s)} \beta_s^{1_{\{g_s(x^{(i)}) = y_i\}}}}{\sum_{j=1}^k v_j^{(s)} \beta_s^{1_{\{g_s(x^{(j)}) = y_j\}}}}, \forall 1 \leq i \leq k$.

Output the hypothesis g that is the “weighted majority vote” of each round of hypothesis:

$$g(x) := \text{sign}\left(\sum_{s=1}^t \log\left(\frac{1}{\beta_s}\right) g_s(x)\right).$$

If the weak learner makes few mistakes, then γ_s is close to 0, as is β_s , so $\log(1/\beta_s)$ is large. So the “weight” of each “voter” is directly related to its number of correct classifications. Also, if the weak learner makes many mistakes, then γ_s is close to 1, β_s is much larger than 1, so $\log(1/\beta_s)$ is negative, i.e. g makes the opposite recommendation of g_s . That is, the boosting algorithm still gains something when the weak learner is very wrong, and this is reflected in the bound below.

Theorem 3.15. *Let g be the output of Algorithm 3.14. Define $\varepsilon_s := \frac{1}{2} - \gamma_s$ for all $1 \leq s \leq t$. Then the average number of mis-classifications of g satisfies*

$$\frac{1}{k} |\{1 \leq i \leq k : g(x^{(i)}) \neq y_i\}| \leq e^{-2 \sum_{s=1}^t \varepsilon_s^2}.$$

Exercise 3.16. Explain why taking the expected value of this inequality does not guarantee PAC learning.

Proof. For any $\beta \geq 0$ and for any $a \in [0, 1]$, we have $\beta^a \leq 1 - (1 - \beta)a$. (To see this, note that the second derivative in a is positive on the left, so that $a \mapsto \beta^a$ is convex, and both quantities agree when $a = 0$ and when $a = 1$. So convexity implies the inequality.) Now, let

$$z_s := \sum_{j=1}^k v_j^{(s)} \beta_s^{[1_{\{g_s(x^{(j)})=y_j\}}]}, \quad \forall 1 \leq s \leq t.$$

Using this inequality and the definition of γ_s we have

$$z_s \leq \sum_{j=1}^k v_j^{(s)} [1 - (1 - \beta_s) 1_{\{g_s(x^{(j)})=y_j\}}] = 1 - (1 - \beta_s)(1 - \gamma_s). \quad (\ddagger)$$

By induction, note that, for all $1 \leq i \leq k$,

$$v_i^{(t+1)} = v_i^{(1)} \prod_{s=1}^t \frac{\beta_s^{[1_{\{g_s(x^{(i)})=y_i\}}]}}{z_s}. \quad (*)$$

If the output g mis-classifies $x^{(i)}$ for some $1 \leq i \leq k$, then

$$\begin{aligned} y_i \sum_{s=1}^t \log\left(\frac{1}{\beta_s}\right) g_s(x^{(i)}) \leq 0 &\Leftrightarrow \prod_{s=1}^t \beta_s^{-y_i g_s(x^{(i)})} \leq 1 \\ &\Leftrightarrow \prod_{s=1}^t \beta_s^{[1-2 \cdot 1_{\{g_s(x^{(i)})=y_i\}}]} \leq 1 \\ &\Leftrightarrow \prod_{s=1}^t \beta_s^{[1_{\{g_s(x^{(i)})=y_i\}}]} \geq \left(\prod_{s=1}^t \beta_s\right)^{1/2}. \quad (**) \end{aligned}$$

Using that $w^{(t+1)} \in \Delta_k$,

$$\begin{aligned}
\prod_{s=1}^t z_s &= \sum_{i=1}^k w_i^{(t+1)} \prod_{s=1}^t z_s \stackrel{(*)}{=} \sum_{i=1}^k w_i^{(1)} \prod_{s=1}^t \beta_s^{[1_{\{g_s(x^{(i)})=y_i\}}]} = \frac{1}{k} \sum_{i=1}^k \prod_{s=1}^t \beta_s^{[1_{\{g_s(x^{(i)})=y_i\}}]} \\
&\geq \frac{1}{k} \sum_{i: g(x^{(i)}) \neq y_i} \prod_{s=1}^t \beta_s^{[1_{\{g_s(x^{(i)})=y_i\}}]} \stackrel{(**)}{\geq} \frac{1}{k} \sum_{i: g(x^{(i)}) \neq y_i} \left(\prod_{s=1}^t \beta_s \right)^{1/2} \\
&= \frac{1}{k} |\{1 \leq i \leq k: g(x^{(i)}) \neq y_i\}| \left(\prod_{s=1}^t \beta_s \right)^{1/2}.
\end{aligned}$$

That is,

$$\frac{1}{k} |\{1 \leq i \leq k: g(x^{(i)}) \neq y_i\}| \leq \prod_{s=1}^t \frac{z_s}{\sqrt{\beta_s}} \stackrel{(\ddagger)}{\leq} \prod_{s=1}^t \frac{[1 - (1 - \beta_s)(1 - \gamma_s)]}{\sqrt{\beta_s}}.$$

Substituting $\beta_s = \gamma_s/[1 - \gamma_s]$ (which happens to minimize the right side over all such choices of β_s) concludes the proof since the s^{th} term is then

$$\beta_s^{-1/2}[1 - (1 - 2\gamma_s)] = \beta_s^{-1/2}2\gamma_s = 2\sqrt{\gamma_s(1 - \gamma_s)} = 2\sqrt{((1/2) - \varepsilon_s)((1/2) + \varepsilon_s)} = \sqrt{1 - 4\varepsilon_s^2}.$$

The final inequality follows since $1 - a \leq e^{-a}$ for all $a \in \mathbb{R}$. \square

Remark 3.17. At each iteration of Algorithm 3.14, the weak learner fails with probability at most δ . So, the weak learner succeeds in all t iterations with probability at least $1 - \delta t$, by the union bound.

In Theorem 3.15, the quantity

$$\frac{1}{k} |\{1 \leq i \leq k: g(x^{(i)}) \neq y_i\}| = \frac{1}{k} \sum_{i=1}^k 1_{\{g(x^{(i)}) \neq y_i\}}$$

is called the **empirical risk**, since it gives the average error on the given sample. We would ideally like to bound $\mathbf{P}(g(X) \neq f(X))$, but since we do not have access directly to the random variable X that generates the samples $x^{(1)}, \dots, x^{(k)}$, we cannot compute $\mathbf{P}(g(X) \neq f(X))$ exactly. We will discuss this issue in more detail in Section 6.

In the case that $f: A \rightarrow B$ is the unknown function with B finite and $|B| > 2$, there are a few ways we can modify Algorithm 3.14; here is one way.

Algorithm 3.18 (Multi-Class Adaptive Boosting (AdaBoost.M1) [FS97]). The input is $0 < \varepsilon < 1/2$, an ε -weak learning algorithm, a number of iterations t , samples $x^{(1)}, \dots, x^{(k)} \in A$, and labels $y_1, \dots, y_k \in \{1, \dots, p\}$. Initialize $v^{(1)} \in \Delta_k$, with $v^{(1)} := \frac{1}{t}(1, \dots, 1)$.

For each $1 \leq s \leq t$, do the following.

- Using the ε -weak learning algorithm on input $\mathbf{P}_{v^{(s)}}$, get hypothesis $g_s: A \rightarrow \{1, \dots, p\}$. (We assume the random samples given to the weak learning algorithm are independent of each other for all $1 \leq s \leq t$.)
- Let $\gamma_s := \sum_{i=1}^k v_i^{(s)} 1_{\{g_s(x^{(i)}) \neq y_i\}}$.
- Let $\beta_s := \frac{\gamma_s}{1 - \gamma_s}$.

- Define $v_i^{(s+1)} := \frac{v_i^{(s)} \beta_s^{[1_{\{g_s(x^{(i)})=y_i\}]}}{\sum_{j=1}^k v_j^{(s)} \beta_s^{[1_{\{g_s(x^{(j)})=y_j\}]}}$, $\forall 1 \leq i \leq k$.

Output the hypothesis g that is the “weighted plurality vote” of each round of hypothesis:

$$g(x) := \operatorname{argmax}_{y \in \{1, \dots, p\}} \sum_{s=1}^t \log \left(\frac{1}{\beta_s} \right) 1_{\{g_s(x)=y\}}.$$

3.4. Occam’s Razor. Let $A = \cup_{n=1}^{\infty} A_n$ be a set, let $\mathcal{F} = \cup_{n=1}^{\infty} \mathcal{F}_n$ be a set of functions from A to B , and let $\mathcal{G} = \cup_{n=1}^{\infty} \mathcal{G}_n$ be a set of functions from A to B .

Let $f \in \mathcal{F}$. A labelled sample S with cardinality $m > 0$ sampled according to f is a set of ordered pairs

$$\{(x^{(1)}, f(x^{(1)}), \dots, (x^{(m)}, f(x^{(m)})) : x^{(i)} \in A, \forall 1 \leq i \leq m\}.$$

We assume in this section that the size of $g \in \mathcal{G}$ and $f \in \mathcal{F}$, denoted $\operatorname{size}(\cdot)$, is well-defined to be the bit-length of an encoding of these functions (e.g. \exists a function $\phi_{\mathcal{F}} : \mathcal{F} \rightarrow \{0, 1\}^{\mathbb{N}}$ such that $[\phi_{\mathcal{F}}(f)]_{\operatorname{size}(f)+k} = 0$ for all $k \geq 1$. That is, the binary encoding of f only uses the coordinates from 1 to $\operatorname{size}(f)$.) We say that $g \in \mathcal{G}$ is **consistent** with $f \in \mathcal{F}$ on S if $f(x^{(i)}) = g(x^{(i)})$ for all $1 \leq i \leq m$.

Definition 3.19 (Occam Algorithm). Let $\alpha \geq 0$ and let $0 \leq \beta < 1$ be constants. We say that L is an (α, β) -**Occam algorithm for \mathcal{F} using \mathcal{G}** if, for $n \geq 1$ and for any input sample S of cardinality m labelled according to $f \in \mathcal{F}_n$, L outputs a hypothesis $g \in \mathcal{G}$ such that

- g is consistent with f , and
- $\operatorname{size}(g) \leq (n \cdot \operatorname{size}(f))^{\alpha} m^{\beta}$.

We say that L is an **efficient (α, β) -Occam algorithm** if its run time is at most a polynomial in n, m and $\operatorname{size}(f)$.

Theorem 3.20 (Occam’s Razor). *Let L be an efficient (α, β) -Occam algorithm for \mathcal{F} using \mathcal{G} . Let $0 < \varepsilon, \delta < 1$, let $f \in \mathcal{F}_n$, and let \mathbf{P} be a probability law on A . Then there is a constant $a > 0$ such that, if L is given as input a random sample S of m examples sampled according to f (and \mathbf{P}), with*

$$m \geq a \left(\frac{1}{\varepsilon} \log(1/\delta) + \left(\frac{(n \cdot \operatorname{size}(f))^{\alpha}}{\varepsilon} \right)^{\frac{1}{1-\beta}} \right).$$

Then with probability at least $1 - \delta$, the output g of L satisfies $\mathbf{P}(f(x) \neq g(x)) < \varepsilon$, and the run time of L is at most polynomial in $n, \operatorname{size}(f), 1/\varepsilon$ and $1/\delta$.

In the following result, we assume that $\mathcal{G}_n = \cup_{m=1}^{\infty} \mathcal{G}_{n,m}$, and when the algorithm L has input a sample S of cardinality m , the output is $g \in \mathcal{G}_{n,m}$.

Theorem 3.21 (Occam’s Razor, Cardinality Version). *Let L be an algorithm such that, for any $n \geq 1$ and for any $f \in \mathcal{F}_n$, if L is given as input a random sample S of m examples drawn according to f , then L runs in time at most polynomial in n, m and $\operatorname{size}(f)$, and outputs $g \in \mathcal{G}_{n,m}$ that is consistent with f on S . Then there is a constant $b > 0$ such that, for any n , for any probability law \mathbf{P} on A , and for any $f \in \mathcal{F}_n$, if L is given as input a random sample of m examples drawn according to f (and \mathbf{P}), where $|\mathcal{G}_{n,m}|$ satisfies*

$$\log |\mathcal{G}_{n,m}| \leq b \varepsilon m - \log(1/\delta),$$

(or equivalently m satisfies $m \geq (1/(b\varepsilon))(\log |\mathcal{G}_{n,m}| + \log(1/\delta))$), then L is guaranteed to find $g \in \mathcal{G}_{n,m}$ that with probability at least $1 - \delta$ satisfies $\mathbf{P}(f(x) \neq g(x)) \leq \varepsilon$.

Proof. We say $g \in \mathcal{G}$ is *bad* if $\mathbf{P}(f(x) \neq g(x)) > \varepsilon$. Since the random samples of S are each independent, the probability that a fixed bad hypothesis g is consistent with f on a random sample of cardinality m is at most $(1 - \varepsilon)^m$. Let $\mathcal{G}' \subseteq \mathcal{G}_{n,m}$ be the subset of all bad hypotheses in $\mathcal{G}_{n,m}$. By the union bound, the probability that there exists $g \in \mathcal{G}'$ that is consistent with f on a random sample of cardinality m is at most $|\mathcal{G}'| (1 - \varepsilon)^m$. We want this quantity to be at most δ . Since $\mathcal{G}' \subseteq \mathcal{G}_{n,m}$, it suffices to show that $|\mathcal{G}_{n,m}| (1 - \varepsilon)^m \leq \delta$. Taking logarithms and using $\log(1/(1 - \varepsilon)) = \Theta(\varepsilon)$ concludes the proof. \square

Proof of Theorem 3.20. Let $\mathcal{G}_{n,m}$ be the set of hypotheses that can be output when the input has cardinality m . Since L is an (α, β) -Occam algorithm, every such hypothesis g has bit-length at most $\text{size}(g) \leq (n \cdot \text{size}(f))^\alpha m^\beta$, so that $|\mathcal{G}_{n,m}| \leq 2^{\text{size}(g) \leq (n \cdot \text{size}(f))^\alpha m^\beta}$. By Theorem 3.21, the output g of L satisfies $\mathbf{P}(f \neq g) < \varepsilon$ with probability at least $1 - \delta$, if

$$\log |\mathcal{G}_{n,m}| \leq b\varepsilon m - \log(1/\delta).$$

That is, we want m to satisfy

$$m \geq b^{-1}\varepsilon^{-1} \log |\mathcal{G}_{n,m}| + b^{-1}\varepsilon^{-1} \log(1/\delta).$$

i.e. it suffices to choose m so that $m \geq 2b^{-1}\varepsilon^{-1} \max(\log |\mathcal{G}_{n,m}|, \log(1/\delta))$. Choosing $a = 2/b$ concludes the proof. \square

3.5. Additional Comments. A slightly different version of the Perceptron Algorithm can be stated as follows

Algorithm 3.22 (Perceptron Algorithm, Version 2).

- Define $w^{(1)} := 0 \in \mathbb{R}^n$ and let $s := 1$
- If there exists some $1 \leq i \leq k$ such that $y_i \neq \text{sign}(\langle w^{(s)}, x^{(i)} \rangle)$, i.e. a mis-classification occurs, define

$$w^{(s+1)} := w^{(s)} + y_i \frac{x^{(i)}}{\|x^{(i)}\|}.$$

- Increase the value of s by one. Repeat the previous step until no such i exists.
- Output $w := w^{(s)}$.

We can then state a version of Theorem 2.9 for Algorithm 3.22.

Theorem 3.23 (Perceptron Algorithm Run Time, Version 2). Let $x^{(1)}, \dots, x^{(k)} \in \mathbb{R}^n$ and let $y_1, \dots, y_k \in \{-1, 1\}$ be given. Assume that there exists $w \in \mathbb{R}^n$ such that

$$\text{sign}(\langle w, x^{(i)} \rangle) = y_i, \quad \forall 1 \leq i \leq k.$$

Assume $\|w\| = 1$. Define

$$\sigma := \min_{1 \leq i \leq k} \left| \left\langle w, \frac{x^{(i)}}{\|x^{(i)}\|} \right\rangle \right|.$$

Then the Perceptron Algorithm 3.22 terminates with a value of s satisfying

$$s \leq \sigma^{-2}.$$

The Winnow Algorithm is similar to the Perceptron Algorithm, though the data is usually assumed to be vectors from $\{0, 1\}^n$, and multiplicative updates are made to the weight vector w , rather than additive updates.

In our analysis of the Perceptron Algorithm 2.8, we defined

$$\beta := \max_{i=1, \dots, k} \|x^{(i)}\|, \quad \theta := \min \left\{ \|w\| : \forall 1 \leq i \leq k, y_i \langle w, x^{(i)} \rangle \geq 1 \right\}.$$

We remarked after Lemma 2.11 that the margin $1/\theta$ measures how wide a symmetric “slab” through the origin can separate the vectors $x^{(1)}, \dots, x^{(k)}$ into their two classes. The **support vector machine** is another linear classifier that takes this analysis into account as follows. Let $\lambda > 0$ and suppose we want to find the $w \in \mathbb{R}^n$ and $z_1, \dots, z_k \in \mathbb{R}$ minimizing

$$\lambda \|w\|^2 + \frac{1}{k} \sum_{i=1}^k z_i.$$

subject to the linear constraints

$$y_i \langle w, x^{(i)} \rangle \geq 1 - z_i, \quad z_i \geq 0, \quad \forall 1 \leq i \leq k.$$

This is a quadratic minimization problem subject to linear constraints.

4. VAPNIS-CHERVONENKIS (VC) THEORY

Definition 4.1 (Metric Space). Let A be a set and let $d: A \times A \rightarrow [0, \infty)$. We say that A is a **metric** on A if:

- $d(x, y) \geq 0$ for all $x, y \in A$, and $d(x, y) = 0$ only when $x = y$.
- $d(x, y) = d(y, x)$ for all $x, y \in A$.
- $d(x, y) \leq d(x, z) + d(z, y)$ for all $x, y, z \in A$.

When d is a metric on A , we refer to (A, d) as a **metric space**. For any $x \in A$ and for any $\varepsilon > 0$, we denote the **open ball with radius ε centered at x** as

$$B(x, \varepsilon) = B_d(x, \varepsilon) := \{y \in A : d(x, y) < \varepsilon\}.$$

Definition 4.2 (ε -net). Let A be a set and let $d: A \times A \rightarrow [0, \infty)$ be a metric on A . An **ε -net** is a subset $\{x^{(i)}\}_{i \in I}$ of A such that

$$\bigcup_{i \in I} B(x^{(i)}, \varepsilon) = A.$$

Proposition 4.3. Let $\|\cdot\|$ be a norm on \mathbb{R}^n . Let $\varepsilon > 0$. Then there exists an ε -net \mathcal{N} in the unit ball $B := \{x \in \mathbb{R}^n : \|x\| < 1\}$ such that

$$|\mathcal{N}| \leq (1 + 2/\varepsilon)^n.$$

Proof. Let $\{x^{(i)}\}_{i=1}^k$ be a set of maximal cardinality such that $\|x^{(i)} - x^{(j)}\| \geq \varepsilon$ for all $1 \leq i, j \leq k$ such that $i \neq j$. It follows that $\{x^{(i)}\}_{i=1}^k$ is an ε -net in B (if $B \cap \left(\bigcup_{i=1}^k B(x^{(i)}, \varepsilon)\right) \neq B$, then there exists $x \in B$ such that $\|x^{(i)} - x\| \geq \varepsilon$ for all $1 \leq i \leq k$, contradicting the maximal cardinality of $\{x^{(i)}\}_{i=1}^k$.) By definition of $\{x^{(i)}\}_{i=1}^k$, the open balls $\{B(x^{(i)}, \varepsilon/2)\}_{i=1}^k$ are disjoint and their union is contained in $B(0, 1 + \varepsilon/2)$. So, comparing the volumes of these two sets, we have

$$k(\varepsilon/2)^n \leq (1 + \varepsilon/2)^n.$$

That is, $k \leq (1 + 2/\varepsilon)^n$. □

Exercise 4.4. Let $\|\cdot\|$ be a norm on \mathbb{R}^n . Let $\varepsilon > 0$. Then any ε -net \mathcal{N} in the unit ball $B := \{x \in \mathbb{R}^n : \|x\| < 1\}$ satisfies

$$(1/\varepsilon)^n \leq |\mathcal{N}| \leq (1 + 2/\varepsilon)^n.$$

Definition 4.5 (Covering Number). Let (A, d) be a metric space and let $\varepsilon > 0$. The ε -covering number of (A, d) , denoted $\mathcal{N}(A, d, \varepsilon)$ is the smallest cardinality of an ε -net in A .

Exercise 4.6. Let $\|\cdot\|$ be the standard norm on \mathbb{R}^n . Let $\varepsilon > 0$. Let $B := \{x \in \mathbb{R}^n : \|x\| < 1\}$. Then

$$(1/\varepsilon)^n \leq |\mathcal{N}(B, \|\cdot\|, \varepsilon)| \leq (1 + 2/\varepsilon)^n.$$

Definition 4.7 (VC Dimension). Let A be a set and let $\mathcal{F} \subseteq \{-1, 1\}^A$ be a class of Boolean functions on A . We say a set $B \subseteq A$ is **shattered** by \mathcal{F} if, for any $g: B \rightarrow \{-1, 1\}$, there exists $f \in \mathcal{F}$ such that $f(x) = g(x)$ for all $x \in B$. The **VC-dimension** of \mathcal{F} , denoted $\text{VCdim}(\mathcal{F})$ is the largest cardinality of a subset that is shattered by \mathcal{F} .

Note that by definition of $\text{VCdim}(\mathcal{F})$, we have $|\mathcal{F}| \geq 2^{\text{VCdim}(\mathcal{F})}$. So, $\text{VCdim}(\mathcal{F})$ is somewhat analogous to the log of the cardinality of \mathcal{F} . Let \mathcal{F} be the set of all boolean functions $f: \{-1, 1\}^n \rightarrow \{-1, 1\}$. Then $|\mathcal{F}| = 2^{2^n}$ and $\text{VCdim}(\mathcal{F}) = 2^n$, so in this case the trivial lower bound on $|\mathcal{F}|$ is sharp. However, this lower bound can be quite far from exact, as we discuss below.

Lemma 4.8 (Pajor's Lemma). Let A be a finite set and let $\mathcal{F} \subseteq \{-1, 1\}^A$ be a class of Boolean functions on A . Then

$$|\mathcal{F}| \leq |\{\emptyset \subseteq A' \subseteq A : A' \text{ is shattered by } \mathcal{F}\}|.$$

Proof. We induct on the size of A . The case $|A| = 1$ is clear (e.g. if $\mathcal{F} = 1$, then the right side is also 1 since \emptyset is counted). Suppose then that the Lemma holds when $|A| = n \geq 1$ and consider the case $|A| = n + 1$. Write $A = \{a\} \cup A_0$, where $|A_0| = n$. Let $\mathcal{F}_1 := \{f \in \mathcal{F} : f(a) = 1\}$, $\mathcal{F}_{-1} := \{f \in \mathcal{F} : f(a) = -1\}$. By the inductive hypothesis,

$$|\mathcal{F}_1| \leq |\{\emptyset \subseteq A' \subseteq A_0 : A' \text{ is shattered by } \mathcal{F}_1|_{A_0}\}| =: c_1.$$

$$|\mathcal{F}_{-1}| \leq |\{\emptyset \subseteq A' \subseteq A_0 : A' \text{ is shattered by } \mathcal{F}_{-1}|_{A_0}\}| =: c_{-1}.$$

So, it remains to show that

$$c_1 + c_{-1} \leq c_0 := |\{\emptyset \subseteq A' \subseteq A : A' \text{ is shattered by } \mathcal{F}\}|.$$

If a set $A' \subseteq A_0$ is shattered by $\mathcal{F}_1|_{A_0}$, then A' is also shattered by \mathcal{F} , since $\mathcal{F}_1 \subseteq \mathcal{F}$. The same goes for \mathcal{F}_{-1} . Suppose now that $A' \subseteq A_0$ is in the intersection of the sets defined by c_1 and c_{-1} . Then $c_1 + c_{-1}$ will count the set A' twice, while c_0 only counts it once. However, the set $\{a\} \cup A'$ shattered by \mathcal{F} , since A' is in both sets corresponding to c_1 and c_{-1} , so this set is also counted by c_0 (but not by c_1 or c_{-1}). The inequality $c_1 + c_{-1} \leq c_0$ follows. □

Lemma 4.9 (Sauer-Shelah). Let A be a finite set with $|A| = n$, and let $\mathcal{F} \subseteq \{-1, 1\}^A$ be a class of Boolean functions on A . Let $v := \text{VCdim}(\mathcal{F})$. Then

$$|\mathcal{F}| \leq \sum_{i=0}^d \binom{n}{i} \leq (en/d)^d.$$

Proof. By Pajor's Lemma 4.8,

$$|\mathcal{F}| \leq |\{\emptyset \subseteq A' \subseteq A: A' \text{ is shattered by } \mathcal{F}\}|.$$

By the definition of $\text{VCdim}(\mathcal{F})$, each set counted on the right has cardinality at most d . So,

$$|\mathcal{F}| \leq |\{\emptyset \subseteq A' \subseteq A: |A'| \leq d\}| = \sum_{i=0}^d \binom{n}{i}.$$

The last inequality follows by the binomial theorem since $d \leq n$, so

$$\sum_{i=0}^d \binom{n}{i} (d/n)^d \leq \sum_{i=0}^d \binom{n}{i} (d/n)^i \leq \sum_{i=0}^n \binom{n}{i} (d/n)^i = (1 + (d/n))^n \leq e^d.$$

□

Exercise 4.10. Show that Sauer-Shelah lemma is sharp for all n, d . That is, find \mathcal{F} with $d := \text{VCdim}(\mathcal{F})$ such that

$$|\mathcal{F}| \leq \sum_{i=0}^d \binom{n}{i} \leq (en/d)^d.$$

(Hint: consider the set of $x \in \{0, 1\}^n$ such that x has at most d entries equal to 1.)

Proposition 4.11. Let \mathcal{F} be the set of linear threshold functions

$$\mathcal{F} := \{f: \mathbb{R}^n \rightarrow \{-1, 1\} \quad : \exists w \in \mathbb{R}^n, t \in \mathbb{R}, \quad f(x) = \text{sign}(\langle w, x \rangle - t), \forall x \in \mathbb{R}^n\}.$$

Then $\text{VCdim}(\mathcal{F}) = n + 1$.

Proof. Consider the set of vectors $B := \{0, e_1, \dots, e_n\}$ with exactly $n - 1$ entries equal to 0, together with the zero vector. Then $|B| = n + 1$. We claim that B is shattered by \mathcal{F} . Indeed, if $g: B \rightarrow \{-1, 1\}$ is given, denote $w := (g(e_1) - g(0), \dots, g(e_n) - g(0)) \in \{0, 1\}^n$, then define $f(x) := \text{sign}(\langle w, x \rangle + g(0))$. Then for any $1 \leq i \leq n$, $\langle w, e_i \rangle + g(0) = g(e_i)$, and $f(0) = g(0)$, so $f(x) = g(x)$ for all $x \in B$. We conclude that $\text{VCdim}(\mathcal{F}) \geq n + 1$.

We now show that $\text{VCdim}(\mathcal{F}) \leq n + 1$. We argue by contradiction. Suppose $B \subseteq \mathbb{R}^n$ is a set of $n + 2$ vectors that is shattered by \mathcal{F} . By Remark 2.5, there is a set $B' \subseteq \mathbb{R}^{n+1}$ of $n + 2$ vectors that is shattered by the set of homogeneous linear threshold functions

$$\mathcal{F}' := \{f: \mathbb{R}^{n+1} \rightarrow \{-1, 1\} \quad : \exists w \in \mathbb{R}^n, \quad f(x) = \text{sign}(\langle w, x \rangle), \forall x \in \mathbb{R}^{n+1}\}.$$

Suppose $B' = \{x^{(1)}, \dots, x^{(n+2)}\} \subseteq \mathbb{R}^{n+1}$. Then there exist constants $\alpha_1, \dots, \alpha_{n+2} \in \mathbb{R}$ such that $\sum_{i=1}^{n+2} \alpha_i x^{(i)} = 0$. Let $I := \{1 \leq i \leq n+2: \alpha_i > 0\}$, $J := \{1 \leq i \leq n+2: \alpha_i < 0\}$. Then

$$\sum_{i \in I} \alpha_i x^{(i)} = - \sum_{i \in J} \alpha_i x^{(i)}. \quad (*)$$

Since B' is shattered by \mathcal{F}' , there is a $w \in \mathbb{R}^{n+1}$ such that $\langle w, x^{(i)} \rangle > 0$ for all $i \in I$ and $\langle w, x^{(i)} \rangle < 0$ for all $i \in J$. So, if I is nonempty,

$$0 < \sum_{i \in I} \alpha_i \langle w, x^{(i)} \rangle = \left\langle w, \sum_{i \in I} \alpha_i x^{(i)} \right\rangle \stackrel{(*)}{=} - \left\langle w, \sum_{i \in J} \alpha_i x^{(i)} \right\rangle = - \sum_{i \in J} \alpha_i \langle w, x^{(i)} \rangle \leq 0.$$

So, we have a contradiction when I is nonempty. Since either I or J must be nonempty, the proof is complete. □

The following concept, sometimes simply called an ε -net in the literature, is quite different from the ε -nets we discussed for metric spaces.

Definition 4.12 (Measure-Theoretic ε -net). Let $0 < \varepsilon < 1$. Let Ω be a set and let \mathbf{P} be a probability law on Ω . Let $A \subseteq \Omega$. A measure theoretic ε -net for A is a set of points $S \subseteq A$ such that, for every $A' \subseteq A$ with $\mathbf{P}(A') > \varepsilon$, there exists $s \in S$ such that $\{s\} \cap A' \neq \emptyset$.

That is, every region with large measure contains some point in S . Since no reference to any metric is made in this definition, it is quite different that an ε -net for metric spaces.

Example 4.13. Let Ω be the set of closed intervals in $[0, 1]$, and for any $\omega \in \Omega$, let $\mathbf{P}(\omega)$ be the usual length of ω . Let $\varepsilon > 0$. Then the points $\{i\varepsilon\}_{i=0}^{\lfloor 1/\varepsilon \rfloor}$ is an ε -net for Ω , since any closed interval in $[0, 1]$ of width larger than ε must intersect this set of points.

If on the other hand $\Omega = [0, 1]$ with the same \mathbf{P} , then no finite ε -net exists, since the complement of any finite set has measure 1. However, if $[0, 1]$ has the usual metric on it, then a metric ε -net certainly exists.

Exercise 4.14. Show that both our notions of ε -net agree (up to changing the constant ε) in the following case: Ω is a metric space, \mathbf{P} is a probability law on Ω , $A = \{B(x, r) : x \in \Omega, r > 0\}$ and there exist $a, b, c_1, c_2 > 0$ such that $c_1 r^a \leq \mathbf{P}(B(x, r)) \leq c_2 r^b$ for all $x \in \Omega, r > 0$.

Below, we consider \mathcal{F} to be a subset of $\{0, 1\}$ -valued functions on A . Let \mathbf{P} be a probability law on A . Let $f, g \in \mathcal{F}$. Since $f = 1_{\{f=1\}}$, we can identify f with the set where it is 1 and extend set operations to functions in \mathcal{F} . For example, $f \cap g := 1_{\{f=1\} \cap \{g=1\}}$ and $f \Delta g := 1_{\{f=1\} \Delta \{g=1\}}$, where Δ denotes symmetric difference. Then $f \cap g = g \cap f$ and $f \Delta g = g \Delta f$. Also, we can define $\mathbf{P}(f) := \mathbf{P}(f = 1) = \mathbf{E}f$, so that \mathbf{P} can be extended to a probability law on $\{0, 1\}^A$. The notion of measure-theoretic ε -net for a set of boolean functions is then well-defined using this probability law \mathbf{P} . Define now

$$D(f) := \{f \Delta g : g \in \mathcal{F}\}.$$

Exercise 4.15. For any $f \in \mathcal{F}$,

$$\text{VCdim}(\mathcal{F}) = \text{VCdim}(D(f)).$$

Lemma 4.16. Let $\varepsilon > 0$. Suppose $f \in \mathcal{F}$ is a function to be learned by an algorithm and S is a (measure-theoretic) ε -net for $D(f)$. Suppose the algorithm outputs a hypothesis $g \in \mathcal{F}$ such that $f(s) = g(s)$ for all $s \in S$. Then

$$\mathbf{P}(f \neq g) < \varepsilon.$$

Proof. Since $f(s) = g(s)$ for all $s \in S$, $(g \Delta f)(s) = 0$ for all $s \in S$, so $(g \Delta f) \cap \{s\} = \emptyset$ for all $s \in S$. Since $g \Delta f \in D(f)$, and S is an ε -net for $D(f)$, we conclude by the definition of ε -net that $\mathbf{P}(g \Delta f) < \varepsilon$. That is, $\mathbf{P}(f \neq g) < \varepsilon$. \square

So, if $f \in \mathcal{F}$ is the function to be learned, the above lemma shows that creating an ε -net for $D(f)$ with high probability is sufficient to PAC learn \mathcal{F} . In retrospect, the Occam's Razor bound in Theorem 3.20 used this fact implicitly.

Proposition 4.17. Let \mathbf{P} be a probability law on A . Let $\mathcal{F} \subseteq \{0, 1\}^A$. Let $d := \text{VCdim}(\mathcal{F})$. Let $0 < \varepsilon, \delta < 1$. Let S be a random sample from A of size m where

$$m \geq 100(\varepsilon^{-1} \log(1/\delta) + d\varepsilon^{-1} \log(1/\varepsilon)).$$

Then with probability at least $1 - \delta$, S is a (measure-theoretic) ε -net in \mathcal{F} .

Proof. Let S_1 be a random sample of size m from A . Let C be the event that S_1 does not form an ε -net in \mathcal{F} . Our goal is to upper bound $\mathbf{P}(C)$. If C occurs, then there exists $h \in \mathcal{F}$ with $\mathbf{P}h \geq \varepsilon$ and such that $h \cap S_1 = \emptyset$. Let S_2 be a random size of size m from A that is independent of S_1 . Since $\mathbf{P}h \geq \varepsilon$, if X is the number of times that S_2 intersects h , then

$$\begin{aligned} \mathbf{P}(X < \varepsilon m/2) &\leq P(|X - \varepsilon m| < \varepsilon m/2) \leq 4\varepsilon^{-2}m^{-2}\mathrm{Var}(X) \\ &= 4\varepsilon^{-2}m^{-2}m\varepsilon(1 - \varepsilon) = 4\varepsilon^{-1}m^{-1}(1 - \varepsilon). \end{aligned}$$

So, if $\varepsilon < 1/2$ and $m > 10/\varepsilon$, $\mathbf{P}(X < \varepsilon m/2) < 1/2$. Let C' be the event that C occurs (so there exists $h \in \mathcal{F}$ with $\mathbf{P}h \geq \varepsilon$, $h \cap S_1 = \emptyset$) and $|S_2 \cap h| > \varepsilon m/2$. We have shown that $\mathbf{P}(C'|C) \geq 1/2$, and since $C' \subseteq C$, we have $\mathbf{P}(C'|C) = \mathbf{P}(C')/\mathbf{P}(C)$, so

$$\mathbf{P}(C') \geq (1/2)\mathbf{P}(C).$$

So, in order to upper bound $\mathbf{P}(C)$ it suffices to upper bound $\mathbf{P}(C')$.

Let S be a random sample from A of size $2m$ and let $h \in \mathcal{F}$ with $\mathbf{P}h \geq \varepsilon$. Suppose $|S \cap h| > \varepsilon m/2$. Then, choose two disjoint $T_1, T_2 \subseteq S$ each of cardinality m , uniformly at random among all such partitions of S into two equal sized sets. Then since (T_1, T_2) is equal in distribution to (S_1, S_2) , we have

$$\mathbf{P}(C') = \mathbf{P}(\exists \tilde{h} \in \mathcal{F}: \mathbf{P}(\tilde{h}) \geq \varepsilon, \tilde{h} \cap T_1 = \emptyset).$$

The probability $\mathbf{P}(\tilde{h} \cap T_1 = \emptyset)$ only depends on $\tilde{h}|_{T_1}$. From the union bound and Lemma 4.9,

$$\mathbf{P}(C') \leq \sum_{\tilde{h} \in \mathcal{F}|_{T_1}: \mathbf{P}(\tilde{h}) \geq \varepsilon} \mathbf{P}(\tilde{h} \cap T_1 = \emptyset) \leq (em/d)^d \mathbf{P}(h \cap T_1 = \emptyset). \quad (*)$$

The probability $\mathbf{P}(h \cap T_1 = \emptyset)$ can be computed from the following combinatorial problem. Suppose we have $2m$ cubes sorted into piles U_1, U_2 each of size m , and we label $\ell \geq \varepsilon m/2$ cubes red, uniformly at random. then $\mathbf{P}(h \cap T_1 = \emptyset)$ is the probability that all red cubes are in T_1 . If ℓ is fixed, this probability is

$$\frac{\binom{m}{\ell}}{\binom{2m}{\ell}} = \frac{m!(2m - \ell)!}{2m!(m - \ell)!} = \prod_{i=0}^{\ell-1} \frac{m - i}{(2m - i)} \leq \prod_{i=0}^{\ell-1} \frac{1}{2} = 2^{-\ell}.$$

So,

$$\mathbf{P}(h \cap T_1 = \emptyset) \leq \sum_{j=\ell}^{\infty} 2^{-j} \leq 2^{-\ell+1} \leq 2^{1-\varepsilon m/2}. \quad (**)$$

Combining (*) and (**) gives

$$\mathbf{P}(C') \leq 2 \left(\frac{em}{d} \right)^d 2^{-\varepsilon m/2}.$$

That is,

$$\log \mathbf{P}(C) \leq \log 4 + d[1 + \log(m/d)] - \log(2)\varepsilon m/2$$

So, choosing $m \geq 10\varepsilon^{-1} \log(1/\delta) + 10d\varepsilon^{-1} \log(1/\varepsilon)$ means $\log \mathbf{P}(C) \leq \log \delta$, as desired. \square

Combining Exercise 4.15 with Propositions 4.16 and 4.17 proves the following Theorem, sometimes called the Fundamental Theorem of Statistical Learning or the Fundamental Theorem of Machine Learning

Theorem 4.18 (Fundamental Theorem of Statistical Learning, Version 1). *Let A be a set. Let $\mathcal{F} \subseteq \{0, 1\}^A$ be a class of boolean functions. Let $d := \text{VCdim}(\mathcal{F})$. Suppose $f \in \mathcal{F}$ is a function to be learned by an algorithm. Let S be a random sample from A of size m where*

$$m \geq 100(\varepsilon^{-1} \log(1/\delta) + d\varepsilon^{-1} \log(1/\varepsilon)).$$

Then with probability at least $1 - \delta$, S is a (measure-theoretic) ε -net for $D(f)$. Suppose the algorithm outputs a hypothesis $g \in \mathcal{F}$ such that $f(s) = g(s)$ for all $s \in S$. Then

$$\mathbf{P}(f \neq g) < \varepsilon.$$

That is, the algorithm can PAC learn \mathcal{F} .

Theorem 4.18 immediately implies that, if an algorithm can output a hypothesis g that agrees with the values of f on the random sample of sufficiently large size, then this algorithm can PAC learn. So, knowing that a function class has a relatively small VC-dimension immediately implies that it can be PAC learned. One major caveat of Theorem 4.18 is that **it does not guarantee efficient PAC learnability**. In fact, the task of finding the hypothesis $g \in \mathcal{F}$ that agrees with f might be computationally hard. For example, if \mathcal{F} is the class of 3-term DNF formulae on n variables, then $\text{VCdim}(\mathcal{F}) \leq \log_2 |\mathcal{F}| \leq \log_2(3^{3n}) \leq 6n$, but we know from Theorem 3.7 that \mathcal{F} is not efficiently PAC learnable (in its own function class).

Theorem 4.19 (Fundamental Theorem of Statistical Learning, Version 2). *Let A be a set. Let $\mathcal{F}, \mathcal{G} \subseteq \{0, 1\}^A$ be two classes of boolean functions. Let $d := \text{VCdim}(\mathcal{G})$. Suppose $f \in \mathcal{F}$ is a function to be learned by an algorithm. Let S be a random sample from A of size m where*

$$m \geq 100(\varepsilon^{-1} \log(1/\delta) + d\varepsilon^{-1} \log(1/\varepsilon)).$$

Then with probability at least $1 - \delta$, S is a (measure-theoretic) ε -net for $D(f)$. Suppose the algorithm outputs a hypothesis $g \in \mathcal{G}$ such that $f(s) = g(s)$ for all $s \in S$. Then

$$\mathbf{P}(f \neq g) < \varepsilon.$$

That is, the algorithm can PAC learn \mathcal{F} , viewed as a subset of \mathcal{G} .

4.1. Applications of the Fundamental Theorem.

Proposition 4.20. *Let \mathcal{G} be a class of functions on a set A . Let $n \geq 3$. Let \mathcal{F} be the set of linear threshold functions in n elements of \mathcal{G} :*

$$\mathcal{F} := \{f: \mathbb{R}^n \rightarrow \{-1, 1\} \quad : \exists w \in \mathbb{R}^n, t \in \mathbb{R}, \quad f(x) = \text{sign}\left(\sum_{i=1}^n w_i g_i(x) - t\right), \forall x \in A\}.$$

Then $\text{VCdim}(\mathcal{F}) \leq (\dots)$.

Proof. □

Proposition 4.21. *(sample complexity of boosting, with VC-dimension bounds)*

Proof. □

Theorem 4.22 (PAC Learning Linear Threshold Functions).

5. SOME CONCENTRATION OF MEASURE

The mathematical tools contained in this section will be used in Section 6.

5.1. Concentration for Independent Sums. In certain cases, we can make rather strong conclusions about the distribution of sums of i.i.d. random variables, improving upon the laws of large numbers.

Theorem 5.1 (Hoeffding Inequality/ Large Deviation Estimate). *Let X_1, X_2, \dots be independent identically distributed random variables with $\mathbf{P}(X_1 = 1) = \mathbf{P}(X_1 = -1) = 1/2$. Let $a_1, a_2, \dots \in \mathbb{R}$. Then, for any $n \geq 1$,*

$$\mathbf{P}\left(\sum_{i=1}^n a_i X_i \geq t\right) \leq e^{-\frac{t^2}{2\sum_{i=1}^n a_i^2}}, \quad \forall t \geq 0.$$

Consequently,

$$\mathbf{P}\left(\left|\sum_{i=1}^n a_i X_i\right| \geq t\right) \leq 2e^{-\frac{t^2}{2\sum_{i=1}^n a_i^2}}, \quad \forall t \geq 0.$$

Proof. By dividing a_1, \dots, a_n by a constant, we may assume $\sum_{i=1}^n a_i^2 = 1$. Let $\alpha > 0$. Using the (exponential) moment method as in Markov's inequality, and $\alpha t \geq 0$,

$$\mathbf{P}\left(\sum_{i=1}^n a_i X_i \geq t\right) = \mathbf{P}\left(e^{\alpha \sum_{i=1}^n a_i X_i} \geq e^{\alpha t}\right) \leq e^{-\alpha t} \mathbf{E} e^{\alpha \sum_{i=1}^n a_i X_i} = e^{-\alpha t} \prod_{i=1}^n \mathbf{E} e^{\alpha a_i X_i}.$$

The last equality used independence of X_1, X_2, \dots . Using an explicit computation and Exercise 5.2,

$$\mathbf{E} e^{\alpha a_i X_i} = (1/2)(e^{\alpha a_i} + e^{-\alpha a_i}) = \cosh(\alpha a_i) \leq e^{\alpha^2 a_i^2 / 2}, \quad \forall i \geq 1.$$

In summary, for any $t \geq 0$

$$\mathbf{P}\left(\sum_{i=1}^n a_i X_i \geq t\right) \leq e^{-\alpha t} e^{\alpha^2 \sum_{i=1}^n a_i^2 / 2} = e^{-\alpha t + \alpha^2 / 2}.$$

Since $\alpha > 0$ is arbitrary, we choose α to minimize the right side. This minimum occurs when $\alpha = t$, so that $-\alpha t + \alpha^2 / 2 = -t^2 / 2$, giving the first desired bound. The final bound follows by writing $\mathbf{P}(|\sum_{i=1}^n a_i X_i| \geq t) = \mathbf{P}(\sum_{i=1}^n a_i X_i \geq t) + \mathbf{P}(-\sum_{i=1}^n a_i X_i \geq t)$ and then applying the first inequality twice. \square

Exercise 5.2. Show that $\cosh(x) \leq e^{x^2/2}, \forall x \in \mathbb{R}$.

In particular, Hoeffding's inequality implies that

$$\mathbf{P}\left(\frac{1}{n} \left| \sum_{i=1}^n X_i \right| \geq t\right) \leq 2e^{-nt^2/2}, \quad \forall t \geq 0.$$

This inequality is much stronger than either Markov's or Cheyshev's inequality, since they only respectively imply that

$$\mathbf{P}\left(\frac{1}{n} \left| \sum_{i=1}^n X_i \right| \geq t\right) \leq \frac{1}{t}, \quad \mathbf{P}\left(\frac{1}{n} \left| \sum_{i=1}^n X_i \right| \geq t\right) \leq \frac{1}{nt^2}, \quad \forall t \geq 0.$$

Note also that Hoeffding's inequality gives a quantitative bound for any fixed $n \geq 1$, unlike the (non-quantitative) limit theorems which only hold as $n \rightarrow \infty$.

Exercise 5.3 (Chernoff Inequality). Let $0 < p < 1$. Let X_1, X_2, \dots be independent identically distributed random variables with $\mathbf{P}(X_1 = 1) = p$ and $\mathbf{P}(X_1 = 0) = 1 - p$ for any $i \geq 1$. Then for any $n \geq 1$

$$\mathbf{P}\left(\frac{1}{n} \sum_{i=1}^n X_i \geq t\right) \leq e^{-np} \left(\frac{ep}{t}\right)^{tn}, \quad \forall t \geq p.$$

Prove the same estimate for $\mathbf{P}\left(\frac{1}{n} \sum_{i=1}^n X_i \leq t\right)$ for any $t \leq p$. (Hint: $1 + x \leq e^x$ for any $x \in \mathbb{R}$, so $1 + (e^\alpha - 1)p \leq e^{(e^\alpha - 1)p}$.)

Exercise 5.4. For any natural number n and a parameter $0 < p < 1$, define an Erdős-Renyi graph on n vertices with parameter p to be a random graph (V, E) on a (deterministic) vertex set V of n vertices (thus (V, E) is a random variable taking values in the discrete space of all $2^{\binom{n}{2}}$ possible undirected graphs one can place on V) such that the events $\{i, j\} \in E$ for unordered pairs with $i, j \in V$ are independent and each occur with probability p .

Suppose we have an Erdős-Renyi random graph $G = (V, E)$ on n vertices with parameter $0 < p < 1$. Define $d := p(n - 1)$.

- Show that d is the expected degree of each vertex in G . (The degree of a vertex $v \in V$ is the number of vertices connected to v by an edge in E .)
- Show that there exists a constant $c > 0$ such that the following holds. Assume $p \geq \frac{c \log n}{n}$. Then with probability larger than .9, all vertices of G have degrees in the range $(.9d, 1.1d)$. (Hint: first consider a single vertex, then use the union bound over all vertices.)

5.2. Concentration for Lipschitz Functions. One way to phrase the general question in the subject of concentration of measure is: how far is a random variable from its mean value? Hoeffding's Inequality says that linear functions of mean zero ± 1 valued independent random variables are exponentially close to their mean value. A similar statement can be made for bounded random variables (see Theorem 5.8 below). In order to answer the general question, we next consider Lipschitz functions of i.i.d. random variables. We focus on the Gaussian setting for simplicity.

For any $x = (x_1, \dots, x_n) \in \mathbb{R}^n$, we denote $\|x\| := (x_1^2 + \dots + x_n^2)^{1/2}$.

Theorem 5.5 (Concentration of measure for Gaussians, Lipschitz function form). Let $f: \mathbb{R}^n \rightarrow \mathbb{R}$. Suppose that for all $x, y \in \mathbb{R}^n$, $|f(x) - f(y)| \leq \|x - y\|$, so that f is 1-Lipschitz. Let $X = (X_1, \dots, X_n)$ be a mean zero Gaussian random vector with identity covariance matrix. Then for all $t > 0$,

$$\mathbf{P}(x \in \mathbb{R}^n: |f(x) - \mathbf{E}f(X)| \geq t) \leq 2e^{-2t^2/\pi^2}.$$

Proof. We assume that f all partial derivatives of f exist and are continuous. Let $Y = (Y_1, \dots, Y_n)$ be another mean zero Gaussian random vector with identity covariance matrix, such that Y and X are independent. Let $0 \leq \theta \leq \pi/2$ and define

$$Z_\theta := X \sin \theta + Y \cos \theta.$$

By rotation invariance of a Gaussian random vector, Z_θ and $\frac{d}{d\theta} Z_\theta = X \cos \theta - Y \sin \theta$ have the same joint distribution as X and Y (since the vectors $(\sin \theta, \cos \theta)$ and $(\cos \theta, -\sin \theta)$ are

orthogonal in \mathbb{R}^2 .) Let $\phi: \mathbb{R} \rightarrow [0, \infty)$ be a convex function. Using then Jensen's Inequality, then the Chain Rule, then Jensen's inequality and Fubini's Theorem,

$$\begin{aligned} \mathbf{E}\phi(f(X) - \mathbf{E}f(Y)) &\leq \mathbf{E}\phi(f(X) - f(Y)) = \mathbf{E}\phi\left(\int_0^{\pi/2} \frac{d}{d\theta}f(Z_\theta)d\theta\right) \\ &= \mathbf{E}\phi\left(\int_0^{\pi/2} \langle(\nabla f)(Z_\theta), \frac{d}{d\theta}Z_\theta\rangle d\theta\right) = \mathbf{E}\phi\left(\frac{1}{\pi/2} \int_0^{\pi/2} \frac{\pi}{2} \langle(\nabla f)(Z_\theta), \frac{d}{d\theta}Z_\theta\rangle d\theta\right) \\ &\leq \mathbf{E}\frac{1}{\pi/2} \int_0^{\pi/2} \phi\left(\frac{\pi}{2} \langle(\nabla f)(Z_\theta), \frac{d}{d\theta}Z_\theta\rangle\right) d\theta = \frac{1}{\pi/2} \int_0^{\pi/2} \mathbf{E}\phi\left(\frac{\pi}{2} \langle(\nabla f)(Z_\theta), \frac{d}{d\theta}Z_\theta\rangle\right) d\theta \\ &= \frac{1}{\pi/2} \int_0^{\pi/2} \mathbf{E}\phi\left(\frac{\pi}{2} \langle(\nabla f)(X), Y\rangle\right) d\theta = \mathbf{E}\phi\left(\frac{\pi}{2} \langle(\nabla f)(X), Y\rangle\right) \end{aligned}$$

Let $\alpha \in \mathbb{R}$ and let $\phi(x) := e^{\alpha x}$ for all $x \in \mathbb{R}$. Then using independence in Y and Fubini's Theorem,

$$\mathbf{E}\exp(\alpha[f(X) - \mathbf{E}f(Y)]) \leq \mathbf{E}\exp\left(\alpha\frac{\pi}{2} \sum_{i=1}^n \frac{\partial f}{\partial x_i}(X)Y_i\right) = \mathbf{E}_X \prod_{i=1}^n \mathbf{E}_Y \exp\left(\alpha\frac{\pi}{2} \frac{\partial f}{\partial x_i}(X)Y_i\right).$$

Using an explicit computation, for any $s \in \mathbb{R}$ and for any $1 \leq i \leq n$,

$$\mathbf{E}_Y e^{sY_i} = \int_{-\infty}^{\infty} e^{sy} e^{-y^2/2} \frac{dy}{\sqrt{2\pi}} = e^{s^2/2} \int_{-\infty}^{\infty} e^{-(y-s)^2/2} \frac{dy}{\sqrt{2\pi}} = e^{s^2/2}.$$

So, applying this inequality with $s = \alpha\frac{\pi}{2} \frac{\partial f}{\partial x_i}(X)$ for each $1 \leq i \leq n$,

$$\mathbf{E}\exp(\alpha[f(X) - \mathbf{E}f(Y)]) \leq \mathbf{E}\exp\left(\alpha^2 \frac{\pi^2}{8} \sum_{i=1}^n \left(\frac{\partial f}{\partial x_i}(X)\right)^2\right) \leq \exp\left(\alpha^2 \frac{\pi^2}{8}\right).$$

(Since f is 1-Lipschitz, $|\langle \nabla f(x), y \rangle| \leq 1$ for all $x, y \in \mathbb{R}^n$ with $\|y\| \leq 1$. In particular, using $y := \nabla f(x) / \|\nabla f(x)\|$, we get $\|\nabla f(x)\| \leq 1$.) So,

$$\begin{aligned} \mathbf{P}(f(X) - \mathbf{E}f(Y) > t) &= \mathbf{P}(\exp(\alpha[f(X) - \mathbf{E}f(Y)]) > e^{\alpha t}) \\ &\leq e^{-\alpha t} \exp\left(\alpha^2 \frac{\pi^2}{8}\right) = \exp\left(-\alpha t + \alpha^2 \frac{\pi^2}{8}\right). \end{aligned}$$

The minimum α occurs when $\alpha = 4t/\pi^2$, so making this choice of α , we get

$$\mathbf{P}(f(X) - \mathbf{E}f(Y) > t) \leq \exp(-2t^2/\pi^2).$$

Similarly, $\mathbf{P}(f(X) - \mathbf{E}f(Y) < -t) \leq \exp(-2t^2/\pi^2)$, so that

$$\begin{aligned} \mathbf{P}(|f(X) - \mathbf{E}f(Y)| > t) &= \mathbf{P}(f(X) - \mathbf{E}f(Y) > t) + \mathbf{P}(f(X) - \mathbf{E}f(Y) < -t) \\ &\leq 2\exp(-2t^2/\pi^2). \end{aligned}$$

□

Theorem 5.6 (Johnson-Lindenstrauss Lemma). *Let $x^{(1)}, \dots, x^{(n)} \in \mathbb{R}^m$. Let $\varepsilon > 0$. Then there exists a linear function $h: \mathbb{R}^m \rightarrow \mathbb{R}^{O(\varepsilon^{-2} \log n)}$ such that*

$$\|x^{(i)} - x^{(j)}\| \leq \|h(x^{(i)}) - h(x^{(j)})\| \leq (1 + \varepsilon) \|x^{(i)} - x^{(j)}\|, \quad \forall 1 \leq i, j \leq n.$$

One proves this via the probabilistic method. By concentration of measure, a random projection does what we require.

Proof. Fix $1 \leq k \leq m$. Let $\Pi: \mathbb{R}^m \rightarrow \mathbb{R}^m$ be the orthogonal projection such that

$$\Pi(z_1, \dots, z_m) := (z_1, \dots, z_k, 0, \dots, 0), \quad \forall (z_1, \dots, z_m) \in \mathbb{R}^m.$$

Let $X = (X_1, \dots, X_m)$ be a standard m -dimensional Gaussian random vector. Define

$$a := \mathbf{E} \|\Pi X\|.$$

We will eventually show that $a \geq 10^{-2}\sqrt{k}$. Observe

$$\mathbf{E} \|\Pi X\|^2 = \mathbf{E} \sum_{i=1}^k X_i^2 = k \mathbf{E} X_1^2 = k. \quad (*)$$

Now, by Theorem 5.5 for the 1-Lipschitz function $x \mapsto \|\Pi x\|$,

$$\begin{aligned} \mathbf{E} \|\Pi X\|^4 &= \int_0^\infty 4u^3 \mathbf{P}(\|\Pi X\| \geq u) du \\ &= \int_0^{2a} 4u^3 \mathbf{P}(\|\Pi X\| \geq u) du + \int_{2a}^\infty 4u^3 \mathbf{P}(\|\Pi X\| \geq u) du \\ &\leq \int_0^{2a} 4u^3 du + \int_{2a}^\infty 4u^3 \mathbf{P}(|\|\Pi X\| - a| > u/2) du \\ &\leq 16a^4 + 8 \int_{2a}^\infty u^3 e^{-u^2/2\pi^2} du = 16a^4 + 8(2\pi^2)(2a^2 + \pi^2)e^{-2a^2/\pi^2} \leq 16a^4 + 2\pi^4 \\ &\leq 16a^4 + 200k^2 \leq 216 \left(\int_{\mathbb{R}^m} \|\Pi x\|^2 \gamma_m(x) dx \right)^2, \text{ using Jensen's inequality and } (*). \end{aligned}$$

So, if $Z := \|\Pi X\|$ is a random variable, we have shown that $\mathbf{E} Z^4 < c(\mathbf{E} Z^2)^2$ where $c := 216$. So, using Hölder's Inequality, for $p = 3/2$, $q = 3$,

$$\mathbf{E} Z^2 = \mathbf{E}(Z^{2/3} Z^{4/3}) \leq (\mathbf{E} Z)^{2/3} (\mathbf{E} Z^4)^{1/3} \leq (\mathbf{E} Z)^{2/3} c^{1/3} (\mathbf{E} Z^2)^{2/3}.$$

Using this inequality and (*),

$$\mathbf{E} Z \geq c^{-1/2} \sqrt{\mathbf{E} Z^2} \geq 216^{-1/2} \sqrt{k}. \quad (**)$$

In summary, $a \geq 2^{-4}\sqrt{k}$ for a defined above.

Let A be an $m \times m$ matrix of i.i.d. standard Gaussian random variables. Fix $x^{(0)} \in \mathbb{R}^m$ with $\|x^{(0)}\| = 1$. By rotation invariance of the Gaussian measure, A and AQ have the same distribution where Q is a fixed $m \times m$ orthogonal matrix, so if we choose Q so that $Q(1, 0, \dots, 0)^T = x^{(0)}$, we get

$$\begin{aligned} \mathbf{P}(A \in \mathbb{R}^{m \times m}: \|\Pi A x^{(0)}\|_2 - a \geq \varepsilon a) &= \mathbf{P}(A \in \mathbb{R}^{m \times m}: \|\Pi A(1, 0, \dots, 0)^T\|_2 - a \geq \varepsilon a) \\ &= \mathbf{P}(X \in \mathbb{R}^m: \|\Pi X\| - a \geq \varepsilon a). \end{aligned}$$

So, by Theorem 5.5 applied to the 1-Lipschitz function $x \mapsto \|\Pi x\|$, and using $a \geq 2^{-4}\sqrt{k}$, for any $\varepsilon > 0$, and for any

$$\mathbf{P}(A \in \mathbb{R}^{m \times m}: \|\Pi A x^{(0)}\|_2 - a \geq \varepsilon a) \leq 2e^{-2\varepsilon^2 a^2/\pi^2} \leq 2e^{-2^{-10}k\varepsilon^2}.$$

Let $x^{(1)}, \dots, x^{(n)}$ be n points in \mathbb{R}^m . If $k \geq 2^{12}\varepsilon^{-2} \log n$, the union bound shows that

$$\mathbf{P}\left(A \in \mathbb{R}^{m \times m}: \exists i \neq j: \left| \left\| \Pi A \left(\frac{x^{(i)} - x^{(j)}}{\|x^{(i)} - x^{(j)}\|} \right) \right\| - a \right| \geq \varepsilon a \right) \leq \binom{n}{2} 2e^{-2^{-10}k\varepsilon^2} < 1.$$

For any $1 \leq i \leq n$, define $y_i := \Pi A x^{(i)} / (a(1 - \varepsilon))$. Then $\exists A \in \mathbb{R}^{n \times m}$ such that

$$1 \leq \left\| \frac{y^{(i)} - y^{(j)}}{\|x^{(i)} - x^{(j)}\|} \right\| \leq \frac{1 + \varepsilon}{1 - \varepsilon} \leq 1 + 3\varepsilon, \quad \forall 1 \leq i, j \leq n.$$

So, our required embedding is $h := \frac{\Pi A}{a(1 - \varepsilon)}$, so that $h(x^{(i)}) = y^{(i)}$ for all $1 \leq i \leq n$. Note that h is linear and its nonzero entries form a rectangular matrix of i.i.d. Gaussians. Also, we can choose $k := \lceil 2^{12} \varepsilon^{-2} \log n \rceil$. (In fact, if we choose k to be slightly larger, then the probability becomes exponentially small, so essentially all A satisfies our desired property, hence essentially all linear projections $h: \mathbb{R}^n \rightarrow \mathbb{R}^{O(\varepsilon^{-2} \log n)}$ satisfy our desired property.) \square

Exercise 5.7. High-dimensional geometry is much different than low-dimensional geometry, as this exercise demonstrates.

- Show that “most” of the mass of an n -dimensional Gaussian is concentrated on the sphere of radius \sqrt{n} centered at the origin. That is, if X_1, \dots, X_n are n i.i.d. standard Gaussian random variables, then

$$\lim_{n \rightarrow \infty} \mathbf{P}(\sqrt{X_1^2 + \dots + X_n^2} \in (n + \sqrt{3n}, n - \sqrt{3n})) \geq 2/3.$$

In fact, you should be able to compute the limit exactly.

- Generally, “most” of the mass of a high-dimensional convex body is concentrated near the surface of the body. Let Vol_n denote the usual volume in \mathbb{R}^n (so that the volume of a unit square $[0, 1]^n$ is 1.) For example, show that, for any $\varepsilon > 0$,

$$\lim_{n \rightarrow \infty} \text{Vol}_n\left(\left[-\frac{1}{2}(1 - \varepsilon), \frac{1}{2}(1 - \varepsilon)\right]^n\right) = 0.$$

- Let $B_n := \{x \in \mathbb{R}^n : \|x\| \leq 1\}$ be the unit ball centered at the origin. Show that

$$\lim_{n \rightarrow \infty} \text{Vol}_n(B_n) = 0.$$

- Let $C_n = \{x \in \{[-1/2, 1/2]^n : \exists y \in \{-1/2, 1/2\}^n \text{ such that } \|x - y\| \leq 1/2\}\}$ be the union of balls of radius $1/2$ centered at the corners of the hypercube $[-1/2, 1/2]^n$. Let $D_n := \{x \in \mathbb{R}^n : \|x\| \leq r\}$ be a ball of radius r centered at the origin, where r is chosen to be as large as possible so that D_n does not intersect the interior of C_n . (Put another way, D_n is tangent to the balls C_n .) Find

$$\lim_{n \rightarrow \infty} \text{Vol}_n(D_n).$$

Before you do the computation, try to guess what the answer should be.

5.3. Additional Comments. Hoeffding’s inequality in Theorem 5.1 can be generalized to the following statement.

Theorem 5.8 (Hoeffding Inequality/ Large Deviation Estimate). For all $i \geq 1$, let $a_i < b_i$ be real numbers. Let X_1, X_2, \dots be independent random variables with $\mathbf{P}(X_i \in [a_i, b_i]) = 1$. Then, for any $n \geq 1$,

$$\mathbf{P}\left(\sum_{i=1}^n X_i - \mathbf{E}\left(\sum_{j=1}^n X_j\right) \geq t\right) \leq e^{-\frac{2t^2}{\sum_{i=1}^n (b_i - a_i)^2}}, \quad \forall t \geq 0.$$

Consequently,

$$\mathbf{P}\left(\left|\sum_{i=1}^n X_i - \mathbf{E}\left(\sum_{j=1}^n X_j\right)\right| \geq t\right) \leq 2e^{-\frac{2t^2}{\sum_{i=1}^n (b_i - a_i)^2}}, \quad \forall t \geq 0.$$

The proof of Theorem 5.8 imitates that of Theorem 5.1, while using the following Lemma.

Lemma 5.9 (Hoeffding's Lemma). *Let $a < b$ be real numbers. Let X be a random variable with $\mathbf{P}(X \in [a, b]) = 1$. Then for any $\alpha \in \mathbb{R}$,*

$$\mathbf{E}e^{\alpha(X - \mathbf{E}X)} \leq e^{\frac{1}{8}\alpha^2(b-a)^2}.$$

Proof. By replacing X with $X - \mathbf{E}X$, we may assume that $\mathbf{E}X = 0$ and instead prove that

$$\mathbf{E}e^{\alpha X} \leq e^{\frac{1}{8}\alpha^2(b-a)^2}.$$

Since $a \leq X \leq b$ with probability one, there is a random $Y \in [0, 1]$ such that $X =: aY + b(1 - Y)$. That is, this equality holds when $Y := (X - b)/(a - b)$. By convexity of $x \mapsto e^{\alpha x}$,

$$e^{\alpha X} = e^{\alpha(aY + b(1-Y))} \leq Y e^{\alpha a} + (1 - Y)e^{\alpha b} = \frac{X - b}{a - b} e^{\alpha a} + \frac{a - X}{a - b} e^{\alpha b}.$$

Let $\gamma := -b/(a - b)$, $c := \alpha(a - b)$, $f(c) := -\gamma c + \log(1 - \gamma + \gamma e^c)$. Taking expectations of both sides and using $\mathbf{E}X = 0$ we get

$$\mathbf{E}e^{\alpha X} \leq -\frac{b}{a - b} e^{\alpha a} + \frac{a}{a - b} e^{\alpha b} = e^{f(c)}.$$

Note that $f(0) = 0$, $f'(c) = -\gamma + \frac{\gamma e^c}{1 - \gamma + \gamma e^c}$, $f'(0) = 0$, and

$$f''(c) = \frac{\gamma e^c}{1 - \gamma + \gamma e^c} - \frac{\gamma^2 e^{2c}}{(1 - \gamma + \gamma e^c)^2} = \frac{\gamma e^c}{1 - \gamma + \gamma e^c} \left(1 - \frac{\gamma e^c}{1 - \gamma + \gamma e^c}\right)$$

So, if $s := \frac{\gamma e^c}{1 - \gamma + \gamma e^c}$, we have $f''(c) = s(1 - s) \leq 1/4$. So, by Taylor's Theorem (with error term), for any $c \in \mathbb{R}$, there exists c_0 between 0 and c such that

$$f(c) = f(0) + cf'(0) + \frac{c^2}{2} f''(c_0) = \frac{c^2}{2} f''(c_0) \leq \frac{c^2}{8}.$$

In conclusion $\mathbf{E}e^{\alpha X} \leq e^{c^2/8}$. □

Theorem 5.5 can be generalized to uniformly log-concave densities on Euclidean space (see Ledoux, "The Concentration of Measure Phenomenon," Proposition 2.18)

Theorem 5.10 (Concentration of measure for Log-Concave Measures, Lipschitz function form). *Let $f: \mathbb{R}^n \rightarrow \mathbb{R}$. Suppose that for all $x, y \in \mathbb{R}^n$, $|f(x) - f(y)| \leq \|x - y\|$, so that f is 1-Lipschitz. Let $u: \mathbb{R}^n \rightarrow \mathbb{R}$ be a function such that $e^{-u(x)}$ is a probability density on \mathbb{R}^n . Assume there exists $c > 0$ such that the Hessian of u satisfies $\text{Hess}(u)(x) \geq cI$, in the matrix sense. (That is, all eigenvalues of the Hessian of u are bounded below by c , for all $x \in \mathbb{R}^n$.) Let X have distribution e^{-u} . Then, for all $t > 0$,*

$$\mathbf{P}(x \in \mathbb{R}^n: |f(x) - \mathbf{E}f(X)| \geq t) \leq 2e^{-ct^2/2}.$$

Note that Hoeffding's Inequality 5.8 provides the same bound when $\mathbf{P}(X_1 = 1) = \mathbf{P}(X_1 = -1) = 1$, or when $\mathbf{P}(X_1 = 1) = \varepsilon, \mathbf{P}(X_1 = -1) = 1 - \varepsilon$ with $0 < \varepsilon < 1$ arbitrary. In the latter case, former case, Hoeffding's inequality is fairly sharp, but in the latter case, Hoeffding's inequality is not quite sharp. Put another way, Hoeffding's inequality does not account for the variance of the random variables. Bennett's inequality below does account for the variance of the random variables, with a potentially worse decay than Hoeffding's inequality for large values of t .

Theorem 5.11 (Bennett's Inequality). *Let $c > 0$. Let X_1, X_2, \dots be independent random variables with $X_i \leq c$. For any $s > 0$, define $h(s) := (1 + s) \log(1 + s) - s$. Define $\sigma^2 := \sum_{i=1}^n \text{Var} X_i^2$ and assume $0 < \sigma < \infty$. Then, for any $n \geq 1$,*

$$\mathbf{P}\left(\sum_{i=1}^n X_i - \mathbf{E}\left(\sum_{j=1}^n X_j\right) \geq t\right) \leq e^{-\frac{\sigma^2}{c^2} h\left(\frac{ct}{\sigma^2}\right)} \leq e^{-\frac{t^2}{2+2ct/(3\sigma^2)}}, \quad \forall t \geq 0.$$

Consequently, if $|X_i| \leq c$ for all $1 \leq i \leq n$,

$$\mathbf{P}\left(\left|\sum_{i=1}^n X_i - \mathbf{E}\left(\sum_{j=1}^n X_j\right)\right| \geq t\right) \leq 2e^{-\frac{2t^2}{\sum_{i=1}^n (b_i - a_i)^2}}, \quad \forall t \geq 0.$$

Proof. Without loss of generality, we may assume that $\mathbf{E}X_i = 0$ and $X_i \leq 1$ for all $1 \leq i \leq n$. For any $s \in \mathbb{R}$, define $\phi(s) := e^s - s - 1$. Note that

$$\phi(\alpha) \leq \alpha^2/2, \quad \text{if } s < 0, \quad \phi(\alpha x) \leq \alpha^2 \phi(x), \quad \text{if } x > 0, \alpha \in [0, 1]. \quad (*)$$

The second inequality follows e.g. from the power series expansion of ϕ . Using the definition of ϕ and $\mathbf{E}X_i = 0$ for all $1 \leq i \leq n$, for any $\alpha \in [0, 1]$,

$$\begin{aligned} \mathbf{E}e^{\alpha X_i} &= 1 + \alpha \mathbf{E}X_i + \mathbf{E}\phi(\alpha X_i) = 1 + \mathbf{E}\phi(\alpha X_i) \\ &\stackrel{(*)}{\leq} 1 + \mathbf{E}\phi(\alpha \max(X_i, 0)) + (\alpha^2/2) \mathbf{E}[\max(-X_i, 0)]^2 \end{aligned}$$

Using now (*) and the bound $X_i \leq 1$ for all $1 \leq i \leq n$, and also $\psi(\alpha) \geq \alpha^2/2$ for $\alpha > 0$,

$$\mathbf{E}e^{\alpha X_i} \leq 1 + \phi(\alpha) \mathbf{E}[\max(X_i, 0)]^2 + (\alpha^2/2) \mathbf{E}[\max(-X_i, 0)]^2 \leq 1 + \phi(\alpha) \mathbf{E}X_i^2 \leq e^{\phi(\alpha) \mathbf{E}X_i^2}.$$

The proof now proceeds as in Hoeffding's inequality, Theorem 5.1. For any $t \geq 0$

$$\mathbf{P}\left(\sum_{i=1}^n X_i \geq t\right) \leq e^{-\alpha t} \mathbf{E}e^{\alpha \sum_{i=1}^n X_i} \leq e^{-\alpha t + \phi(\alpha) \sum_{i=1}^n \mathbf{E}X_i^2} = e^{-\alpha t + \phi(\alpha) \sigma^2}.$$

Since $\alpha > 0$ is arbitrary, we choose α to minimize the right side. This minimum occurs when $\frac{t}{\sigma^2} = \phi'(\alpha) = e^\alpha - 1$, so that

$$\alpha = \log\left(1 + \frac{t}{\sigma^2}\right).$$

At this value of α , $\phi(\alpha) = \frac{t}{\sigma^2} - \alpha$, so

$$-\alpha t + \sigma^2 \phi(\alpha) = -t \log\left(1 + \frac{t}{\sigma^2}\right) + t - \sigma^2 \log\left(1 + \frac{t}{\sigma^2}\right) = \frac{t}{\sigma^2} h(t/\sigma^2).$$

, giving the first desired bound. The final bound follows by writing $\mathbf{P}(|\sum_{i=1}^n a_i X_i| \geq t) = \mathbf{P}(\sum_{i=1}^n a_i X_i \geq t) + \mathbf{P}(-\sum_{i=1}^n a_i X_i \geq t)$ and then applying the first inequality twice.

$$h(s) \geq s^2/(2 + 2s/3)$$

□

6. EMPIRICAL RISK MINIMIZATION (ERM) AND CONCENTRATION

Problem 6.1 (Statistical Supervised Learning Problem). Let A, B be sets. Let $f: A \rightarrow B$ be an unknown function. Let \mathbf{P} be an unknown probability distribution on A . The goal of the learning problem is to approximately determine the function f on all of A using a small number of sample values of f on A . Let $X^{(1)}, \dots, X^{(k)}$ be a random sample of size k (i.e. a sequence of independent random variables in A distributed according to \mathbf{P}) and let $Y^{(1)}, \dots, Y^{(k)} \in B$. It is known that

$$f(X^{(i)}) = Y^{(i)}, \quad \forall 1 \leq i \leq k.$$

The goal is to output a deterministic function $g: A \rightarrow B$ that minimizes the prediction error

$$\mathbf{P}(f(X^{(1)}) \neq g(X^{(1)})).$$

Since the probability distribution \mathbf{P} is unknown, it is generally impossible to exactly minimize the prediction error. So, the goal is often restated as minimizing the **empirical risk** or **empirical error** defined as

$$\frac{1}{k} \left| i \in \{1, \dots, k\} : g(X^{(i)}) \neq Y^{(i)} \right|. \quad (*)$$

The task of minimizing the quantity $(*)$ is called **empirical risk minimization (ERM)**. We can equivalently write the empirical error as

$$\frac{1}{k} \sum_{i=1}^k 1_{g(X^{(i)}) \neq Y^{(i)}}.$$

A basic question is then: how close is the empirical error to the true error? For example, given $t > 0$, can we bound

$$\mathbf{P} \left(\left| \frac{1}{k} \sum_{i=1}^k 1_{g(X^{(i)}) \neq Y^{(i)}} - \mathbf{P}(g(X^{(1)}) \neq Y^{(1)}) \right| > t \right)?$$

Since

$$\mathbf{E} \frac{1}{k} \sum_{i=1}^k 1_{g(X^{(i)}) \neq Y^{(i)}} = \mathbf{P}(g(X^{(1)}) \neq Y^{(1)})$$

, we get a bound from Hoeffding's Inequality 5.8, namely

$$\mathbf{P} \left(\left| \frac{1}{k} \sum_{i=1}^k 1_{g(X^{(i)}) \neq Y^{(i)}} - \mathbf{P}(g(X^{(1)}) \neq Y^{(1)}) \right| > t \right) \leq 2e^{-2t^2k}$$

Therefore:

Proposition 6.2. *For any $\varepsilon > 0$ and for any $g: A \rightarrow B$, with probability at least $1 - \varepsilon$, we have*

$$\left| \frac{1}{k} \sum_{i=1}^k 1_{g(X^{(i)}) \neq Y^{(i)}} - \mathbf{P}(g(X^{(1)}) \neq Y^{(1)}) \right| \leq \frac{1}{\sqrt{k}} \sqrt{\log(2/\varepsilon)}.$$

This Proposition is however unsatisfactory. If G minimizes the empirical risk, then G is in fact a random function (since it depends on the random variables $X^{(1)}, \dots, X^{(k)}$). On the other hand, if \bar{g} minimizes the actual risk, then \bar{g} is deterministic. So, in order to compare the minimizers of the empirical and actual risk, we need a bound as in Proposition 6.2 that is *uniform* over all possible minimizers of the empirical risk. For example, suppose that g comes from a class of functions \mathcal{G} . In order to compare the empirical and actual risk, we need to bound

$$\mathbf{P}\left(\sup_{g \in \mathcal{G}} \left| \frac{1}{k} \sum_{i=1}^k 1_{g(X^{(i)}) \neq Y^{(i)}} - \mathbf{P}(g(X^{(1)}) \neq Y^{(1)}) \right| > t\right).$$

If we can show this probability is small, then we can in fact conclude that the empirical risk minimum is close to the actual risk minimum. For example:

Proposition 6.3. *Suppose \mathcal{G} is a finite class of functions from A to B . Let $f: A \rightarrow B$. Let $X^{(1)}, \dots, X^{(k)}$ be a random sample of size k and for any $1 \leq i \leq k$, let $Y^{(i)} := X^{(i)}$. Let G be a random element of \mathcal{G} that minimizes the empirical risk $ER(g) := \frac{1}{k} \sum_{i=1}^k 1_{g(X^{(i)}) \neq Y^{(i)}}$ among all $g \in \mathcal{G}$. Let \bar{g} minimize the risk $\mathbf{P}(g(X^{(1)}) \neq Y^{(1)})$ among all $g \in \mathcal{G}$. Let $\varepsilon > 0$. Then, with probability at least $1 - \varepsilon$,*

$$|ER(g) - \mathbf{P}(\bar{g}(X^{(1)}) \neq Y^{(1)})| \leq \dots$$

Proof. □

6.1. Gaussian Processes.

Theorem 6.4. (Slepian's Lemma) *Let (X_1, \dots, X_n) and (Y_1, \dots, Y_n) be n -dimensional Gaussian random vectors such that $EX_i = EY_i = 0$ for $i = 1, \dots, n$, $EX_i^2 = EY_i^2 = 1$ for $i = 1, \dots, n$. Assume that $EX_i X_j \leq EY_i Y_j$ for $i, j \in \{1, \dots, n\}$, $i \neq j$. Then*

$$P(X_1 > 0, \dots, X_n > 0) \leq P(Y_1 > 0, \dots, Y_n > 0)$$

More generally, for any $\alpha_1, \dots, \alpha_n \in \mathbb{R}$,

$$P(X_1 > \alpha_1, \dots, X_n > \alpha_n) \leq P(Y_1 > \alpha_1, \dots, Y_n > \alpha_n)$$

Proof. Let $\{r_{ij}\}_{i,j=1}^n$ be a symmetric positive definite matrix. Define

$$(2\pi)^{-n/2} |\det r|^{-1/2} \int_0^\infty \dots \int_0^\infty e^{-x^T r^{-1} x/2} dx =: \int_0^\infty \dots \int_0^\infty g(x, r) dx =: f(r)$$

In the special case that $r_{ij} = EX_i X_j$, we have from the definition of the multivariate normal,

$$P(X_1 > 0, \dots, X_n > 0) = (2\pi)^{-n/2} |\det r|^{-1/2} \int_0^\infty \dots \int_0^\infty e^{-x^T r^{-1} x/2} dx$$

From Theorem 6.5,

$$\int e^{-i\langle y, x \rangle} e^{-y^T r y/2} dy = |\det r|^{-1/2} (2\pi)^{n/2} e^{-x^T r^{-1} x/2}$$

So, $g(x, r) = (2\pi)^n \int e^{-i\langle y, x \rangle} e^{-y^T r y / 2} dy$. And for $i \neq j$, $\partial g / \partial r_{ij} = \partial^2 g / \partial x_i \partial x_j$. So, by differentiating under the integral sign and then integrating by parts,

$$\begin{aligned} \frac{\partial f}{r_{12}} &= \int_0^\infty \cdots \int_0^\infty \frac{\partial^2 g}{\partial x_1 \partial x_2} dx \\ &= \int_0^\infty \cdots \int_0^\infty g(0, 0, x_3, \dots, x_n) dx_3 \cdots dx_n \geq 0 \quad (*) \end{aligned}$$

For $\lambda \in [0, 1]$ and $\varepsilon > 0$, let $r_{ij} = \lambda EY_i Y_j - (1 - \lambda) EX_i X_j + \varepsilon \delta_{ij}$. Then r is symmetric positive semidefinite, so (*) and our assumption implies that

$$\frac{\partial f}{\partial \lambda} = \sum_{i \neq j} \frac{\partial f}{\partial r_{ij}} \frac{\partial r_{ij}}{\partial \lambda} = \sum_{i \neq j} \frac{\partial f}{\partial r_{ij}} (EY_i Y_j - EX_i X_j) \geq 0$$

Integrating this inequality for $\lambda \in [0, 1]$ and then letting $\varepsilon \rightarrow 0$ concludes the theorem. \square

Theorem 6.5. (Gaussian Fourier Transform) $\mathcal{F}(e^{-\pi|x|^2}) = e^{-\pi|\xi|^2}$

Proof. It suffices to prove the one dimensional identity $\int_{\mathbb{R}} e^{-\pi(x^2+2ixy)} dx = e^{-\pi y^2}$. In this case we write $\int_{\mathbb{R}} e^{-\pi(x^2+2ixy)} dx = e^{-\pi y^2} \int_{\mathbb{R}} e^{-\pi(x+iy)^2} dx$ and then shift the contour. \square

Theorem 6.6. (Slepian's Inequality) Let (X_1, \dots, X_n) and (Y_1, \dots, Y_n) be n -dimensional Gaussian random vectors such that $EX_i = EY_i = 0$ for $i = 1, \dots, n$, and $EX_i^2 = EY_i^2 = 1$ for $i = 1, \dots, n$. Assume that $E(X_i - X_j)^2 \leq E(Y_i - Y_j)^2$ for $i, j \in \{1, \dots, n\}$, $i \neq j$. Then for all $\alpha \in \mathbb{R}$,

$$P(\sup_{i=1, \dots, n} X_i > \lambda) \leq P(\sup_{i=1, \dots, n} Y_i > \lambda)$$

In particular, $E \sup_{i=1, \dots, n} X_i \leq E \sup_{i=1, \dots, n} Y_i$.

Proof. By our assumption, $EX_i X_j \geq EY_i Y_j$ for $i, j \in \{1, \dots, n\}$, $i \neq j$. For $i = 1, \dots, n$ let $f_i: \mathbb{R} \rightarrow [0, \infty)$ be a non-increasing smooth bounded function. Let $h(x) = \prod_{i=1}^n f_i(x_i)$. For $i \neq j$, $\partial^2 h / \partial x_i \partial x_j = f'_i(x_i) f'_j(x_j) \geq 0$. Let (Z_1, \dots, Z_n) be a mean zero Gaussian random vector with unit variances, and with positive definite covariance matrix r . Let $f(r) = Eh(Z_1, \dots, Z_n)$. As in the proof of Slepian's Lemma, Theorem 6.4, for $i \neq j$,

$$\frac{\partial f}{\partial r_{ij}} = \int h(x) \frac{\partial g}{\partial r_{ij}} dx = \int \frac{\partial^2 h}{\partial x_i \partial x_j} g \geq 0.$$

We therefore conclude that $E \prod_{i=1}^n f_i(X_i) \geq E \prod_{i=1}^n f_i(Y_i)$. Fix $\lambda \in \mathbb{R}$. Let f_i approach $1_{(-\infty, \lambda]}$, so that $P(\sup_{i=1, \dots, n} X_i < \lambda) \geq P(\sup_{i=1, \dots, n} Y_i < \lambda)$, proving the theorem. \square

Theorem 6.7. (Sudakov-Fernique Inequality [Cha05]) Let $X = (X_1, \dots, X_n)$ and $Y = (Y_1, \dots, Y_n)$ be n -dimensional Gaussian random vectors such that $EX_i = EY_i = 0$ for $i = 1, \dots, n$. Assume that $E(X_i - X_j)^2 \leq E(Y_i - Y_j)^2$ for $i, j \in \{1, \dots, n\}$, $i \neq j$. Then

$$E \sup_{i=1, \dots, n} X_i \leq E \sup_{i=1, \dots, n} Y_i$$

Proof. Let g, X_1, \dots, X_n be Gaussians with $Eg^2 = \tau^2$.

$$\begin{aligned} EgF(g) &= \frac{1}{\sqrt{2\pi\tau}} \int_{\mathbb{R}} te^{-t^2/(2\tau^2)} F(t) dt = \frac{\tau^2}{\sqrt{2\pi\tau}} \int_{\mathbb{R}} \left(-\frac{d}{dt}\right) e^{-t^2/(2\tau^2)} F(t) dt \\ &= \frac{\tau^2}{\sqrt{2\pi\tau}} \int_{\mathbb{R}} F'(t) e^{-t^2/(2\tau^2)} dt = Eg^2 EF'(g) \end{aligned}$$

Let $X'_i := X_i - g \frac{EgX_i}{Eg^2}$. Then $EX'_i g = 0$ for $i = 1, \dots, n$, so (X'_1, \dots, X'_n) is independent of g . So, conditioned on (X'_1, \dots, X'_n) , we have

$$\begin{aligned} E(gF(X)|X'_1, \dots, X'_n) &= Eg^2 E\left(\frac{dF}{dg} | X'_1, \dots, X'_n\right) \\ &= \sum_{i=1}^n Eg^2 E\left(\frac{\partial F}{\partial x_i} \frac{\partial x_i}{\partial g} | X'_1, \dots, X'_n\right) \\ &= \sum_{i=1}^n EgX_i E\left(\frac{\partial F}{\partial x_i} | X'_1, \dots, X'_n\right) \end{aligned}$$

So, by integrating out the conditioning,

$$EgF(X) = \sum_{i=1}^n EgX_i E\left(\frac{\partial F}{\partial x_i}\right)$$

In particular,

$$E(X_j F(X)) = \sum_{i=1}^n EX_i X_j E\left(\frac{\partial F}{\partial x_i}\right) \quad (*)$$

Let $F(x) = F_\beta(x) := \frac{1}{\beta} \log(\sum_{i=1}^n e^{\beta x_i})$. We may assume that X and Y are independent. For $0 \leq t \leq 1$, let $Z := \sqrt{1-t}X + \sqrt{t}Y$, and let $f(t) := E(F_\beta(Z))$. By the chain rule, we have

$$f'(t) = E \sum_{i=1}^n \frac{\partial F}{\partial x_i}(Z) \left(\frac{Y_i}{2\sqrt{t}} - \frac{X_i}{2\sqrt{1-t}} \right)$$

Using (*) and the chain rule,

$$E\left(\frac{\partial F}{\partial x_i}(Z) X_i\right) = \sqrt{1-t} \sum_{j=1}^n E(X_i X_j) E\left(\frac{\partial^2 F}{\partial x_i \partial x_j}(Z)\right)$$

$$E\left(\frac{\partial F}{\partial x_i}(Z) Y_i\right) = \sqrt{t} \sum_{j=1}^n E(Y_i Y_j) E\left(\frac{\partial^2 F}{\partial x_i \partial x_j}(Z)\right)$$

So, combining these equalities,

$$f'(t) = \frac{1}{2} \sum_{i,j=1}^n E\left(\frac{\partial^2 F}{\partial x_i \partial x_j}(Z)\right) [E(Y_i Y_j) - E(X_i X_j)]$$

Now, by the definition of F , $\partial F/\partial x_i =: p_i(x) = e^{\beta x_i}/(\sum_{i=1}^n e^{\beta x_i})$. So, for fixed x , the numbers $p_i(x)$, $i = 1, \dots, n$ are nonnegative and sum to 1. Observe that

$$\frac{\partial^2 F}{\partial x_i \partial x_j}(x) = \begin{cases} \beta(p_i(x) - p_i(x)^2) & , i = j \\ -\beta p_i(x)p_j(x) & , i \neq j \end{cases}$$

So,

$$\begin{aligned} & \sum_{i,j=1}^n \frac{\partial^2 F}{\partial x_i \partial x_j}(x)[E(Y_i Y_j) - E(X_i X_j)] \\ &= \beta \sum_{i=1}^n p_i(x)[E(Y_i Y_i) - E(X_i X_i)] - \beta \sum_{i,j=1}^n p_i(x)p_j(x)[E(Y_i Y_j) - E(X_i X_j)] \end{aligned}$$

Also, since $\sum_{i=1}^n p_i(x) = 1$, we have

$$\begin{aligned} & \sum_{i=1}^n p_i(x)[E(Y_i Y_i) - E(X_i X_i)] \\ &= \frac{1}{2} \sum_{i,j=1}^n p_i(x)p_j(x)[E(Y_i Y_i) - E(X_i X_i) + E(Y_j Y_j) - E(X_j X_j)] \end{aligned}$$

So, combining these equalities,

$$\begin{aligned} & \sum_{i,j=1}^n \frac{\partial^2 F}{\partial x_i \partial x_j}(x)[E(Y_i Y_j) - E(X_i X_j)] \\ &= \frac{\beta}{2} \sum_{i,j=1}^n p_i(x)p_j(x)[(E(Y_i Y_i) + E(Y_j Y_j) - 2E(Y_i Y_j)) \\ & \quad - (E(X_i X_i) + E(X_j X_j) - 2E(X_i X_j))] \\ &= \frac{\beta}{2} \sum_{i,j=1}^n p_i(x)p_j(x)[E(Y_i - Y_j)^2 - E(X_i - X_j)^2] \end{aligned}$$

That is, $f'(t) \geq 0$ for all $0 \leq t \leq 1$. So,

$$E(F(Y)) = f(1) \geq f(0) = E(F(X)) \quad (**)$$

Now,

$$\begin{aligned} \max_{i=1,\dots,n} x_i &= \beta^{-1} \log e^{\beta \max_{i=1,\dots,n} x_i} \leq \beta^{-1} \log e^{\beta \sum_{i=1}^n x_i} \\ &\leq \beta^{-1} \log(n e^{\beta \max_{i=1,\dots,n} x_i}) = \beta^{-1} \log n + \max_{i=1,\dots,n} x_i \end{aligned}$$

That is, $\max_{i=1,\dots,n} x_i \leq F_\beta(x) \leq \beta^{-1} \log n + \max_{i=1,\dots,n} x_i$. So, letting $\beta \rightarrow \infty$ in $(**)$ concludes the theorem. \square

Theorem 6.8. (Talagrand’s Majorizing Measures Theorem) Suppose (Y_1, \dots, Y_n) is a Gaussian random vector. Let $X = \{1, \dots, n\}$. For $i, j \in X$, define $d(i, j) := \sqrt{E(Y_i - Y_j)^2}$. Let \mathcal{P}_X be the set of Borel probability measures on X . Define

$$\gamma_2(X, d) := \inf_{\mu \in \mathcal{P}_X} \sup_{x \in X} \int_0^\infty \sqrt{\log \left(\frac{1}{\mu(B(x, r))} \right)} dr$$

Then $c\gamma_2(X, d) \leq E \sup_{i \in X} Y_i \leq C\gamma_2(X, d)$.

6.2. Sub-Gaussian Processes.

Definition 6.9 (Sub-Gaussian Random Variable). A real-valued random variable X is said to be **sub-Gaussian** if there exist $c_1, c_2 > 0$ such that

$$\mathbf{P}(|X| > t) \leq c_1 e^{-c_2 t^2}, \quad \forall t > 0.$$

Exercise 6.10. Let X be a real-valued random variable with mean zero. Then the following are equivalent

- $\exists a > 0$ such that, for all $t > 0$, $\mathbf{E}e^{tX} \leq e^{t^2 a^2 / 2}$.
- $\exists b > 0$ such that, for all $t > 0$, $\mathbf{P}(|X| > t) \leq 2e^{-bt^2}$.
- $\exists c > 0$ such that $\mathbf{E}e^{cX^2} \leq 2$.

For all $t \in \mathbb{R}$, define

$$\psi_2(t) := e^{t^2} - 1.$$

For a random variable X define

$$\|X\|_{\psi_2} := \inf\{t > 0: \mathbf{E}\psi_2(|X|/t) \leq 1\}.$$

6.3. General Empirical Processes.

6.4. Additional Comments.

7. CONCENTRATION OF EMPIRICAL PROCESSES

8. DEEP LEARNING

Theorem 8.1 (PAC Learning Neural Networks).

9. APPENDIX: BASICS OF COMPLEXITY THEORY

A Turing machine is a standard model of computation introduced by Turing in 1937. Informally, a one-tape Turing machine is a computing device with a finite state control device, a tape (with cells indexed by the positive integers), and a tape head that points to (or scans) a given cell on the tape. At a given time (indexed by the positive integers), the machine changes its state, writes a symbol on the cell to which is currently pointing, and then moves the tape head one cell to the left or one cell to the right. Time is then increased by one. The action of the machine at a given time is a function of its current state and of the symbol that is currently scanned by the tape head. More formally:

Definition 9.1 (Turing Machine). A **one-tape Turing machine** is defined by

- Σ := a finite set (“alphabet”), together with a blank symbol $\{\square\}$.
- Q := a finite set of “control states” $\cup \{q_{\text{accept}}, q_{\text{reject}}, q_{\text{start}}\}$.

- $\delta: \Sigma \times Q \rightarrow \Sigma \times Q \times \{\leftarrow, \rightarrow\}$, the “transition function.”

A **configuration** of a Turing machine consists of

- the symbols on the tape up to and including its rightmost non-blank symbol,
- the current state of the control device (an element of Q), and
- the position of the tape head (an element of the positive integers).

So, if $i + 1, n$ are positive integers, if $\sigma_1, \dots, \sigma_n \in \Sigma$, and if $q \in Q$, the string

$$\sigma_1 \sigma_2 \cdots \sigma_i q \sigma_{i+1} \sigma_{i+2} \cdots \sigma_n$$

represents the configuration of the Turing machine where the finite state control is in state q , the tape head is currently scanning position $i + 1$ on the tape, and the contents of the tape are $\sigma_1 \cdots \sigma_n$.

Let $x_1, \dots, x_n \in \Sigma \setminus \{\square\}$ and let $x = x_1 x_2 \cdots x_n$ be the corresponding string. A Turing machine then **computes** on the input x by doing the following.

- The Turing machine is initialized to the initial configuration $C_0 := q_{\text{start}} x_1 x_2 \cdots x_n$.
- At any time step $i \geq 1$, the Turing machine applies the transition function δ to the previous configuration C_{i-1} to obtain the next configuration C_i . So, if $C_{i-1} = \sigma_1 \sigma_2 \cdots \sigma_i q \sigma_{i+1} \sigma_{i+2} \cdots \sigma_n$, we have $\delta(\sigma_{i+1}, q) =: (\sigma', q', a)$ for some $\sigma' \in \Sigma$, $q' \in Q$ and $a \in \{\leftarrow, \rightarrow\}$. We then define

$$C_i := \begin{cases} \sigma_1 \sigma_2 \cdots \sigma_{i-1} q' \sigma_i \sigma' \sigma_{i+2} \sigma_{i+3} \cdots \sigma_n & , \text{ if } a = \leftarrow \\ \sigma_1 \sigma_2 \cdots \sigma_{i-1} \sigma_i \sigma' q' \sigma_{i+2} \sigma_{i+3} \cdots \sigma_n & , \text{ if } a = \rightarrow . \end{cases}$$

- If at any time step $k \geq 1$ the Turing machine enters a halting state (q_{accept} or q_{reject}), we say the machine **halts** and **accepts** (or **rejects**) input x in k steps.

If the machine halts in k steps, the computation of the machine on input x can be described as a sequence of configurations C_0, C_1, \dots, C_k .

Remark 9.2. Modern computers do not directly implement Turing machines. Modern computers more closely resemble RAM machines, or von Neumann machines.

Remark 9.3. A multi-tape Turing Machine is defined similarly to the above, except it has one single input tape (where only the input state is written), one single output tape, a constant number of other (“work”) tapes, the function δ is then a function of all of the tapes and the current state of the finite control. When the configuration changes, each work tape head overwrites the cell it is currently scanning, and each work tape head moves left or right one cell.

Let Σ be a finite set (or “alphabet”) and let Σ^* denote the set of finite strings of elements of Σ .

Definition 9.4 (Decided). A set (or “language”) $S \subseteq \Sigma^*$ is **recognized** (or **decided**) by a Turing Machine M if:

- If $x \in S$, then $M(x)$ accepts.
- If $x \notin S$, then $M(x)$ rejects.

Definition 9.5 (Accepted). A set (or “language”) $S \subseteq \Sigma^*$ is **accepted** by a Turing Machine M if:

- If $x \in S$, then $M(x)$ accepts.

- If $x \notin S$, then $M(x)$ does not accept.

Let $\mathbb{N} := \{0, 1, 2, \dots\}$. For any $x \in \Sigma^*$, let $|x|$ denote the length of the string x .

Definition 9.6 (Time Complexity). Let $f: \mathbb{N} \rightarrow \mathbb{N}$. We say a set $S \subseteq \Sigma^*$ satisfies $S \in \text{TIME}(f(n))$ if there exists a multitape Turing Machine M such that

- M recognizes S , and
- For all $x \in S$, $M(x)$ halts within $f(|x|)$ steps.

Definition 9.7 (Space Complexity). Let $f: \mathbb{N} \rightarrow \mathbb{N}$. We say a set $S \subseteq \Sigma^*$ satisfies $S \in \text{SPACE}(f(n))$ if there exists a multitape Turing Machine M such that

- M recognizes S , and
- For all $x \in S$, $M(x)$ uses no more than $f(|x|)$ squares of its *work tapes*. That is, $M(x)$ scans at most $f(|x|)$ cells of its work tapes.

Definition 9.8 (Time Complexity for Functions). Let $f: \mathbb{N} \rightarrow \mathbb{N}$. We say a function $g: \Sigma^* \rightarrow \Sigma^*$ satisfies $g \in \text{FTIME}(f(n))$ if there exists a multitape Turing Machine M such that

- For all $x \in \Sigma^*$, $M(x)$ writes $g(x)$ on its output tape, and
- $M(x)$ halts within $f(|x|)$ steps.

Definition 9.9 (Space Complexity for Functions). Let $f: \mathbb{N} \rightarrow \mathbb{N}$. We say a function $g: \Sigma^* \rightarrow \Sigma^*$ satisfies $g \in \text{FSPACE}(f(n))$ if there exists a multitape Turing Machine M such that

- For all $x \in \Sigma^*$, $M(x)$ writes $g(x)$ on its output tape, and
- $M(x)$ uses no more than $f(|x|)$ squares of its *work tapes*.

Definition 9.10 (Complexity Class P). We define

$$\mathbf{P} := \cup_{k \geq 1} \text{TIME}(n^k).$$

Definition 9.11 (Complexity Class NP). \mathbf{NP} is the class of sets $S \subseteq \Sigma^*$ such that there exists a Turing Machine V (a “verifier”) and a polynomial $p: \mathbb{N} \rightarrow \mathbb{N}$ such that $x \in S$ if and only if: \exists a string $y \in \Sigma^*$ of length at most $p(|x|)$ such that $V(x, y)$ accepts in time at most $p(|x|)$.

Definition 9.12 (NP-hard and NP-complete). A set $S \subseteq \Sigma^*$ is \mathbf{NP} -hard if for all $T \in \mathbf{NP}$, there exists a logspace function $g: \Sigma^* \rightarrow \Sigma^*$ such that $x \in S$ if and only if $g(x) \in T$. And S is \mathbf{NP} -complete if $S \in \mathbf{NP}$ and S is \mathbf{NP} -hard.

The SAT problem is the following. Let n

Definition 9.13 (Satisfiability Problem SAT). Let $f: \{-1, 1\}^n \rightarrow \{-1, 1\}$ be an unknown function with query access. Decide whether or not there exists $x_1, \dots, x_n \in \{-1, 1\}$ such that $f(x_1, \dots, x_n) = 1$.

Theorem 9.14 (Cook-Levin, 1971). SAT is \mathbf{NP} -complete.

\mathbf{NP} can be equivalently defined using non-deterministic Turing Machines. A **non-deterministic Turing machine** is defined by repeating the definition of a Turing machine, except the function δ is allowed to take multiple values. Then the computation of a non-deterministic Turing machine can be viewed as a tree of configurations, rather than a sequence of configurations,

since one configuration can transition to multiple different configurations. More specifically, we associate a directed edge from configuration C to configuration C' if δ maps configuration C to configuration C' . A non-deterministic Turing machine accepts $x \in \Sigma^*$ in time $t > 0, t \in \mathbb{Z}$ if there exists a path from the initial configuration to an accepting configuration of length at most t . A non-deterministic Turing machine M **accepts** a set $S \subseteq \Sigma^*$ if: $x \in S$ if and only if M accepts x .

Intuitively, a non-deterministic Turing machine is a Turing machine that is allowed to make arbitrary choices at each step of its computation.

A nondeterministic Turing machine can solve SAT in linear time using a binary tree of configurations (Exercise).

Definition 9.15 (Non-deterministic Time Complexity). Let $f: \mathbb{N} \rightarrow \mathbb{N}$. We say a set $S \subseteq \Sigma^*$ satisfies $S \in \text{NTIME}(f(n))$ if there exists a nondeterministic Turing Machine M such that

- M recognizes S , and
- For all $x \in S$, when M has input x , the computation of M has no path longer than $f(|x|)$.

An equivalent definition of **NP** is then

$$\mathbf{NP} := \cup_{k \geq 1} \text{NTIME}(n^k).$$

Definition 9.16 (Complexity Class #P). A function $f: \{0, 1\}^* \rightarrow \mathbb{N}$ is in **#P** if \exists a polynomial $p: \mathbb{N} \rightarrow \mathbb{N}$ and a polynomial time Turing Machine M such that, for all $x \in \{0, 1\}^*$,

$$f(x) = \left| \left\{ y \in \{0, 1\}^{p(|x|)} : M(x, y) = 1 \right\} \right|.$$

Equivalently, f is in **#P** if there is a polynomial time non-deterministic Turing machine M such that, for all $x \in \{0, 1\}^*$, $f(x)$ is equal to the number of accepting paths of M on input x .

Informally, **#P** problems count the number of solutions to problems in **NP**.

Definition 9.17 (Complexity Class RP). **RP** is the class of sets $S \subseteq \Sigma^*$ such that there exists a non-deterministic polynomial time Turing Machine M such that

- M accepts S , and
- If $x \in S$, then at least $1/2$ of all computation paths of $M(x)$ accept.

Informally, **RP** is a Turing machine that can use randomness at each of its computation steps.

10. APPENDIX: SOME FUNCTIONAL ANALYSIS

Below, we consider either \mathbb{R} or \mathbb{C} as scalars. That is, we will be dealing with vector spaces over the fields \mathbb{R} or \mathbb{C} , and we will only distinguish them where necessary. A **topological vector space** is a Hausdorff topological space X that is also a vector space, such that: the map $(x, y) \mapsto x - y$ from $X \times X \rightarrow X$ is continuous, and the map $(\alpha, x) \mapsto \alpha x$ from $\{\text{scalars}\} \times X \rightarrow X$ is continuous.

A **normed linear space** X is a vector space (over \mathbb{R} or \mathbb{C}) with a **norm** $\|\cdot\|$. A norm is a function $\|\cdot\|: X \rightarrow [0, \infty)$ such that $\|x\| \geq 0$ with equality if and only if $x = 0$, $\|\alpha x\| = |\alpha| \|x\|$ for α a scalar, and $\|x + y\| \leq \|x\| + \|y\|$. Using the norm, we see that

$d(x, y) := \|x - y\|$ is a metric, whose open balls define the metric topology on X . We refer to this topology as the norm topology, or strong topology. Using the triangle inequality, one can show that a normed linear space is also a topological vector space. A **Banach space** is a normed linear space that is complete with respect to the norm topology. A vector subspace $Y \subseteq X$ is simply called a **subspace**, and will be closed unless otherwise stated. Also, unless otherwise stated, \overline{A} denotes the closure of A with respect to the norm topology.

An **inner product** is a function $(\cdot, \cdot): X \times X \rightarrow \{\text{scalars}\}$ that is linear in the first argument, conjugate linear in the second argument, Hermitian symmetric, and positive definite. An inner product space is a vector space with an inner product. Defining $\|x\| := (x, x)^{1/2}$ shows that an inner product space is a normed linear space (using the Cauchy-Schwarz inequality). A **Hilbert space** is an inner product space that is complete in the norm topology. Cauchy-Schwarz shows that the inner product is continuous with respect to the norm topology, since

$$|(u, v) - (u_0, v_0)| \leq |(u - u_0, v)| + |(u_0, v - v_0)| \leq \|u - u_0\| \|v\| + \|v - v_0\| \|u_0\|$$

One can show that $\|x - y\|^2 = \|x\|^2 - 2\Re(x, y) + \|y\|^2$. In Hilbert space we have the **parallelogram law** $\|u + v\|^2 + \|u - v\|^2 = 2\|u\|^2 + 2\|v\|^2$, and the **polarization identity** $(u, v) = \frac{1}{4} \sum_k i^k \|u + i^k v\|^2$, where $k = 0, 2$ for X over \mathbb{R} , and $k = 0, 1, 2, 3$ for X over \mathbb{C} . Actually, polarization holds in a Banach space if and only if it is a Hilbert space.

A continuous linear function $L: (X, \|\cdot\|_X) \rightarrow (Y, \|\cdot\|_Y)$ is called a **linear operator**. One can show that L is uniformly continuous if and only if it is continuous if and only if it is continuous at zero if and only if it is **bounded**, i.e. $\|L(x)\|_Y \leq M \|x\|_X$. Showing that continuity at zero implies boundedness involves a scaling argument. Showing that boundedness implies uniform continuity follows by linearity.

The least such M such that $\|L(x)\|_Y \leq M \|x\|_X$ is called the **operator norm** of L , and it is denoted by $\|L\|$. Note that the set of bounded linear maps $B(X, Y)$ from X to Y is a normed linear space with respect to the operator norm topology. If Y is complete, then $B(X, Y)$ is a Banach space. (Fix $x \in X$, observe $L_n(x) \rightarrow L(x)$, L achieves linearity using subsequential arguments, etc.) Essentially by definition, we have $\|L\| := \sup_{\|x\| \leq 1} \|L(x)\|$. If X, Y are Hilbert spaces, $\|L\| = \sup_{\|x\| \leq 1, \|y\| \leq 1} |\langle L(x), y \rangle|$. For $L: H \rightarrow H$ a linear operator on a Hilbert space H , we can apply Thm. 10.2 to define the **adjoint** L^* . Here $L^*(u) = v$ is the unique v such that $(L(w), u) = (w, L^*(u))$. Moreover, $\|L\| = \|L^*\|$ (using the formula for $\|L\|$ from a few previous lines). More directly, we can observe that $\|A^*y\|^2 = (AA^*y, y) \leq \|A\| \|A^*y\| \|y\|$, so $\|A^*\| \leq \|A\|$. Then, using that $A^{**} = A$ (take complex conjugates of the definition of A^*) we see that $\|A\| \leq \|A^*\|$.

A **linear functional** on X is a bounded map from X to scalars (\mathbb{C} or \mathbb{R}). The space of linear functionals is called the **dual space** of X , and is denoted by X^* . The norm of $x^* \in X^*$ is given by $\|x^*\| := \sup_{\|x\| \leq 1} |x^*(x)|$.

A **semi-norm** N on a vector space X is a function $N: X \rightarrow \mathbb{R}$ such that: $\forall x \in X, N(x) \geq 0, N(\alpha x) = |\alpha| N(x)$, and $N(x + y) \leq N(x) + N(y)$. A collection \mathcal{N} of seminorms is called **separating** if $N(x) = 0 \forall N \in \mathcal{N}$ if and only if $x = 0$. For a seminorm N , define $S_N(a, \rho) := \{x \in X: N(x - a) < \rho\}$. $S_N(a, \rho)$ is called the **open N ball** of radius ρ centered at a . A set of the form $S = S_{N_1}(a, \rho_1) \cap \cdots \cap S_{N_j}(a, \rho_j)$ is called an **open \mathcal{N} ball** centered at a . Given X and \mathcal{N} , let $X^{\mathcal{N}}$ denote X with the topology given by open \mathcal{N} balls. One can check that $X^{\mathcal{N}}$ is Hausdorff if and only if \mathcal{N} is separating, and $X^{\mathcal{N}}$ is a topological vector

space when \mathcal{N} is separating. Sometimes, $X^{\mathcal{N}}$ is called a **Fréchet space**. For example, consider the **Schwartz class**, $\mathcal{S}(\mathbb{R}) = \{f \in C^\infty : x^n f^{(k)} \in L_\infty, \forall n, k \geq 0\}$, equipped with the seminorms $\|f\|_{n,k} = \sup_{x \in \mathbb{R}} |x^n f^{(k)}(x)|$. Here \mathcal{N} is separating since $\|f\|_{0,0} = 0$ implies $f = 0$.

A topological vector space is called **locally convex** if X has a neighborhood base at zero consisting of convex sets. That is, for any neighborhood U of 0, $\exists V$ convex such that $0 \in V \subseteq U$. It takes some work to show: if X is a topological vector space, and \mathcal{N} is a family of seminorms, then $X^{\mathcal{N}} = X$ if and only if X is locally convex. As an example of the above we define the **weak topology** of X , denoted by $\sigma(X, X^*)$, as the topology on X induced by the seminorms $x \mapsto \|L(x)\|$, $L \in X^*$. In particular, we can form a basis of neighborhoods of x_0 using the seminorms:

$$U(x_0; \rho, x_1^*, \dots, x_n^*) := \{x \in X : |x_j^*(x) - x_j^*(x_0)| < \rho \text{ for } j = 1, \dots, n\}$$

Note that $x_n \rightarrow x$ weakly (i.e. with respect to $\sigma(X, X^*)$) if and only if $\forall x^* \in X^*$, $x^*(x_n) \rightarrow x^*(x)$. A sequence $\{x_n\}_{n \geq 1}$ such that $x^*(x_n)$ converges for all $x^* \in X^*$ is called **weakly Cauchy**.

Analogously, we can define the **weak* topology** on X^* , denoted by $\sigma(X^*, X)$. Here a basis of neighborhoods of $x_0^* \in X^*$ is given by

$$U(x_0^*, \varepsilon, x_1, \dots, x_n) := \{x^* \in X^* : |x^*(x_j) - x_0^*(x_j)| < \varepsilon \text{ for } j = 1, \dots, n\}$$

For X a normed linear space, define the **canonical map** $\iota: X \rightarrow X^{**}$ by $\iota(x)(x^*) = x^*(x)$. One can check that ι is linear, and Thm. 10.5 shows that ι is an isometry:

$$\|\iota(x)\| = \sup_{\|x^*\| \leq 1} |\iota(x)(x^*)| = \sup_{\|x^*\| \leq 1} |x^*(x)| = \|x\|$$

In particular, $\iota: (X, \sigma(X, X^*)) \rightarrow (\iota(X), \sigma(X^{**}, X^*))$ is a homeomorphism, from X with the weak topology, to $\iota(X) \subseteq X^{**}$ with the weak* topology (of X^{**} acting on X^*). A Banach space is said to be **reflexive** if $\iota: X \rightarrow X^{**}$ is onto.

Let $K \subseteq X$ with X a topological vector space. A point $x \in K$ is called an **extreme point** if it is not contained in the interior of any **segment** $\{ty + (1-t)z : t \in [0, 1], y, z \in K\} \subseteq K$. By interior we mean $\{ty + (1-t)z : t \in (0, 1), y, z \in K\}$. If K is compact and convex, then a **face** $F \subseteq K$ is a subset such that: if $y \in K$ is in the interior of a segment in K , then the whole segment is in F . Let K be a nonempty compact, convex set. Denote the set of extreme points of K by $\mathcal{E}(K)$. Define \widehat{K} as the intersection of all compact, convex sets containing $\mathcal{E}(K)$. By definition, $\widehat{K} \subseteq K$ and $\widehat{K} \neq \emptyset$. An **extreme face** F of K is a face of K where the only faces of F are \emptyset and F .

Let μ be a probability measure. A subset $H \subseteq L_1(\mu)$ is called **uniformly integrable** if, for all $\varepsilon > 0$, $\exists \eta > 0$ with $\sup \left\{ \int_A |f| d\mu : \mu(A) \leq \eta, f \in H \right\} \leq \varepsilon$. (Uniform integrability has a slightly different definition in the probability literature.)

An **associative algebra** \mathcal{A} over a field F is a vector space over F with a bilinear, associative multiplication: $(ab)c = a(bc)$, $a(b+c) = ab+ac$, $(a+b)c = ac+bc$, and $a(\lambda c) = (\lambda a)c = \lambda(ac)$. A **Banach algebra** is a real or complex Banach space that is an associative algebra such that $\|ab\| \leq \|a\| \|b\|$. Examples include: $(C(X), \mathbb{C})$ for X a topological space, $B(V, V)$, and $L_1(\mathbb{R}^n)$ with convolution. Note that the latter has no identity. However, given a (complex) Banach algebra without identity, if we let $B := \{(a, \alpha) : a \in \mathcal{A}, \alpha \in \mathbb{C}\} = \mathcal{A} \oplus \mathbb{C}$,

and we define $(a, \alpha)(b, \beta) := (ab + \alpha b + \beta a, \alpha\beta)$ and $\|(a, \alpha)\| := \|a\| + |\alpha|$, then B has identity $e = (0, 1)$.

Let \mathcal{A} be a Banach algebra with identity. Using power series (and that $\|a^n\| \leq \|a\|^n$), we see: if $\|a\| < 1$, then $1 - a$ is invertible (let $(1 - a)^{-1} := \sum_{n \geq 0} a^n$, so $\|(1 - a)^{-1}\| \leq \frac{1}{1 - \|a\|}$). Using similar arguments, we see that: the **invertible elements** \mathcal{U} form an open set, and the inverse map is continuous from \mathcal{U} to itself. For $a \in \mathcal{U}$ and $\|x - a\| < \|a^{-1}\|^{-1}$, note that $\|a^{-1}x - 1\| = \|a^{-1}(x - a)\| \leq \|a^{-1}\| \|x - a\| < 1$, so $1 - (1 - a^{-1}x) = a^{-1}x$ has inverse b , so $1 = (ba^{-1})x$ and $a^{-1}xb = 1$, so $xb = a$, $x(ba^{-1}) = 1$, i.e. $(ba^{-1}) = x^{-1}$.

From now on, we assume that all Banach algebras \mathcal{A} are complex and have an identity. Let $x \in \mathcal{A}$, and for convenience, λ denotes $\lambda \cdot 1 \in \mathcal{A}$. The **spectrum** of x is $\sigma(x) := \{\lambda \in \mathbb{C} : x - \lambda \text{ is not invertible}\}$. The **resolvent set** of x is $\rho(x) := \{\lambda \in \mathbb{C} : x - \lambda \text{ is invertible}\}$. The **resolvent** of x is the function $R(\lambda) = (x - \lambda)^{-1}$, defined for $\lambda \in \rho(x)$. The **spectral radius** is $r(x) := \sup\{|\lambda| : \lambda \in \sigma(x)\}$.

From our analysis of the invertible elements, we see that $r(a) \leq \|a\|$. Thus, $\sigma(a)$ is bounded. In fact, it is compact, since $\rho(a)$ is open ($\lambda \in \mathbb{C} \mapsto a - \lambda \in \mathcal{A}$ is continuous, and $\rho(a) = \{\lambda : a - \lambda \in \mathcal{U}\}$), so $\sigma(a) = \rho(a)^c$ is closed. Also, $\sigma(a) \neq \emptyset$. To see this, define a **weakly analytic** function as a function ϕ from an open set $V \subseteq \mathbb{C}$ to a complex Banach space \mathcal{A} such that $\xi \circ \phi$ is analytic for every $\xi \in \mathcal{A}^*$. Using power series, one can show: $R(\lambda) = (a - \lambda)^{-1}$ is weakly analytic on $\rho(a)$ and $\|R(\lambda)\| \rightarrow 0$ as $\lambda \rightarrow \infty$. So, if $\sigma(a)$ is empty, Liouville's Theorem (for bounded analytic functions) says that $\xi((a - \lambda)^{-1}) \equiv 0$, so $(a - \lambda)^{-1} \equiv 0 \forall \lambda$, which is a contradiction.

For the power series argument, let $\lambda_0 \in \rho(a)$. Write $a - \lambda = (1 - \lambda_0)(a - (a - \lambda_0)^{-1}(\lambda - \lambda_0))$. Then $a - \lambda$ is invertible if $\|(a - \lambda_0)^{-1}(\lambda - \lambda_0)\| < 1$. Choose λ so the latter condition holds. Then $(a - \lambda)^{-1} = \sum_{n \geq 0} (a - \lambda_0)^{-n} (\lambda - \lambda_0)^n (a - \lambda_0)^{-1}$. So $\xi((a - \lambda)^{-1}) = \sum_{n \geq 0} \xi((a - \lambda_0)^{-n-1}) (\lambda - \lambda_0)^n$, i.e. $\xi(R(\lambda))$ is analytic (in particular, for $\lambda \in \rho(a)$.) Now, $(a - \lambda)^{-1} = [\lambda(\lambda^{-1}a - 1)]^{-1} = \lambda^{-1}(\lambda^{-1}a - 1)^{-1}$ and $\|(\lambda^{-1}a - 1)^{-1}\| \leq (1 - |\lambda|^{-1}\|a\|)^{-1} \rightarrow 1$ as $\lambda \rightarrow \infty$, so $\|R(\lambda)\| \rightarrow 0$ as $\lambda \rightarrow \infty$.

Below, \mathcal{B} denotes a commutative Banach algebra with identity. A **character** α of \mathcal{B} is a nonzero multiplicative linear function on \mathcal{B} , i.e. $\alpha(ab) = \alpha(a)\alpha(b)$. (A character is not assumed to be bounded a priori, but this can be proven. See Thm. 10.23(9).) The **spectrum** of \mathcal{B} is the set $\tilde{\mathcal{B}}$ of all characters of \mathcal{B} . For $a \in \mathcal{B}$, $\alpha \in \tilde{\mathcal{B}}$ define $\hat{a}(\alpha) := \alpha(a)$. The map $a \mapsto \hat{a}$ from \mathcal{B} into $C(\tilde{\mathcal{B}})$ is called the **Gelfand map** or canonical map.

An **ideal** $I \subseteq \mathcal{B}$ is closed under addition (within I) and multiplication by elements of \mathcal{B} . An ideal is called **maximal** if $I \neq \mathcal{B}$ and I is not contained in any larger proper ideal. Define the **singular elements** \mathcal{S} as the union of all maximal ideals in \mathcal{B} . Define the **radical** of \mathcal{B} as the intersection of all maximal ideals in \mathcal{B} . Since $I \subseteq \bar{I} \subseteq \mathcal{S}$ is a proper ideal, we see that maximal ideals are closed, so the radical is closed. \mathcal{B} is called **semisimple** if its radical is $\{0\}$.

An **involution** on a Banach algebra \mathcal{B} is a map $\mathcal{B} \rightarrow \mathcal{B}$ written $a \mapsto a^*$ such that: $a^{**} = a$, $(a + b)^* = a^* + b^*$, $(\lambda a)^* = \bar{\lambda}a^*$, and $(ab)^* = b^*a^*$. Note that $1^* = 1$, since applying $*$ to $1 \cdot 1^* = 1^*$ gives $1^{**} \cdot 1^* = 1^{**}$, so since $1^{**} = 1$, we get $1 \cdot 1^* = 1$, so $1^* = 1$. We say a is **Hermitian** if $a = a^*$, **strongly positive** if $a = b^*b$ for some b , **positive** if $\sigma(a) \subseteq [0, \infty)$ and **real** if $\sigma(a) \subseteq \mathbb{R}$. An involution is **symmetric** if $1 + a^*a$ is invertible for all $a \in \mathcal{B}$.

A Banach $*$ algebra \mathcal{B} is called **$*$ multiplicative** if $\|a^*a\| = \|a^*\| \|a\|$, **$*$ isometric** if $\|a^*\| = \|a\|$, and **$*$ quadratic** if $\|a^*a\| = \|a\|^2$. One can show that the first two together

are equivalent to the third. The forward direction is clear. For the reverse, observe that $\|a\|^2 = \|a^*a\| \leq \|a^*\| \|a\|$, so $\|a\| \leq \|a^*\|$ (and also for $a = b^*$), so $\|a\| = \|a^*\|$. Also, $\|a^*a\| = \|a\|^2 = \|a^*\| \|a\|$. A B^* **algebra** is a quadratic $*$ algebra. (In modern terminology a C^* algebra is a multiplicative $*$ algebra. From the equivalences just shown, a B^* algebra is also a C^* algebra, and no one uses the term B^* algebra anymore).

Let $A, B \in B(H)$, i.e. let A, B be bounded linear operators on a Hilbert space H . Let $*$ denote the adjoint operation. Summarizing some properties of $B(H)$, we have: A^* is linear and bounded with $\|A^*\| = \|A\|$, $A^{**} = A$, $(\alpha A + \beta B)^* = \bar{\alpha}A^* + \bar{\beta}B^*$, $\|AB\| \leq \|A\| \|B\|$, and $(AB)^* = B^*A^*$. Also note that $\|A^*A\| = \|A\|^2$. This follows since $\|A^*A\| \leq \|A^*\| \|A\| = \|A\|^2$, and $\|Ax\|^2 = (A^*Ax, x) \leq \|A^*A\| \|x\|^2$. From above, we saw that $B(X)$ is a Banach algebra whenever X is a Banach space. Thus, $B(H)$ is a B^* algebra.

In these notes, a C^* -**algebra** on a Hilbert space H is a subalgebra \mathcal{A} of $B(H)$ which is closed in norm and such that $A \in \mathcal{A}$ implies $A^* \in \mathcal{A}$. A subalgebra closed under taking adjoints is called a $*$ **subalgebra** of $B(H)$. Note that $B(H)$ is a C^* -algebra. A **maximal abelian self-adjoint (m.a.s.a.) algebra** on H is a commutative algebra $\mathcal{A} \subseteq B(H)$ which is not contained in any larger commutative subalgebra, and such that \mathcal{A} is a $*$ -subalgebra. Let $S \subseteq B(H)$. Define $S' = \{A \in B(H) : AB = BA, \forall B \in S\}$. S' is then a subalgebra of $B(H)$ for any set S , and we call S' the **commutator algebra** of S .

Let (X, μ) be a measure space. Let $f \in L_\infty(\mu)$. Define $M_f: L_2(\mu) \rightarrow L_2(\mu)$ by $M_f(g) := fg$. Since $fg \in L_2$ for $g \in L_2$, M_f is everywhere defined, and $\|M_f g\|_2^2 \leq \|f\|_\infty^2 \|g\|_2^2$. Therefore, $\|M_f\| \leq \|f\|_\infty$. Note also that $M_{fg} = M_f M_g$, $M_{\alpha f + \beta g} = \alpha M_f + \beta M_g$, $M_{\bar{f}} = M_f^*$.

Unless otherwise specified, in this Section for a measure space (X, μ) , we assume: every measurable set in X of positive measure contains a subset of finite strictly positive measure. (That is, μ has no infinite atoms). Under this assumption, $\|M_f\| = \|f\|_\infty$, which can be proven by considering the indicator function where $|f| = a$ for $\|f\|_\infty > a > 0$ (if such an a exists). The **multiplication algebra**, denoted by $\mathcal{M}(X, \mu)$, of (X, μ) is the algebra of operators on $L_2(X, \mu)$ consisting of all $M_f, f \in L_\infty$.

Let $D(w, \varepsilon) = \{z \in \mathbb{C} : |z - w| < \varepsilon\}$. If $f \in L_\infty(X, \mu)$, define the **essential range** of f as $\{w \in \mathbb{C} : \mu(f^{-1}(D(w, \varepsilon))) > 0 \text{ for all } \varepsilon > 0\}$. For (X, μ) with no infinite atoms, an exercise shows: $\sigma(M_f) = \text{essential range of } f$. Let \mathcal{A} be a subalgebra of $B(H)$. A vector $x \in H$ is called a **cyclic vector** for \mathcal{A} if $Ax := \{Ax : A \in \mathcal{A}\}$ is dense in H . A **unitary operator** $U: H \rightarrow K$ between two Hilbert spaces is a linear surjective operator such that $\|Ux\| = \|x\| \forall x \in H$. To emphasize the surjectivity, we write $U: H \twoheadrightarrow K$.

A bounded operator $A: H \rightarrow H$ is called: **normal** if $A^*A = AA^*$, **Hermitian** if $A = A^*$, **unitary** if A is onto and $\|Ax\| = \|x\| \forall x \in H$, and **orthogonal** if H is real and A is unitary. Suppose H is a Hilbert space, $A: H \rightarrow H$ is linear, and $(Ax, x) = 0$ for all $x \in H$. Then (a) if H is complex, then $A = 0$ (b) if H is real and $A^* = A$ then $A = 0$. To prove this, we observe:

$$(A(x+y), x+y) - (A(x-y), (x-y)) = 2(Ax, y) + 2(Ay, x)$$

So $(Ax, y) + (Ay, x) = 0$. If H is real and $A = A^*$, this yields $(Ax, y) = 0$ for all x, y , so $A = 0$. For H complex, we can get the same conclusion by writing $(Ax, y) + (Ay, x) = 0$ and substituting $x \mapsto ix$ to get $i(Ax, y) - i(Ay, x) = 0$, etc.

Using the above we see: a linear function is **unitary** if and only if U is bounded and $UU^* = U^*U = id$. If the latter holds, then $\|Ux\|^2 = (U^*Ux, x) = \|x\|^2$, and $U(U^*x) = x$, so U is onto. In the reverse direction, if U is unitary, then $((U^*U - id)x, x) = \|Ux\|^2 - \|x\|^2 = 0$,

so from the above we see that $U^*U - id = 0$. Using surjectivity of U (given x , let y such that $Uy = x$) we have $UU^*x = UU^*Uy = Uy = x$, so $UU^* = id$.

Let (X, μ) be a σ -finite measure space, and let $f \in L_\infty$. Using that $M_f^* = M_{\bar{f}}$, one can see that: (1) M_f is normal, (2) M_f is Hermitian if and only if f is real a.e. (i.e. $f = \bar{f}$ a.e.), and (3) M_f is unitary if and only if $|f| = 1$ a.e. (i.e. $f\bar{f} = 1$ a.e.). A sequence A_n of operators on a Banach space B **converges strongly** to a bounded operator A if $A_n x \rightarrow Ax$ for each $x \in B$. A_n **converges weakly** if $\langle A_n x, y \rangle \rightarrow \langle Ax, y \rangle \forall x \in B, y \in B^*$. If B is a Hilbert space, weak convergence is therefore: $\langle A_n x, y \rangle \rightarrow \langle Ax, y \rangle \forall x, y \in H$.

Let X be a set and let \mathcal{S} be a σ -field in X . A projection valued measure on \mathcal{S} is a function $E(\cdot)$ from \mathcal{S} to projections on a Hilbert space H such that: (1) $E(\emptyset) = 0$, (2) $E(X) = id$, (3) $E(A \cap B) = E(A)E(B)$, $A, B \in \mathcal{S}$ and (4) If A_1, A_2, \dots is a disjoint sequence in \mathcal{S} then

$$E(\cup_{n=1}^{\infty} A_n) = \sum_{n=1}^{\infty} E(A_n) \quad , \text{ a strongly convergent sum}$$

Let (X, \mathcal{S}) be a measurable space and $E(\cdot)$ a projection valued measure on \mathcal{S} with values in $B(H)$. If $f = \sum_{j=1}^n a_j \chi_{B_j}$ is a simple complex valued measurable function on X , define the integral with respect to E as

$$\int f dE := \sum_{j=1}^n a_j E(B_j)$$

We now consider unbounded functions on Banach spaces. Let $T: X \rightarrow Y$ be a (possibly unbounded) function. If $\mathcal{D} \subseteq X$ is dense, and $T: \mathcal{D} \rightarrow Y$, we say that T is **densely defined**. If T is a function from X to Y with domain \mathcal{D} , the **graph** of T is $G_T := \{(x, Tx) : x \in \mathcal{D}\}$. We say that T is **closed** if $G_T \subseteq X \times Y$ is closed.

Let H be a Hilbert space. Let $T: H \rightarrow H$ be linear and densely defined with domain \mathcal{D} . Define \mathcal{D}_{T^*} as follows: $y \in \mathcal{D}_{T^*}$ if and only if the map $x \mapsto (Tx, y)$ is continuous from \mathcal{D} to \mathbb{C} . For such a y , \exists a unique $y^* \in H$ such that $(Tx, y) = (x, y^*)$ (by Riesz representation, Thm. 10.2). We define $T^*y := y^*$. Thus,

$$(Tx, y) = (x, T^*y) \quad \forall x \in \mathcal{D}_T, y \in \mathcal{D}_{T^*}$$

Write $A \subseteq B$ if $A = B$ on their common domains, and $\mathcal{D}_A \subseteq \mathcal{D}_B$. An important property is: if $A \subseteq B$, then $B^* \subseteq A^*$. (Since $B \supseteq A$, there are more x 's to check that $x \mapsto (Bx, y)$ is continuous, so less y 's will satisfy this condition, so $\mathcal{D}_{B^*} \subseteq \mathcal{D}_{A^*}$).

Let A be densely defined in H . A is **symmetric** if $A \subseteq A^*$ (i.e. $(Ax, y) = (x, Ay) \forall x, y \in \mathcal{D}_A$). A is called **self-adjoint** if $A = A^*$ (i.e. $\mathcal{D}_A = \mathcal{D}_{A^*}$, and the operators agree on this domain.) A linear operator $A: H \rightarrow K$ between Banach spaces H, K is **compact** if the image of every bounded set has compact closure.

A **semigroup of operators** on a Banach space B is a function $s \mapsto T_s$ from $[0, \infty)$ to bounded operators on B such that: $T_0 = id$ and $T_{t+s} = T_t T_s$ for $s, t \geq 0$. A semigroup is called **strongly continuous** if for each $x \in B$, the function $t \mapsto T_t x$ is continuous from $[0, \infty)$ into B . Note that norm continuity of a semigroup (i.e. continuity of $\|T_t - T_s\|$) is stronger than strong continuity, since $\|T_t x - T_s x\| \leq \|T_t - T_s\| \|x\| \rightarrow 0$ as $t \rightarrow s$. A semigroup of operators is called a **contraction semigroup** if $\|T_t\| \leq 1$ for all $t \geq 0$. If $T_t: B \rightarrow B$ is a semigroup of linear operators, define $Af := \lim_{h \downarrow 0} \frac{T_h - id}{h} f$ with domain

$\mathcal{D}_A := \{f \in B : Af \text{ exists}\}$. Then A is called the **infinitesimal generator** of the semigroup T_t . From the definition, we can see that A is a linear function on \mathcal{D}_A .

For $a \leq t \leq b$ let $u_t: [a, b] \rightarrow B$. We say u_t is **strongly continuous** at the point t if $\lim_{h \rightarrow 0} u_{t+h} = u_t$. If $\lim_{h \rightarrow 0} \frac{u_{t+h} - u_t}{h} = f$, then u_t is **strongly differentiable** at the point t and $f =: du_t/dt$. Let $a = t_0 < t_1 < \dots < t_n = b$ and $\delta := \max_{1 \leq k \leq n} |t_k - t_{k-1}|$. If $\lim_{\delta \rightarrow 0} \sum_{k=0}^n u_{t_k} (t_k - t_{k-1})$ exists, then the function u_t is said to be **strongly integrable** on the interval $[a, b]$, and the limit is denoted by $\int_a^b u_t dt$.

If u_t is strongly continuous on $[a, b]$, then by mimicking the theory of the Riemann integral, u_t is strongly integrable on $[a, b]$. Let $T: B \rightarrow C$, $u_t: [a, b] \rightarrow B$. If u_t is strongly differentiable, then Tu_t is also strongly differentiable and $\frac{d(Tu_t)}{dt} = T\left(\frac{du_t}{dt}\right)$. If u_t is strongly integrable on $[a, b]$, then Tu_t is also strongly integrable on $[a, b]$ and $\int_a^b Tu_t dt = T\left(\int_a^b u_t dt\right)$. If u_t is strongly integrable on the interval $[a, a+h]$ and strongly continuous from the right at a , then $\lim_{h \downarrow 0} \frac{1}{h} \int_a^{a+h} u_t dt = u_a$.

If du_t/dt is strongly continuous on $[a, b]$, then $\int_a^b \frac{du_t}{dt} dt = u_b - u_a$. To see this, let ξ be a linear functional. One of our properties above shows that $d(\xi u_t)/dt$ is Riemann integrable. So, using our other properties and the fundamental theorem of calculus for Riemann integrals,

$$\xi \left(\int_a^b \frac{du_t}{dt} dt \right) = \int_a^b \xi \left(\frac{du_t}{dt} \right) dt = \int_a^b \frac{d(\xi(u_t))}{dt} dt = \xi(u_b) - \xi(u_a)$$

Since the weak topology is Hausdorff, we conclude that $\int_a^b \frac{du_t}{dt} dt = u_b - u_a$, as desired. If u_t is strongly integrable on $[a, b]$ then u_{t-h} is strongly integrable on $[a+h, b+h]$ and $\int_{a+h}^{b+h} u_{t-h} dt = \int_a^b u_t dt$.

Let $A: B \rightarrow B$ be a bounded operator. Define the **exponential function** by $e^A := \lim_{N \rightarrow \infty} \sum_{n=0}^N \frac{1}{n!} A^n$. Since $\|A^n/n!\| \leq \|A\|^n/n!$, the limit in the definition of e^A exists, and $\|e^A\| \leq e^{\|A\|}$. Note that $e^{c \cdot id} = e^c \cdot id$. Also, if $AA' = A'A$, then $e^A e^{A'} = e^{A+A'}$ (using the same proof as in the case of the usual exponential function). We now show that $\| \frac{e^{tA} - id}{t} - A \| \rightarrow 0$ as $t \rightarrow 0$. Indeed, $\| e^{tA} - id - tA \| \leq \sum_{n=2}^{\infty} \frac{\|A\|^n}{n!} t^n = e^{t\|A\|} - 1 - t\|A\|$. So, dividing by t and letting $t \rightarrow 0$ gives $\| \frac{e^{tA} - id}{t} - A \| \leq \frac{e^{t\|A\|} - 1}{t} - \|A\| \rightarrow 0$ as $t \rightarrow 0$. Now, if $AA' = A'A$, and if for $t \geq 0$ $\|e^{tA}\| \leq 1$, $\|e^{tB}\| \leq 1$, then for any $f \in B$, $\|e^{tA}f - e^{tB}f\| \leq t\|Af - Bf\|$. To see this, note that

$$e^{tA}f - e^{tA'}f = (e^{\frac{tA}{n}})^n f - (e^{\frac{tA'}{n}})^n f = \left(\sum_{k=1}^n e^{\frac{k-1}{n}tA} e^{\frac{n-k}{n}tB} \right) (e^{\frac{tA}{n}}f - e^{\frac{tA'}{n}}f)$$

So, $\|e^{tA}f - e^{tA'}f\| \leq n\|e^{tA/n}f - e^{tA'/n}f\|$. But $\frac{e^{tA/n} - id}{t/n} \rightarrow A$ and $\frac{e^{tA'/n} - id}{t/n} \rightarrow A'$ as $n \rightarrow \infty$, so $n(e^{tA/n} - e^{tA'/n}) \rightarrow t(A - B)$ as $n \rightarrow \infty$. So, letting $n \rightarrow \infty$ gives $\|e^{tA}f - e^{tA'}f\| \leq t\|A - B\|$, as desired.

Theorem 10.1. (Hilbert space projections) Let H be a Hilbert space, $W \subseteq H$ a closed convex set, $u \in H$, $M \subseteq H$ a closed subspace. Define $M^\perp := \{h \in H : (h, m) = 0, \forall m \in M\}$.

- $\exists v \in W$ with $\|u - v\| = \inf_{w \in W} \|u - w\|$.
- Every $u \in H$ can be uniquely written as $u = m + v$, $m \in M$, $v \in M^\perp$. (We therefore write $H = M \oplus M^\perp$)
- $(M^\perp)^\perp = M$

Proof of (a): Let $a := \inf_{w \in W} \|u - w\|$. Let $\{w_n\} \subseteq W$ be a minimizing sequence. The parallelogram law says

$$\|2u - (w_n + w_m)\|^2 + \|w_n - w_m\|^2 = 2(\|u - w_m\|^2 + \|u - w_n\|^2) \rightarrow 4a^2 \quad (*)$$

as $m, n \rightarrow \infty$. But $\frac{1}{2}(w_n + w_m) \in M$, so $4\|u - \frac{1}{2}(w_n + w_m)\|^2 \geq 4a^2$, by definition of a . Then from the left side of (*), $\|w_n - w_m\|^2 \rightarrow 0$, so $\{w_n\}$ is Cauchy, i.e. $v := \lim w_n$ exists, and continuity of the norm gives $\|u - v\| = a$.

(b): First observe that M^\perp is closed and $M \cap M^\perp = 0$ by definition of M^\perp . Uniqueness follows since $M \cap M^\perp = 0$. To get existence, use part (a) to find $v \in M$ with $\|u - v\| = \inf_{m \in M} \|u - m\|$. Let $m \in M$ with $\|m\| = 1$. Then $v + (u - v, m)m \in M$, and by definition of v ,

$$\|u - v\|^2 \leq \|u - v - (u - v, m)m\|^2 = \|u - v\|^2 - |(u - v, m)|^2$$

Thus $(u - v, m) = 0$, i.e. $u - v \in M^\perp$, so $u = v + (u - v)$.

(c): By definition, $M \subseteq M^{\perp\perp}$. For $u \in M^{\perp\perp}$, apply part (b) to get $u = m + m^\perp$, so $0 = m^\perp + (m - u)$, $m^\perp \in M^\perp$, $m - u \in M^{\perp\perp}$. Apply part (b) again to M^\perp , so $H = M^\perp \oplus M^{\perp\perp}$. By uniqueness of this decomposition for $0 \in H$, we conclude $m^\perp = 0$, $m - u = 0$, so $u = m \in M$, i.e. $M^{\perp\perp} = M$. \square

Theorem 10.2. (Riesz Representation Theorem, Hilbert space version) Let ℓ be a continuous linear functional on a Hilbert space H . Then \exists unique $v \in H$ with $\ell(u) = (u, v)$. Also, $\|\ell\| = \|v\|$.

Proof. Uniqueness is clear. For existence, if $\ell = 0$ take $v = 0$. Otherwise let $M = \{u: \ell(u) = 0\}$. Observe that M is a closed subspace and $M \neq H$. So we can let $w \neq 0$, $w \in M^\perp$, via Thm. 10.1(b). Then $\ell(w) \neq 0$. Let $v = (\overline{\ell(w)} / \|w\|^2)w$. Then $\ell(u - (\ell(u)/\ell(w))w) = 0$, so $u - (\ell(u)/\ell(w))w \in M$, and $v \in M^\perp$ so

$$\langle u, v \rangle = \left\langle u - \left(u - \frac{\ell(u)}{\ell(w)}w\right), v \right\rangle = \left\langle \frac{\ell(u)}{\ell(w)}w, \frac{\overline{\ell(w)}}{\|w\|^2}w \right\rangle = \ell(u)$$

Finally, Cauchy-Schwarz shows $|\ell(u)| = |(u, v)| \leq \|u\| \|v\|$ and $|(v, v)| = |\ell(v)| \leq \|\ell\| \|v\|$. \square

Theorem 10.3. (Hahn-Banach, Abstract Form) Let X be a real vector space, and let $p: X \rightarrow \mathbb{R}$ be a function such that $p(x + x') \leq p(x) + p(x')$ and $p(tx) = tp(x)$ ($x, x' \in X, t \geq 0, t \in \mathbb{R}$). Let f be a linear functional on a subspace Y of X with $f(y) \leq p(y)$ for all $y \in Y$. Then \exists a linear functional F on X with $F(y) = f(y) \forall y \in Y$, and such that $F(x) \leq p(x) \forall x \in X$.

(Hahn-Banach, Functional Form) Let $Y \subseteq X$ be a vector subspace of the normed linear space X . Let $y^* \in Y^*$. Then $\exists x^* \in X^*$ with $\|x^*\| = \|y^*\|$ and $x^*(y) = y^*(y), \forall y \in Y$.

Proof of Abstract Form: Trick: use Zorn's Lemma, and then extend by one dimension. For any two linear functionals $(f_1, \text{Dom}(f_1)), (f_2, \text{Dom}(f_2))$ that are dominated by p , write $f_1 \leq f_2$ if $f_1 = f_2$ on $\text{Dom}(f_1) \cap \text{Dom}(f_2)$, and $\text{Dom}(f_1) \subseteq \text{Dom}(f_2)$. Thus, $f_1 \leq f_2$ if and only if f_2 extends f_1 . Any chain of such extensions has an upper bound F . (For $x \in \cup_\alpha \text{Dom}(f_\alpha)$, $x \in \text{Dom}(f_{\alpha_0})$ for some α_0 , so define $F(x) := f_{\alpha_0}(x)$.) By construction, F is then linear, since $x, y \in \text{Dom}(F)$ implies $x, y \in \text{Dom}(f_{\alpha_1})$ for some α_1 . By Zorn's Lemma, we can let (f_0, Y_0) be a maximal extension. It remains to prove $Y_0 = X$, by contradiction.

Assume $y_1 \in X$, $y_1 \notin Y_0$. Let $Y_1 = \text{span}\{y_1, Y_0\}$. A simple argument shows: every $x \in Y_1$ can be written uniquely as $x = y + cy_1$, $y \in Y_0, c \in \mathbb{R}$. Define f_1 on Y_1 by $f_1(y + cy_1) := f_0(y) + ck$ for some $k \in \mathbb{R}$, chosen below. It remains to show that $f_1 \leq p$ on Y_1 . Let $y, y' \in Y_0$ and observe:

$$f_0(y') - f_0(y) = f_0(y' - y) \leq p(y' - y) \leq p(y' + y_1) + p(-y_1 - y)$$

That is,

$$-p(-y_1 - y) - f_0(y) \leq p(y' + y_1) - f_0(y') \quad (*)$$

Let k be any number between: the sup of the left side of $(*)$ over $y \in Y_0$, and the inf of the right side of $(*)$ over $y' \in Y_0$. We need three cases. If $c = 0$, then $f_1(x) = f_0(x) \leq p(x)$. If $c > 0$, then for $x = y + cy_1$, the choice of k for $y' = c^{-1}y$ gives

$$f_1(x) = f_0(y) + ck \leq f_0(y) + c(p(c^{-1}y) + y_1 - f_0(c^{-1}y)) = p(y + cy_1) = p(x)$$

If $c < 0$, then for $x = y + cy_1$, the choice of k for $y' = c^{-1}y$ gives

$$f_1(x) = f_0(y) + ck \leq f_0(y) + c(-p(y_1 - c^{-1}y) - f_0(c^{-1}y)) = p(y + cy_1) = p(x)$$

So, $f_1(x) \leq p(x)$, contradicting the maximality of f_0 .

Proof of Functional Form: If the scalars are real, apply part (a) to $p(x) = \|y^*\| \|x\|$, $f = y^*$. If the scalars are complex, extend $\Re(y^*)$ to $F: X \rightarrow \mathbb{R}$ via part (a), then define $x^*(x) := F(x) - iF(ix)$. One can check that x^* is complex linear. Also, for $y \in Y$

$$(\Re y^*)(iy) + i(\Im y^*)(iy) = y^*(iy) = iy^*(y) = -(\Im y^*)(y) + i(\Re y^*)(y)$$

so $(\Re y^*)(iy) = -(\Im y^*)(y)$, implying that

$$x^*(y) = (\Re y^*)(y) - i(\Re y^*)(iy) = (\Re y^*)(y) + i(\Im y^*)(y) = y^*(y)$$

so that x^* extends y^* . If $x^*(x) = re^{i\theta}$, then $\Im(x^*(e^{-i\theta}x)) = 0 = -F(ie^{-i\theta}x)$, so

$$|x^*(x)| = x^*(e^{-i\theta}x) = F(e^{-i\theta}x) \leq \|y^*\| \|e^{-i\theta}x\| = \|y^*\| \|x\|$$

□

Theorem 10.4. (Hahn-Banach, Geometric form, a) Let Y be a closed subspace of a normed linear space X . Let $x_0 \in X$, $x_0 \notin Y$. Then $\exists x^* \in X^*$ such that $x^*(Y) = 0$ and $x^*(x_0) = 1$, with $\|x^*\| = d(x_0, Y)^{-1}$.

(Hahn-Banach, Geometric form, b) Let X be a real normed linear space, $C \subseteq X$ a convex open set containing 0. If $x_0 \in X$, $x_0 \notin C$, then $\exists \phi: X \rightarrow \mathbb{R}$ continuous with $\phi(x_0) = 1$ and $\phi(C) < 1$

(Hahn-Banach, Geometric form, c) Let X be a locally convex linear space (that is, 0 has a neighborhood basis of convex sets), $A, B \subseteq X$ disjoint, convex, with A open. Then \exists a continuous linear functional f on X and $\alpha \in \mathbb{R}$ such that $\Re(f(A)) < \alpha$ and $\Re(f(B)) \geq \alpha$.

(Hahn-Banach, Geometric form, d) Let X be a locally convex space, $A, B \subseteq X$ disjoint closed convex sets, A compact. Then \exists a continuous linear functional f on X and $\alpha \in \mathbb{R}$ such that $\Re(f(A)) < \alpha$ and $\Re(f(B)) > \alpha$.

Proof of (a): Let $d = d(x_0, Y)$, $Z = \text{span}\{x_0, Y\}$. As above, every $x \in Z$ can be written uniquely as $x = y + cx_0$. Define $z^*(x) := c$. Then z^* is linear and for $c \neq 0$, and $\|x\| = |c| \|c^{-1}y + x_0\| \geq |c| d = d |z^*(x)|$. Thus, $\|z^*\| \leq d^{-1}$. Letting $\|x_0 - y_n\| \rightarrow d$, $\{y_n\} \in Y$ shows $1 = z^*(x_0 - y_n) \leq \|z^*\| \|x_0 - y_n\| \rightarrow d \|z^*\|$, so $\|z^*\| \geq d^{-1}$. Finally, apply Thm. 10.3(b) to extend z^* .

Proof of (b): (Sketch) On the one-dimensional space $\text{span}\{x_0\}$, define $\phi(tx_0) = t$. Observe that ϕ is dominated by the norm induced by C ($\|y\|_C := \inf\{r > 0: y \in rC\}$). Apply Thm. 10.3(a).

Proof of (c): (Sketch) Assume scalars are real. Let $a_0 \in A, b_0 \in B, x_0 := b_0 - a_0, C := A - B + x_0$. Then C is convex, open, and it contains zero. Also, $x_0 \notin C$. Apply part (b) to C and x_0 . Observe $\phi(b_0) = \phi(a_0) + 1$ and $\phi(a) < \phi(b) + \phi(a_0) - \phi(b_0) + 1, a \in A, b \in B$, so $\phi(a) < \phi(b)$. Let $\alpha = \inf_{b \in B} \phi(b)$, so $\phi(a) \leq \alpha \leq \phi(b)$. Since A is open, $\phi(a) = \alpha$ cannot occur, so $\phi(a) < \alpha$.

Proof of (d): We claim that $C := B - A$ is closed. Given the claim, note that $0 \notin C$, so C^c contains an open neighborhood U of 0. So apply part (c) to U and $B - A$ to get some real nonconstant continuous linear functional with $f(B - A) \geq c$ and $f(U) \leq c$. Let $x \in X$ such that $f(x) = 1$. For $\alpha > 0$ small, $\alpha x \in U$. Since $f(\alpha x) = \alpha, f(U) \supseteq [0, \varepsilon]$ for some small ε . Therefore, $f(B) - f(A) = f(B - A) \geq \varepsilon$, as desired.

We now prove that $C = B - A$ is closed. Let $d \in \overline{B - A}, U$ a neighborhood of $d, A_U := \{a \in A: a \in B - U\}$. Note that $A_U \neq \emptyset$ since $d \in \overline{B - A}$. (Let $d_n \rightarrow d, d_n \in B - A$. Without loss of generality, $d_n \in U$. Then $d_n = b_n - a_n$, so $a_n = b_n - d_n, d_n \in U$. Since A is compact, after taking a subsequence we may assume that $a_n \rightarrow a, a \in A$, so $b_n \rightarrow b$ as well. Since B is closed, $b \in B$. Now $d = \lim(b_n - a_n) = b - a, b \in B, a \in A$, so $a = b - d \in A_U$, as desired.)

Observe $U \subseteq V$ implies $A_U \subseteq A_V$. So, any finite subset of the sets $\{A_U\}$ has the nonempty intersection. So, the compact sets $\{\overline{A_U}\}$ have a common element a_0 . (If not, then $\cap_U A_U$ is empty. Write $A_U =: A \setminus B_U$. Since $\cap_U A_U = \cap_U (A \setminus B_U) = A \setminus (\cup_U B_U), \cup_U B_U = A$. But B_U is open and A is compact, so $A = \cup_{i=1}^n B_{U_i}$, implying $\cap_{i=1}^n A_{U_i} = \emptyset$, a contradiction.)

Let N be a neighborhood of 0. Then $U = N + d$ is a neighborhood of d , so $a_0 \in B - U = B - N - d$, so $(N + a_0) \cap (B - N - d) \neq \emptyset$. Shifting an N from the right to the left, we get $(N + N + a_0) \cap (B - d) \neq \emptyset$. Let M be a neighborhood of 0. By continuity of addition, there exists an open neighborhood N of 0 such that $N + N \subseteq M$. So, any neighborhood of a_0 intersects $B - d$. Since B is closed, $B - d$ is closed. Therefore, by the limit point definition of closedness, $a_0 \in B - d$, so $d \in B - a_0 \subseteq B - A$, as desired. \square

Theorem 10.5. (Hahn-Banach Corollaries) Let X be a normed linear space. If $x_0 \in X, x_0 \neq 0$, then $\exists x^* \in X^*$ with $\|x^*\| = 1$ and $x^*(x_0) = \|x_0\|$. Thus, setting $x_0 = x - x'$, we see that the weak topology is Hausdorff. Also, $\|x_0\| = \sup_{\|x^*\| \leq 1} |x^*(x_0)|$.

Proof. First, use Thm. 10.4(a) with $Y = 0$. The formula for $\|x_0\|$ follows, since $\sup_{\|x^*\| \leq 1} |x^*(x_0)| \leq \sup_{\|x^*\| \leq 1} \|x^*\| \|x_0\| = \|x_0\|$. \square

Theorem 10.6. (Uniform Boundedness Principle/ Banach-Steinhaus) Let $\{L_\alpha\}$ be a set of bounded linear operators from a Banach space X into a normed linear space Y . Assume the L_α are pointwise bounded, i.e. $\|L_\alpha(x)\| \leq C_x$ for all α . Then $\exists C$ with $\|L_\alpha\| \leq C$ for all α (C independent of x).

Proof. Let $F_n := \{x \in X: \|L_\alpha(x)\| \leq n, \forall \alpha\}$. Since $F_n = \cap_\alpha (\|\cdot\| \circ L_\alpha)^{-1}[0, n]$, F_n is closed. By assumption, $\cup_n F_n = X$. Since X is a complete metric space, Baire's Category Theorem shows that some F_N contains a nonempty open ball $B = B(b, 2r) \subseteq X$. Then $\|L_\alpha(x)\| \leq N$ for all α and for all $x \in B$. Dilating and translating B gives our result. Specifically, for $\|x\| \leq 1$ we have $rx + b \in B$, so

$$\|L_\alpha(x)\| = r^{-1} \|L_\alpha(rx)\| \leq r^{-1} (\|L_\alpha(rx + b)\| + \|L_\alpha(b)\|) \leq r^{-1}(N + C_b)$$

i.e. $\|L_\alpha\| \leq r^{-1}(N + C_b)$. □

Theorem 10.7. (Interior Mapping (Open Mapping) Principle)

- (a) Let $L: X \rightarrow Y$ be a continuous linear operator between Banach spaces, which is onto. Then L is open (i.e. its inverse is continuous, if it exists).
- (b) Suppose $\overline{L(B_X)}$ contains an open ball. ($B_X := \{x: \|x\| \leq 1\} = B(0, 1)$). Then $\exists r > 0$ with $L(B_X) \supseteq r \cdot B_Y$, and (by scaling) $L(X) = Y$.

Proof of (a): Trick: Baire, show L is open at zero by approximations, then translate. Let $B(0, n) = \{x: \|x\| < n\}$. Write $Y = L(X) = L(\cup_{n \geq 1} \overline{B(0, n)}) = \cup_{n \geq 1} \overline{L(B(0, n))}$. Thus $Y = \cup_{n \geq 1} \overline{L(B(0, n))}$. Baire's Category Theorem says some $\overline{L(B(0, n))}$ contains an open ball. By scaling (and using continuity) $\overline{L(B(0, n))} = (2n)\overline{L(B(0, 1/2))}$, so $\overline{L(B(0, 1/2))}$ also contains an open ball V . A continuity argument shows that $V - V \subseteq \overline{L(B(0, 1))}$. (Let $v, v' \in \overline{L(B(0, 1/2))}$, let $v_n, v'_n \in L(B(0, 1/2))$ with $v_n \rightarrow v, v'_n \rightarrow v'$. Then $v_n - v'_n \in L(B(0, 1))$, $v_n - v'_n \rightarrow v - v' \in \overline{L(B(0, 1))}$.) Since $V - V$ is open, it is also a neighborhood containing zero. Thus, $\overline{L(B(0, 1))} \supseteq B(0, s)$ for some $s > 0$. Using continuity again,

$$\overline{L(B(0, t))} \supseteq B(0, st) \quad \forall t > 0 \quad (*)$$

We now show that $\overline{L(B(0, c))} \supseteq B(0, sc/2)$ for all $c > 0$. Let $y \in B(0, sc/2)$. Using (*) with $t = c/2$ gives $x_1 \in \overline{L(B(0, c/2))}$ with $\|y - L(x_1)\| < 2^{-2}sc$. Suppose x_1, \dots, x_{n-1} satisfy $\|y - L(x_1 + \dots + x_{n-1})\| < 2^{-n}sc$. Then $y - L(x_1 + \dots + x_{n-1}) \in B(0, 2^{-n}sc)$, so (*) with $t = 2^{-n}c$ gives $x_n \in \overline{L(B(0, 2^{-n}c))}$ with $\|y - L(x_1 + \dots + x_n)\| < 2^{-(n+1)}sc$. Since the norms of the x_n decrease exponentially, $x = \sum_{n=1}^{\infty} x_n$ exists and $\|x\| \leq c$. By definition of the x_i and since L is continuous we get $y = L(x)$, $x \in \overline{L(B(0, c))}$, as desired.

So, $\overline{L(B(0, 2c))} \supseteq B(0, sc/2)$. By translating this result, we see: for U open, each $w \in L(U)$ has a neighborhood contained in $L(U)$, as desired.

Proof of (b): Argue as in part (a). □

Theorem 10.8. (Closed graph theorem) Let $L: X \rightarrow Y$ be a linear function between Banach spaces. Suppose the graph $G := \{(x, L(x)): x \in X\}$ of L is closed in $X \times Y$. (If $x_n \rightarrow x$ and $f(x_n) \rightarrow y$, then $x \in \text{Dom}(L)$ and $f(x) = y$). Then L is bounded.

Proof. Trick: Use a projection. Define the Banach space $X \oplus Y$ via the norm $\|(x, y)\| := \|x\|_X + \|y\|_Y$. By assumption, $G \subseteq X \oplus Y$ is closed. In particular, G is a Banach space. Define $P: G \rightarrow X$ by $P((x, L(x))) = x$, and observe that P is linear, continuous and bijective. Therefore, $P^{-1}(x) = (x, L(x))$ is continuous, from the Interior Mapping Principle (Thm. 10.7(a)). Finally, note that the projection $\Pi: X \oplus Y \rightarrow Y$ defined by $\Pi(x, y) := y$ is bounded, so $\Pi|_G$ is bounded, so $(\Pi|_G) \circ P^{-1} = L$ is bounded. □

Theorem 10.9. (Choquet) Let X be a metrizable, compact, convex subset of a locally convex space E . Let $x_0 \in X$. Then \exists a probability measure $\mu = \mu_{x_0}$ on X that is supported on the extreme points $\mathcal{E}(X)$ of X , and such that $x_0 = \int_{\mathcal{E}(X)} x d\mu(x)$ weakly. (That is, for any continuous linear functional F on X , $F(x_0) = \int_{\mathcal{E}(X)} F(x) d\mu(x)$).

Proof. Strategy: consider $X^* \subseteq C(X)$, construct a strictly convex function, apply Hahn-Banach. By Thm. 10.10, $C(X)$ is separable. Let $A \subseteq C(X)$ be the set of continuous affine functions. Let $\{h_n\}_{n=1}^{\infty} \subseteq A$ be dense on the unit sphere of A . Observe that $f = \sum_{n \geq 1} 2^{-n} h_n^2$

converges uniformly and is strictly convex and nonnegative. Let $B = \text{span}\{A, f\}$. For $g \in C(X)$, define $\bar{g}(x) := \inf\{h(x) : h \in A, h \geq g\}$ (this is the upper concave envelope of the graph of g). Since $\bar{g} + \bar{u} \leq \overline{g+u}$, $p(g) := \bar{g}(x_0)$ is a subadditive function on $C(X)$, which is also homogeneous. Define $\ell: B \rightarrow \mathbb{R}$ by $\ell(h+rf) = h(x_0) + rf(x_0)$. On B , note that $\ell \leq p$. (For $r > 0$, $h+rf = \overline{h+rf}$, and for $r < 0$, $h+rf$ is concave, so $\overline{h+rf} = h+rf \geq \overline{h+rf}$).

Since $\ell \leq p$ on B , Hahn-Banach (Thm. 10.3(a)) gives $m \in C(X)^*$ which extends ℓ and remains dominated by p . Since $m \leq p$, m is nonpositive on nonpositive functions, and Riesz's Representation Theorem shows $m(f) = \int f d\mu$ for μ a regular Borel (probability) measure (using $1 \in A$, so $m(1) = 1$). Now, $f \leq \bar{f}$ so $\mu(f) \leq \mu(\bar{f})$. Conversely, if $h \in A$ and $h \geq f$, then $h \geq \bar{f}$, so $h(x_0) = m(h) = \mu(h) \geq \mu(\bar{f})$. So, taking the infimum over $h \in A, h \geq f$, we get $\bar{f}(x_0) \geq \mu(\bar{f})$. But $\mu(f) = m(f) = \bar{f}(x_0)$, so $\mu(f) \geq \mu(\bar{f})$. Thus, $\mu(f) = \mu(\bar{f})$, i.e. μ vanishes outside $E := \{x : f(x) = \bar{f}(x)\}$. Finally, E is supported on the extreme points of X , since if $x = (1/2)(y+z)$, then strict convexity of f implies $f(x) < (1/2)(f(y) + f(z)) \leq (1/2)(\bar{f}(y) + \bar{f}(z)) \leq \bar{f}(z)$. (And actually, E is equal to the extreme points). \square

Theorem 10.10. (Separability Conditions) *Let X be a locally compact separable metric space, with μ a Borel measure on X . For $1 \leq p < \infty$ we have*

- (a) $(C_c(X), \|\cdot\|_{sup})$ is a separable normed linear space
- (b) $L_p(X, \mu)$ is separable

Theorem 10.11. (Characterization of Locally Convex Spaces) *Let X be a topological linear space. Let \mathcal{N} be the family of continuous semi-norms on X . Let $X^{\mathcal{N}}$ be the Fréchet space formed by X and \mathcal{N} . Then $X^{\mathcal{N}} = X$ iff X is locally convex*

Theorem 10.12. (Alaoglu Theorem/ Banach-Alaoglu) *Let X be a normed linear space. Then the unit ball $B_{X^*} = \{x^* : \|x^*\| \leq 1\}$ of X^* is weak* compact.*

Proof. Let A be the set of scalar valued functions ξ on X with $\|\xi(x)\| \leq \|x\|$ for all $x \in X$. Equivalently, $A = \prod_{x \in X} \{\lambda \in \{\text{scalars}\} : |\lambda| \leq \|x\|\}$. Then A with the product topology is compact by Tychonoff's Theorem. By the definition of the product topology, a basic open neighborhood of some ξ_0 is $\{\xi : |\xi(x_j) - \xi_0(x_j)| < \varepsilon, j = 1, \dots, n\}$. Now for fixed $x \in X$, the projection map $\xi \mapsto \xi(x)$ is continuous, from A (with the product topology) to scalars. (Given ξ in the inverse image of a small open interval, ξ is contained in an open set in A). Consider the natural embedding $B_{X^*} \subseteq A$. By the definition of the weak* topology, the topology induced by $B_{X^*} \subseteq A$ is exactly the weak* topology.

Putting everything together, let $x, y \in X$, α, β scalars, and observe: $\xi(\alpha x + \beta y) - \alpha\xi(x) - \beta\xi(y)$ is a continuous function of ξ , from A to scalars (since it is a composition of continuous functions). Therefore, $\{\xi : \xi(\alpha x + \beta y) - \alpha\xi(x) - \beta\xi(y) = 0\}$ is closed in A (being an inverse image of zero). Therefore,

$$B_{X^*} = \cap_{x,y,\alpha,\beta} \{\xi : \xi(\alpha x + \beta y) - \alpha\xi(x) - \beta\xi(y) = 0\}$$

is closed in the compact set A . \square

Remark 10.13. If X is separable then $(B_{X^*}, \sigma(X^*, X))$ is metrizable. Using the homeomorphism $\iota: (X, \sigma(X, X^*)) \rightarrow (X^{**}, \sigma(X^{**}, X^*))$, we see: if X^* is separable then $(B_X, \sigma(X, X^*))$ is metrizable. To prove the first assertion, let $\{x_n\}_{n \geq 1}$ be dense in B_X . Then, define $\rho(x_1^*, x_2^*) = \sum_{n \geq 1} 2^{-n} |x_1^*(x_n) - x_2^*(x_n)|$. Essentially by definition (using open balls), ρ is

$\sigma(X^*, X)$ continuous. That is, a point x in an open ρ ball U has $x \in V \subseteq U$, V a basic open set in the weak* topology. And the reverse inclusion of open sets holds by scaling appropriately. So, for X separable, B_{X^*} is sequentially compact.

Theorem 10.14. (Properties of Faces)

- (a) $\cap\{\text{faces}\} = \text{face}$
- (b) *A face of a face is a face.*
- (c) *For $A: H \rightarrow \tilde{H}$ continuous and linear, if K is compact and convex, then $A(K)$ is too.*
- (d) *If \tilde{F} is a face of $\tilde{K} := A(K)$, then $F := K \cap A^{-1}(\tilde{F})$ is a face of K .*
- (e) *If $\tilde{F} \subsetneq \tilde{K} := A(K)$, then $F := K \cap A^{-1}(\tilde{F}) \neq K$.*
- (f) *Let X be a topological vector space such that X^* separates points of X (e.g. X is normed linear, using Thm. 10.5, or X is locally convex, using Thm. 10.11). Then a nonempty extreme face is an extreme point.*

Proof of (f): (Properties (a),(b),(c) and (e) are routine. For (d), apply A to a segment containing a point of F .) View X as a vector space over \mathbb{R} . Let $\xi \in X^*$ be real valued. Let $F \subseteq K$ be an extreme face. By definition of a face, F is compact, convex. Since $\xi(F)$ is compact, convex from (c), $\xi(F) = [a, b]$. Now, $\{a\}$ is a face of $\xi(F)$, so $F \cap \xi^{-1}(a)$ is a nonempty face of F . Since F is extreme, $F = F \cap \xi^{-1}(a)$, i.e. $\xi(F) = \{a\}$. If F had two distinct points, we could choose ξ which separates them, by assumption on X^* . Therefore, F has only one point. \square

Theorem 10.15. (Krein-Milman) *Let X be a locally convex topological vector space. If K is compact and convex, then $\hat{K} = K$.*

Proof. We will use Thm. 10.14 multiple times, without further note. Recall that \hat{K} is the intersection of all compact, convex sets containing the extreme points of K . So, by definition, $\hat{K} \subseteq K$. We will see below that $\hat{K} \neq \emptyset$. For ξ a real continuous linear functional, $\xi(\hat{K}) \subseteq \xi(K) = [a, b]$. Now, $K \cap \xi^{-1}(a)$ is a face of K . Let T be the set of non-empty faces of $K \cap \xi^{-1}(a)$. For $t_1, t_2 \in T$, write $t_1 \leq t_2$ if $t_2 \subseteq t_1$ (i.e. we reverse the inclusion). Every chain has a nonempty upper bound, since the intersection of nested compact sets is nonempty. (Recall: we take the definition of compactness, take the contrapositive, and then take complements and apply de Morgan's law. Finally, use the finite intersection property.) By Zorn's Lemma, \exists a nonempty extreme face in $K \cap \xi^{-1}(a)$, which is therefore an extreme point. Thus, $(K \cap \xi^{-1}(a)) \cap \hat{K} \neq \emptyset$, so $a \in \xi(\hat{K})$, so the endpoints of $\xi(K)$ are in $\xi(\hat{K})$, so (by convexity of \hat{K}), $\xi(K) \subseteq \xi(\hat{K})$. Thus, $\xi(K) = \xi(\hat{K})$ for all real linear functionals ξ , so $\hat{K} = K$ from Hahn-Banach in Geometric Form (Thm. 10.4(d)). (If $\hat{K} \subsetneq K$, let $x \in K \setminus \hat{K}$, so Hahn-Banach gives ξ with $\xi(\hat{K}) < \alpha, \xi(x) > \alpha$.) \square

Remark 10.16. Given X, Y , we consider whether or not $Y = X^*$. If so, then $B_Y = B_{X^*}$ is compact and convex, so this theorem says $\widehat{B_Y} = B_Y$. Thus, analyzing the extreme points of B_Y (existence, nonexistence, finiteness, etc.) can tell us whether or not $Y = X^*$ for any X . For example, the unit ball of $C([0, 1], \mathbb{R})$ only has extreme points ± 1 , and that of $L_1([0, 1], \mathbb{R})$ has no extreme points.

Theorem 10.17. (Weak Topology Equivalence)

- (a) The $\sigma(X, X^*)$ topology on X is the weakest topology such that all norm-continuous linear functionals are still continuous. Similarly, the $\sigma(X^*, X)$ topology on X^* is the weakest topology which makes all functionals in $\iota(X) \subseteq X^{**}$ continuous.
- (b) The space of all continuous linear functionals on $(X, \sigma(X, X^*))$ equals X^* .
- (c) Let ϕ_0, \dots, ϕ_n be linear forms on a linear space X (i.e. ignore any topology). Then the following are equivalent: (i) $\phi_0 \in \text{span}\{\phi_j\}_{j=1}^n$, and (ii) $\ker \phi_0 \supseteq \bigcap_{j=1}^n \ker \phi_j$
- (d) The space of all continuous linear functionals on $(X^*, \sigma(X^*, X))$ equals X .

Proof of (a): Apply definitions.

(b): The forward direction follows since the weak topology is weaker than the strong topology. The reverse is part (a).

(c): (i) easily implies (ii). For the reverse, define $\pi: X \rightarrow \mathbb{K}^n$ by $\pi(x) = \{\phi_j(x)\}_{j=1}^n$ ($\mathbb{K} = \text{scalars}$). Then $\ker \pi = \bigcap_{j=1}^n \ker \phi_j$, so (ii) implies that ϕ_0 induces a linear form $\Lambda: \mathbb{K}^n \rightarrow \mathbb{K}$, i.e. $\phi_0 = \Lambda\pi$. (For $z = \{\phi_j(x)\}_{j=1}^n$, z is in the linear space $\pi(X)$, so define $\Lambda(z) = \phi_0(x)$. By (ii), Λ is well-defined. Λ is linear on $\pi(X)$, so extend Λ linearly to \mathbb{K}^n .) From linear algebra, $\phi_0(x) = \sum_{j=1}^n \alpha_j \phi_j(x)$.

(d): Let ϕ be a linear functional on X^* continuous in $\sigma(X^*, X)$. Then $\{x^* \in X^*: |\phi(x^*)| < 1\} \supseteq \{x^* \in X^*: |x_j(x^*)| < \varepsilon, j = 1, \dots, n\}$ for some $\varepsilon > 0$, $x_1, \dots, x_j \in X^*$. In particular, viewing $x_j \in X^{**}$, $\ker \phi \supseteq \bigcap_{i=1}^j \ker(x_j)$. Finally, apply part (c). \square

Theorem 10.18. (Goldstine) *The closed unit ball of X is $\sigma(X^{**}, X^*)$ -dense in the closed unit ball of X^{**} . (Here we identify B_X with $\iota(B_X) \subseteq X^{**}$)*

Proof. Trick: use X^{***} , Hahn-Banach, contradiction. Let $V = \overline{B_X}^{\sigma(X^{**}, X^*)}$. Assume for the sake of contradiction $\exists x^{**} \in X^{**}$ with $\|x^{**}\| \leq 1$, $x^{**} \notin V$. Alaoglu (Thm. 10.12) says V is $\sigma(X^{**}, X^*)$ -compact. So, Hahn-Banach, geometric form (Thm. 10.4(d)) gives ϕ a linear functional on $(X^{**}, \sigma(X^{**}, X^*))$ with $\phi(x^{**}) > \sup\{\phi(v): v \in V\}$. Theorem 10.17(d) says $\phi(x^{**}) = x^{**}(x_0^*)$, for some $x_0^* \in X^*$. But then

$$\|x_0^*\| = \sup_{\substack{x \in X \\ \|x\| \leq 1}} |x_0^*(x)| \leq \sup_{v \in V} |v(x_0^*)| = \sup_{v \in V} |\phi(v)| < \phi(x^{**}) = x^{**}(x_0^*) \leq \|x_0^*\|$$

which is a contradiction. \square

Theorem 10.19. (Reflexive equivalences) *Let X be a Banach space. The following are equivalent:*

- (a) X is reflexive ($\iota(X) = X^{**}$)
- (b) X^* is reflexive
- (c) B_X is $\sigma(X, X^*)$ compact
- (d) Every subspace of X is reflexive
- (e) Every quotient space of X is reflexive

Proof: (a) implies (c): Using ι , we know that $(B_X, \sigma(X, X^*))$ is homeomorphic to $(B_{X^{**}}, \sigma(X^{**}, X^*))$ so apply Alaoglu (Thm. 10.12).

(c) implies (a): $\iota(B_X)$ is compact, so apply Goldstine (Thm. 10.18).

(d) implies (a): True by definition.

(a) implies (d): Let $Y \subseteq X$ be a norm-closed subspace. By, say, Hahn-Banach (Thm. 10.5), Y is $\sigma(X, X^*)$ closed. By (c), $B_Y \subseteq B_X$ is $\sigma(X, X^*)$ compact. By restriction of X^* , B_Y is $\sigma(Y, Y^*)$ compact, so Y is reflexive, since (c) implies (a).

(a) implies (b): By reflexivity, $(X^*, \sigma(X^*, X)) = (X^*, \sigma(X^*, X^{**}))$, so Alaoglu (Thm. 10.12) shows B_{X^*} is $\sigma(X^*, X^{**})$ compact, so use: (c) implies (a).

(b) implies (a): X^* reflexive implies X^{**} is reflexive (since (a) implies (b)) so X is reflexive, using $\iota(X) \subseteq X^{**}$ and that (a) implies (d).

(e) \iff (a): $(X/Y)^* \subseteq X^*$ from Hahn-Banach, so use: (a) \iff (b) \iff (d). \square

Remark 10.20. Let (S, μ) be a σ -finite measure space with infinitely many disjoint sets of positive measure. Then L_1, L_∞ are not reflexive. Since $(L_1)^* = L_\infty$, the previous theorem says, it suffices to show that L_1 is not reflexive. So it suffices to find $\phi \in (L_\infty)^*$ with $\phi \notin \iota(L_1)$. Let $\{E_n\}_{n=1}^\infty$ be disjoint sets of positive measure. Let $Y \subseteq L_\infty$ be the set of functions that are constant on each E_n , with $y \in Y$ satisfying: $\lim_{n \rightarrow \infty} y(E_n)$ exists. Let $y^*(y) = \lim_{n \rightarrow \infty} y(E_n)$. Then y^* has norm 1, and Hahn-Banach (Thm. 10.3) gives an extension $x^* \in (L_\infty)^*$. For the sake of contradiction, assume $\exists g \in L_1$ such that $x^*(f) = \int_S f g d\mu$. In particular, $x^*(f) = \int_S f g d\mu$ for all $f \in Y$. Then $x^*(1_{E_n}) = y^*(1_{E_n}) = 0$, but

$$1 = y^*(1_{\cup_{n=1}^\infty E_n}) = x^*(1_{\cup_{n=1}^\infty E_n}) = \int_S g(1_{\cup_{n=1}^\infty E_n}) = \sum_{n=1}^\infty \int_{E_n} g = \sum_{n=1}^\infty x^*(1_{E_n}) = 0$$

a contradiction. So, no such g exists, i.e. $x^* \notin \iota(L_1)$.

Theorem 10.21. (Principle of Local Reflexivity) Let X be a Banach space and let $E \subseteq X^{**}$ and $F \subseteq X^*$ be finite dimensional subspaces. Given $\varepsilon > 0$ there exists an operator $T: E \rightarrow X$ such that

- (1) $\|T\| \|T^{-1}|_{T(E)}\| \leq 1 + \varepsilon$
- (2) $T|_{E \cap X} = id$
- (3) $f(Te) = e(f)$ for all $f \in F$ and $e \in E$

Proof. We begin with the following Lemma:

Let $\{A_j\}_{j=1}^N$ be bounded, norm-open convex subsets of X and let \tilde{A}_j be the norm interior of the $\sigma(X^{**}, X^*)$ -closure of A_j in X^{**} .

- (a) If $\cap_{j=1}^N \tilde{A}_j \neq \emptyset$ then $\cap_{j=1}^N A_j \neq \emptyset$
- (b) If we have a map $T: X \rightarrow Y$ with Y a finite dimensional Banach space then $T^{**}(\cap_{j=1}^N \tilde{A}_j) = T(\cap_{j=1}^N A_j)$.

Proof of Lemma, (a): We prove the contrapositive statement. Let $X_N = \oplus_{i=1}^N (X)$ and let $A = \{\{x_j\}_{j=1}^N \in X_N : x_j \in A_j, j = 1, \dots, N\}$. Then $A \subseteq X_N$ is bounded, norm-open and convex. If $\cap_{j=1}^N A_j = \emptyset$ then $A \cap V = \emptyset$ for $V = \{\{x_j\}_{j=1}^N \in X_N : x_j = x_1 \text{ for } j = 1, \dots, N\}$. Let $\tilde{A} = \{\{x_j^{**}\}_{j=1}^N \in X_N^{**} : x_j^{**} \in \tilde{A}_j, j = 1, \dots, N\}$ and let $V^{**} = \{\{x_j^{**}\}_{j=1}^N \in X_N^{**} : x_j^{**} = x_1^{**} \text{ for } j = 1, \dots, N\}$.

If $A \cap V = \emptyset$ then since V is a closed subspace, Hahn-Banach, Thm. 10.4(a) says there exists $\phi = (\phi_1, \dots, \phi_N) \in X_N^*$ with $\phi|_V = 0$ and $\phi(a) > 0$ for all $a \in A$. By Goldstine, Thm. 10.18, A is $\sigma(X^{**}, X^*)$ dense in \tilde{A} , so $\phi(a^{**}) \geq 0$ for all $a^{**} \in \tilde{A}$. But \tilde{A} is open so $\phi(\tilde{A})$ is open, so $\phi(a^{**}) > 0$ for all $a^{**} \in \tilde{A}$. But B_V is $\sigma(X^{**}, X^*)$ dense in $B_{V^{**}}$, so $\phi(V^{**}) = 0$, so $\tilde{A} \cap V^{**} = \emptyset$, i.e. $\cap_{j=1}^N \tilde{A}_j = \emptyset$.

Proof of Lemma, (b): Note that $T^{**}(\cap_{j=1}^N \tilde{A}_j)$ and $T(\cap_{j=1}^N A_j)$ are open, convex sets and since $\iota(A_j) \subseteq \tilde{A}_j$, we have $T(\cap_{j=1}^N A_j) \subseteq T^{**}(\cap_{j=1}^N \tilde{A}_j)$. Suppose this containment is strict. Using Hahn-Banach, Thm. 10.4(d), there exists $p \in T^{**}(\cap_{j=1}^N \tilde{A}_j)$ and there exists a functional ϕ on Y^{**} and $\alpha > \beta$ such that $\phi(p) > \alpha > \beta > \phi(T(\cap_{j=1}^N A_j))$.

Let $z^{***} \in X^{***}$ with $z^{***} = T^{***}(\phi)$. Let $A_j^+ = A_j \cap \{x \in X : z^{***}(x) > \alpha\}$, $\tilde{A}_j^+ = \tilde{A}_j \cap \{x^{**} \in X^{**} : z^{***}(x^{**}) > \alpha\}$. By definition of ϕ, z^{***}, T^* , we have $\cap_{j=1}^N A_j^+ = \emptyset$. However, $\cap_{j=1}^N \tilde{A}_j^+ \neq \emptyset$, contradicting (a) and proving (b). To verify the nonempty intersection, write $p = T^{**}(z)$, so that $z^{***}(z) = T^{***}(\phi)(z) = \phi(T^{**}(z)) = \phi(p) > \alpha$, i.e. $z \in \cap_{j=1}^N \tilde{A}_j^+$.

Proof of Theorem: Let $\dim(E) = n$ and $\dim(E \cap X) = n - k$. Let $\{x_j^{**}, e_j^*\}_{j=1}^n$ be a biorthogonal system in $E \times E^*$ such that $\text{span}\{x_j^{**}\}_{j=k+1}^n = E \cap X$ and $\|x_j^{**}\| = 1$. From biorthogonality, the identity $id: E \rightarrow X^{**}$ can be written as $id(e) = \sum_{j=1}^n e_j^*(e)x_j^{**}$. We will find $x_1, \dots, x_k \in X$ such that $T: E \rightarrow X$ defined by $T(e) = \sum_{j=1}^k e_j^*(e)x_j + \sum_{j=k+1}^n e_j^*(e)x_j^{**}$ satisfies the conclusion of the theorem. Note that property (2) is already satisfied by this T . Let $Z = \oplus_{i=1}^k (X)$ and let $\delta > 0$. Fix the following three sets: $\{f_j\}_{j=1}^M$ a basis of F , $\{x_j^*\}_{j=1}^R \subseteq B_{X^*}$ a set such that for every $e \in E$, $\|e\| \leq (1 + \delta) \sup\{|x_j^*(e)| : j = 1, \dots, R\}$, and $\{e_j\}_{j=1}^N$ a δ -net in B_E .

Write $e_j = \sum_{r=1}^n \lambda_r^j x_r^{**}$. For $j = 1, \dots, N$, define $C_j \subseteq Z$ by

$$C_j := \left\{ \{x_s\}_{s=1}^k : \left\| \sum_{s=1}^k \lambda_s^j x_s + \sum_{s=k+1}^n \lambda_s^j x_s^{**} \right\| < (1 + \delta) \|e_j\| \right. \\ \left. \text{and } \|x_s\| < 1 + \delta, s = 1, \dots, k \right\}$$

Note that the C_j are norm-open, bounded and convex. Let \tilde{C}_j denote the norm interior of the $\sigma(X^{**}, X^*)$ -closure of C_j in X^{**} . Since $\{x_s^{**}\}_{s=1}^k \in \cap_{j=1}^N \tilde{C}_j \subseteq Z^{**}$, part (a) of the Lemma gives $\{z_s\}_{s=1}^k \in \cap_{j=1}^N C_j \neq \emptyset$.

Let $\mathbb{K} = \{\text{scalars}\}$ and define $S: Z \rightarrow \mathbb{K}^{M \cdot k} \oplus \mathbb{K}^{R \cdot k}$ by

$$S(\{x_s\}_{s=1}^k) := \{f_j(x_s), x_k^*(x_s)\}_{j=1, \dots, M, k=1, \dots, R, s=1, \dots, k}$$

From part (b) of the Lemma, there exists $\{z_s\}_{s=1}^k \in \cap_{j=1}^N C_j$ with $S(\{z_s\}_{s=1}^k) = S^{**}(\{z_s^{**}\}_{s=1}^k)$. By the last equality, from our formula for id , by choice of $\{f_j\}$ and by definition of T , property (3) follows. It remains to show that property (1) holds. Using again the equality $S(\{z_s\}_{s=1}^k) = S^{**}(\{z_s^{**}\}_{s=1}^k)$, we see that $x_j^*(Te) = x_j^*(e)$. So, by the defining properties of $\{x_j^*\}$,

$$\|Te\| \geq \sup_{j=1, \dots, R} |x_j^*(Te)| = \sup_{j=1, \dots, R} |x_j^*(e)| \geq (1 + \delta)^{-1} \|e\| \quad (*)$$

Given $e \in B_E$, fix e_j with $\|e - e_j\| \leq \delta$. Recall the formula $e_j = \sum_{r=1}^n \lambda_r^j x_r^{**}$. Since $\lambda_r^j = e_r^*(e_j)$ and since $\{z_s\}_{s=1}^k \subseteq C_j$, the definitions of T and C_j show that $\|Te_j\| \leq (1 + \delta) \|e_j\|$. Therefore $\|Te\| \leq \|Te_j\| + \|T(e - e_j)\| \leq (1 + \delta) \|e_j\| + \delta \|T\| \leq (\|e\| + \delta)(1 + \delta) + \delta \|T\|$, so

$$\|Te\| \leq \|e\| + 2\delta + \delta^2 + \delta \|T\| \quad (**)$$

We crudely estimate $\|T\|$ by $\|T\| \leq (1 + \delta) \sum_{j=1}^n \|e_j^*\| \leq 2 \sum_{j=1}^n \|e_j^*\|$ for $\delta < 1$. So, for δ sufficiently small, (*) and (**) give condition (1), proving the theorem. \square

Theorem 10.22. (Spectral Mapping theorem) Let $a \in \mathcal{A}$, \mathcal{A} a complex Banach algebra with identity. Recall: $\sigma(a) = \{\lambda \in \mathbb{C} : x - \lambda \cdot 1 \text{ is not invertible}\}$, and $\sigma(a)$ is compact and nonempty. If p is a polynomial, then $p(\sigma(a)) = \sigma(p(a))$.

Proof. See Thm. 10.26(4). □

Theorem 10.23. (Facts about ideals, etc.) An ideal $I \subseteq \mathcal{B}$ is closed under addition (within I) and multiplication by elements of \mathcal{B} . An ideal is called **maximal** if $I \neq \mathcal{B}$ and I is not contained in any larger proper ideal. Define the **singular elements** \mathcal{S} as the union of all maximal ideals in \mathcal{B} .

- (1) If $\{0\}$ is the only proper ideal in \mathcal{B} , then \mathcal{B} is a field.
- (2) If I is a maximal ideal in \mathcal{B} , then \mathcal{B}/I is a field.
- (3) For $a \in \mathcal{B}$, a is invertible if and only if a does not belong to any maximal ideal.
- (4) If $\alpha \in \tilde{\mathcal{B}}$, then $\alpha(1) = 1$. (Use $\alpha(\underline{1}) = \alpha(1)\alpha(1)$.)
- (5) If I is a proper ideal in \mathcal{B} , then \bar{I} is a proper ideal. (Let $a_n \rightarrow a$, $a_n, b \in I$, $a \in \bar{I}$. Then $ba = \lim ba_n \in \bar{I}$. Also, $I \subseteq \mathcal{S}$, $1 \notin \mathcal{S}$.)
- (6) If I is a maximal ideal then $I = \bar{I}$. (Follows from (5)).
- (7) Let B be a Banach space, and let K be a closed subspace of B . Then B/K with $\|x + K\| := \inf\{\|y\| : y \in x + K\}$ is a Banach space. Moreover, if B is a Banach algebra with identity and K is a closed proper two sided ideal in B , then B/K is a Banach algebra, with the norm as before.
- (8) For ξ a linear function on a Banach space B , ξ is continuous if and only if $\ker \xi$ is closed.
- (9) Any character α (i.e. $\alpha \in \tilde{\mathcal{B}}$) is continuous.
- (10) (Gelfand-Mazur) The only complex Banach algebra with unit which is a division algebra is \mathbb{C} .
- (11) There is a one to one correspondence between characters and maximal ideals given by $\alpha \mapsto \ker \alpha$

Proof of (7): The triangle inequality is mostly routine. Without loss of generality, suppose $\{x_n + K\}_{n \geq 1}$ is Cauchy with $\|x_n - x_m + K\| < 2^{-\min(n,m)}$. By definition of the norm, we may take $\{k_n\}_{n \geq 1} \subseteq K$ with $\|x_n - x_{n+1} - k_{n+1}\| < 2^{-n+1}$. Let $z_n = x_n + \sum_{j=1}^n k_j$. Then $\{z_n\}_{n \geq 1}$ is Cauchy, so $z_n \rightarrow x$ for some $x \in B$. We claim that $x_n + K \rightarrow x + K$ in B/K . This follows since

$$\|x_n - x + K\| \leq \|x_n - z_n + K\| + \|z_n - x + K\| = \|z_n - x + K\| \leq \|z_n - x\|$$

The second assertion is mostly routine.

Proof of (8): If ξ is continuous, $\xi^{-1}(0) = \ker \xi$ is closed. Conversely, if $K := \ker \xi$ is closed, then from (7), B/K is a Banach space. Define $\zeta: B/K \rightarrow \{\text{scalars}\}$ by $\zeta(x + K) = \xi(x)$ (this is well defined and linear with trivial kernel). For $y, z \in B/K, y, z \neq 0$ we therefore have $\zeta(\zeta(z)y - z\zeta(y)) = 0$, so $z = \frac{\zeta(z)}{\zeta(y)}y$, i.e. B/K is one-dimensional. So, ζ is clearly continuous. Also, $\Pi: B \rightarrow B/K$ defined by $\Pi(x) = x + K$ is also easily continuous, so $\xi = \zeta \circ \Pi$ is continuous.

Proof of (9): For α a character, $I := \{a \in \mathcal{B} : \alpha(a) = 0\} = \ker \alpha$ is an ideal that is proper, since $\alpha(1) = 1$. For any $a \in \mathcal{B}$, write $a = (a - \alpha(a)1) + \alpha(a)1$. This shows that $\dim_{\mathbb{C}}(\mathcal{B}/I) = 1$, so I is maximal, hence closed (by (6)), so α is continuous by (8).

Proof of (10): Let $a \in \mathcal{A}$, $\lambda \in \sigma(a)$. Then $a - \lambda 1$ is not invertible, so $a - \lambda 1 = 0$, so $a = \lambda 1$, i.e. \mathcal{A} is scalar multiples of 1.

Proof of (11): From the proof of (9), we see that $\ker \alpha$ is a maximal ideal. Now, let I be a maximal ideal. I is closed from (6), so \mathcal{B}/I is a field from (2) and a complex Banach algebra from (7). From (10), this means \mathcal{B}/I is isomorphic to \mathbb{C} . Let $\beta: \mathcal{B} \rightarrow \mathcal{B}/I \approx \mathbb{C}$ be the natural (projection) homomorphism. By definition, β is a character, and $I = \ker \beta$, so any maximal ideal is the kernel of some character. This shows surjectivity. For injectivity, if $\ker \alpha = \ker \beta = I$, then $1 \notin I$, and from the proof of (9), we may write any $a \in \mathcal{B}$ as $a = c + \lambda 1$, $c \in I$, $\lambda \in \mathbb{C}$. Then $\alpha(a) = \lambda = \beta(a)$. \square

Theorem 10.24. (Characterization of Characters) *Let \mathcal{B} be a commutative Banach algebra with identity. Then the set of characters $\tilde{\mathcal{B}}$ is a closed subset of the unit ball of \mathcal{B}^* in the weak* topology. In particular, $\tilde{\mathcal{B}}$ is a compact Hausdorff space in this topology.*

Proof. Let $\alpha \in \tilde{\mathcal{B}}$. We first show that $\|\alpha\| \leq 1$. Let $a \in \mathcal{B}$, $\|a\| \leq 1$. Then $\|a^n\| \leq 1$, so $\{a^n\}$ is a bounded set. Since $|\alpha(a^n)| = |\alpha(a)|^n$ we must have $|\alpha(a)| \leq 1$, via Thm. 10.23(9). Therefore, $\tilde{\mathcal{B}}$ is contained in the unit ball of \mathcal{B}^* . Now, using the same reasoning as in Alaoglu's Theorem (Thm. 10.12), $\{\xi \in \mathcal{B}^*: \xi(ab) = \xi(a)\xi(b)\}$ (for $a, b \in \mathcal{B}$ fixed) is closed in the weak* topology. Therefore, $\bigcap_{a, b \in \mathcal{B}} \{\xi \in \mathcal{B}^*: \xi(ab) = \xi(a)\xi(b)\}$ is weak* closed. Similarly, $\{\xi \in \mathcal{B}^*: \xi(1) = 1\}$ is weak* closed, so $\tilde{\mathcal{B}}$ is weak* closed (recall that we defined characters to be nonzero). \square

Theorem 10.25. (Gelfand) *The canonical map is a homomorphism from \mathcal{B} into $C(\tilde{\mathcal{B}})$ with norm at most one. (Recall the definition: $a \mapsto \hat{a}$, where $\hat{a}(\alpha) := \alpha(a)$.)*

Proof. $\widehat{ab}(\alpha) = \alpha(ab) = \alpha(a)\alpha(b) = \hat{a}(\alpha)\hat{b}(\alpha)$. So, the Gelfand map is a homomorphism. For all $\alpha \in \tilde{\mathcal{B}}$ we have $|\hat{a}(\alpha)| = |\alpha(a)| \leq \|a\|$ from Thm. 10.24, so $\|\hat{a}\|_\infty \leq \|a\|$. \square

Theorem 10.26. (Facts about the Gelfand map, etc.)

- (1) *The kernel of the canonical map is the radical of \mathcal{B} (i.e. the intersection of maximal ideals). So, the canonical map is injective if and only if the radical is trivial. (Surjectivity will be dealt with below.)*
- (2) $\widehat{1}(\alpha) = \alpha(1) = 1$ for all $\alpha \in \tilde{\mathcal{B}}$ (Use Thm. 10.23(4))
- (3) $\lambda \in \sigma(a)$ if and only if $\lambda \in \text{range of } \hat{a}$, i.e. $\sigma(a) = \mathcal{R}(\hat{a})$
- (4) *The Spectral mapping theorem (Thm. 10.22) follows from (3).*
- (5) $r(a) = \|\hat{a}\|_\infty \leq \|a\|$, from (3) and Thm. 10.25. Therefore, $r(a+b) \leq r(a) + r(b)$ and $r(ab) \leq r(a)r(b)$.
- (6) *The following are equivalent: $a \in \text{radical}$, $\hat{a} = 0$, $\|\hat{a}\|_\infty = 0$, and $r(a) = 0$. (Apply (1) and (5)).*
- (7) $r(a^n) = r(a)^n$, and $r(a) = \lim_{n \rightarrow \infty} \|a^n\|^{1/n}$
- (8) $\|\hat{a}\|_\infty = \|a\|$ for all $a \in \mathcal{B}$ if and only if $\|a^2\| = \|a\|^2$ for all $a \in \mathcal{B}$.
- (9) *Let \mathcal{B} be a symmetric Banach algebra ($1 + a^*a$ is invertible). If a is Hermitian, then a is real. If a is strongly positive, then a is positive.*
- (10) *Let \mathcal{B} be a commutative $*$ algebra with unit. Then the following are equivalent: (i) \mathcal{B} is symmetric, (ii) Hermitian implies real, (iii) $\widehat{(a^*)}(\alpha) = \overline{\hat{a}(\alpha)}$, and (iv) Every maximal ideal is closed under $*$.*

Proof of (1): If $\hat{a} = 0$ then $\alpha(a) = 0$ for all α . So $a \in \ker \alpha \forall \alpha$, so a is in every maximal ideal by Thm. 10.23(11). We can reverse this argument to see that the radical is in the kernel of the canonical map.

Proof of (3): a is invertible, if and only if a is not in a maximal ideal, if and only if $\hat{a}(\alpha) \neq 0$ for each α (from (1)). So, by negation, $\lambda \in \sigma(a)$, if and only if $a - \lambda 1$ is in some maximal ideal, if and only if there exists some α with $\alpha(a - \lambda 1) = 0$, if and only if $\hat{a}(\alpha) - \lambda \hat{1}(\alpha) = 0$, if and only if $\hat{a}(\alpha) = \lambda$, from (2).

Proof of (4): Using (3) and Thm. 10.25,

$$\sigma(p(\widehat{(a)})) = \mathcal{R}(\widehat{p(a)}) = \mathcal{R}(p(\widehat{a})) = p(\sigma(a))$$

Proof of (7): From the remarks at the beginning of this Section, $\sigma(a)$ is compact, so $\exists \lambda \in \sigma(a)$ with $|\lambda| = r(a)$. So, $\lambda^n \in \sigma(a^n)$ by (4) (Thm. 10.22), so $r(a^n) \geq |\lambda^n| = r(a)^n$. Conversely, $\exists \lambda_0 \in \sigma(a^n)$ with $r(a^n) = |\lambda_0|$. By (4), $\exists \lambda \in \sigma(a)$ with $\lambda^n = \lambda_0$, so $r(a)^n \geq |\lambda|^n = |\lambda_0| = r(a^n)$.

To prove the second assertion, we use power series. Let ξ be a linear functional. For λ small enough, $(1 - \lambda a)^{-1} = \sum_{n \geq 0} a^n \lambda^n$, and $\xi((1 - \lambda a)^{-1}) = \sum_{n \geq 0} \xi(a^n) \lambda^n$. From the discussion near the beginning of this Section, $\xi((1 - \lambda a)^{-1})$ is analytic for $(1/\lambda) \notin \sigma(a)$, so the infinite sum converges for $|\lambda| < 1/r(a)$. For fixed ξ and λ as above, $\{|\lambda^n| |\xi(a^n)| : n = 0, 1, 2, \dots\}$ is therefore a bounded set. So, the Uniform boundedness principle (Thm. 10.6) says $\{\lambda^n a^n\}$ is a bounded set, i.e. $\|\lambda^n a^n\| \leq K$, $K > 0$. So, $\limsup \|a^n\|^{1/n} \leq 1/|\lambda|$ for $r(a) < 1/|\lambda|$, so $\limsup \|a^n\|^{1/n} \leq r(a)$. But $r(a)^n = r(a^n) \leq \|a^n\|$ from (5), so $r(a) \leq \liminf \|a^n\|^{1/n}$.

Proof of (8): From (5), $\|\widehat{a}\|_\infty = \|a\|$ if and only if $r(a) = \|a\|$. Now, if $\|a^2\| = \|a\|^2$ for all a , then $\|a\| = \|a^{2^n}\|^{1/2^n}$, so $\|a\| = \lim_n \|a^{2^n}\|^{1/2^n} = r(a)$ from (7). If $r(a) = \|a\|$ for all a , then $\|a^2\| = r(a^2) = r(a)^2 = \|a\|^2$, from (7).

Proof of (9): Let $a = a^*$. By scaling, it suffices to show $a - i$ is invertible. But $(a - i)(a + i)(1 + a^*a)^{-1} = 1$ and $(1 + a^*a)(a + i)(a - i) = 1$, so $a - i$ is invertible. For the second assertion, let $a = b^*b$. Then $a = a^*$, so by the first assertion, a is real. Finally, for $\alpha < 0$ we write $b^*b - \alpha = -\alpha \left(\left(\frac{b}{\sqrt{-\alpha}} \right)^* \left(\frac{b}{\sqrt{-\alpha}} \right) + 1 \right)$, which is invertible, so $\sigma(a) \subseteq [0, \infty)$.

Proof of (10): (i) implies (ii) follows from (9). To show (ii) implies (iii), let $f = a + a^*$, $g = i(a - a^*)$. Then f, g are Hermitian, so $\sigma(f), \sigma(g) \subseteq \mathbb{R}$. Using the Gelfand map and (3), $\alpha(f), \alpha(g) \in \mathbb{R}$ for a character α , i.e. $\alpha(f) = \overline{\alpha(f)}, \alpha(g) = \overline{\alpha(g)}$. So $\alpha(a) + \alpha(a^*) = \overline{\alpha(a) + \alpha(a^*)}$ and $i(\alpha(a) - \alpha(a^*)) = \overline{-i(\alpha(a) - \alpha(a^*))}$, i.e. $-\alpha(a) + \alpha(a^*) = \overline{\alpha(a) - \alpha(a^*)}$. Adding the first and third equality, $\alpha(a) = \overline{\alpha(a^*)}$.

To show (iii) implies (iv), let I be a maximal ideal. By Thm. 10.23(11), $I = \ker \alpha$ for some character α . For $a \in I$, $0 = \alpha(a) = \overline{\alpha(a)} = \alpha(a^*)$, so $a^* \in I$. To show (iv) implies (i), first observe that (iv) implies (iii). For α a character, let $b = c - \alpha(c)1$, so $\alpha(b) = 0$. Then $b, b^* \in \ker \alpha$, i.e. $\alpha(b^*) = \alpha(c^*) - \alpha((\alpha(c)1)^*) = \alpha(c^*) - \overline{\alpha(c)} = 0$, proving (iii). Now, $\alpha(c^*c) = \alpha(c^*)\alpha(c) = |\alpha(c)|^2$, so $\alpha(1 + c^*c) = 1 + |\alpha(c)|^2 \neq 0$, so $1 + c^*c$ is not in any maximal ideal by Thm. 10.23(11), i.e. it is invertible. \square

Theorem 10.27. (Gelfand Map for a B^* algebra) *If \mathcal{B} is a commutative B^* algebra with identity, then the canonical map (Gelfand map) is an isometric isomorphism onto $C(\tilde{\mathcal{B}})$. So, using Thm. 10.24, \mathcal{B} is isometrically isomorphic to the algebra of complex valued functions on a compact Hausdorff space.*

Proof. Claim 1: If \mathcal{B} is a commutative $*$ -multiplicative Banach algebra with identity, then $\|a\| = r(a)$ for all $a \in \mathcal{B}$. (From the remarks at the beginning of this Section, a B^* algebra is $*$ -multiplicative. By Thm. 10.26(3), $r(a) = \|a\|$ if and only if $\|\widehat{a}\| = \|a\|$, so the Gelfand map is an isometry with complete image.)

Claim 2: A commutative B^* algebra with identity is symmetric and semisimple. (In particular, via Thm. 10.26(1), the Gelfand map is injective).

Claim 3: If \mathcal{B} is commutative, and symmetric (with unit), the image of \mathcal{B} under the canonical map is dense in $C(\widetilde{\mathcal{B}})$. (Combining this with the first claim gives surjectivity, proving the theorem).

Proof of Claim 1: Let b be Hermitian. Then $\|b^2\| = \|b^*b\| = \|b^*\| \|b\| = \|b\|^2$, $\|b^{2^n}\| = \|b\|^{2^n}$, so $r(b) = \|b\|$ from Thm. 10.26(7). For a arbitrary, a^*a is Hermitian, so $\|a^*\| \|a\| = \|a^*a\| = r(a^*a) \leq r(a^*)r(a) \leq \|a^*\| r(a)$, using Thm. 10.26(5) twice. So, $\|a\| \leq r(a)$, and Thm. 10.26(5) says $r(a) \leq \|a\|$, so $\|a\| = r(a)$.

Proof of Claim 2: Semisimplicity (radical = 0) follows from Thm. 10.26(6) and Claim 1. Now, using Thm. 10.26(10), it suffices to show: if $a = a^*$, then a is real. By scaling, it suffices to show $a - i$ is invertible, i.e. $1 + ia$ is invertible, i.e. $1 \notin \sigma(-ia)$, i.e. $\lambda + 1 \notin \sigma(\lambda - ia)$ (for some $\lambda \in \mathbb{R}$). If $\lambda + 1 \in \sigma(\lambda - ia)$ for all $\lambda \in \mathbb{R}$, then since $\sup_{\gamma \in \sigma(b)} |\gamma| =: r(b) \leq \|b\|$ from Thm. 10.26(5), we have

$$(\lambda + 1)^2 \leq \|\lambda - ia\|^2 = \|(\lambda + ia)(\lambda - ia)\| = \|\lambda^2 + a^2\| \leq \lambda^2 + \|a^2\|$$

using $*$ -multiplicativity and the triangle inequality. But then $2\lambda + 1 \leq \|a^2\|$, which is a contradiction for λ large. Therefore, $\lambda + 1 \notin \sigma(\lambda - ia)$ for some $\lambda \in \mathbb{R}$.

Proof of Claim 3: Let $\alpha_1 \neq \alpha_2 \in \widetilde{\mathcal{B}}$, $\beta_1, \beta_2 \in \mathbb{C}$. Let $a \in \mathcal{B}$ such that $\alpha_1(a) \neq \alpha_2(a)$. Choose (using basic linear algebra) λ, μ such that $\lambda\alpha_1(a) + \mu = \beta_1$, $\lambda\alpha_2(a) + \mu = \beta_2$. Let $b := \lambda a + \mu$. Then $\widehat{b}(\alpha_1) = \beta_1$, $\widehat{b}(\alpha_2) = \beta_2$. The claim then follows by Stone-Weierstrass Theorem. Note that $\widetilde{\mathcal{B}}$ is a compact Hausdorff space from Thm. 10.24, the image of the Gelfand map clearly contains constants, and Thm. 10.26(10)(iii) shows that this image is closed under conjugation. \square

Theorem 10.28. (*Facts about maximal abelian self adjoint (m.a.s.a.) algebras, etc.*) (Recall \mathcal{A}' is a set of commutators, and \mathcal{M} denotes the multiplication algebra on L_2 .)

- (1) $\mathcal{A} \subseteq B(H)$ is maximal abelian if and only if $\mathcal{A} = \mathcal{A}'$. If $\mathcal{A} \subseteq B(H)$ is m.a.s.a. then $\mathcal{A} = \mathcal{A}'$.
- (2) A m.a.s.a algebra \mathcal{A} is a C^* algebra (and a B^* algebra).
- (3) Let (X, μ) be a σ -finite measure space. Then $\mathcal{M}(X, \mu)$ is a m.a.s.a. algebra.
- (4) Let \mathcal{A} be any $*$ subalgebra of $B(H)$. Suppose K is a closed subspace of H and P is the projection on K . Then K is invariant under \mathcal{A} if and only if $P \in \mathcal{A}'$.
- (5) If H is separable and \mathcal{A} is m.a.s.a. on H then \mathcal{A} has a cyclic vector. ($\exists z$ with $\mathcal{A}z$ dense in H .)

Proof of (1): Let \mathcal{A} be maximal abelian and let $B \in \mathcal{A}'$. We first note that $id, 0 \in \mathcal{A}$, since the set of operators of the form A or $A + id$, with $A \in \mathcal{A}$ or $A = 0$, contains \mathcal{A} and is abelian. Also, the set of operators of the form $A_0 + A_1B + \cdots + A_nB^n$ is an abelian algebra containing \mathcal{A} , so $B \in \mathcal{A}$, i.e. $\mathcal{A}' \subseteq \mathcal{A}$, so $\mathcal{A} = \mathcal{A}'$. Conversely, if $C \supseteq \mathcal{A}$ is an abelian algebra, then $C \subseteq \mathcal{A}' = \mathcal{A}$, so $C = \mathcal{A}$, i.e. \mathcal{A} is maximal abelian.

Let \mathcal{A} be m.a.s.a., and let $B \in \mathcal{A}'$. We first note that $id, 0 \in \mathcal{A}$, since the set of operators of the form A or $A + id$, with $A \in \mathcal{A}$ or $A = 0$, contains \mathcal{A} and is abelian and self adjoint. Now, for $A \in \mathcal{A}$, $A^* \in \mathcal{A}$, so $AB = BA$, $A^*B = BA^*$, so $B^*A^* = A^*B^*$, $B^*A = AB^*$, so $B^* \in \mathcal{A}'$. Write $A = (A + A^*)/2 + i(-i(A - A^*)/2) =: X + iY$. Then $X, Y \in \mathcal{A}'$, $X^* = X$, $Y^* = Y$. The set of operators of the form $A_0 + A_1X + \cdots + A_nX^n$, $A_j \in \mathcal{A}$ is a commutative self adjoint algebra containing \mathcal{A} , so it is \mathcal{A} . This set of operators also contains X , since $1 \in \mathcal{A}$. Therefore, $X \in \mathcal{A}$. Similarly, $Y \in \mathcal{A}$, so $B \in \mathcal{A}$, i.e. $\mathcal{A}' \subseteq \mathcal{A}$. By commutativity of \mathcal{A} , $\mathcal{A} \subseteq \mathcal{A}'$. Therefore $\mathcal{A} = \mathcal{A}'$.

Proof of (2): Let $A_n \in \mathcal{A}$, $A_n \rightarrow A$. For any $B \in \mathcal{A}$ we have $AB - BA = \lim(A_nB - BA_n) = 0$, so $A \in \mathcal{A}' = \mathcal{A}$, from (1). Also, by definition of m.a.s.a., $A \in \mathcal{A}$ implies $A^* \in \mathcal{A}$. Finally, $\|A^*A\| = \|A\|^2$, from the remarks at the beginning of this Section.

Proof of (3): Assume $\mu(X) < \infty$. Let $T \in (\mathcal{M}(X, \mu))'$. Let $g = T(1)$. If $f \in L_\infty$ then $TM_f1 = M_fT1$, so $T(f) = fg$, i.e. $Tf = M_gf$ for $f \in L_\infty$. By approximating appropriate level sets with indicator functions (as in the proof that $\|M_f\| = \|f\|_\infty$), we see that $\|g\|_\infty \leq \|T\|$. Since M_g and T are both bounded on L_2 , the equation $T|_{L_\infty} = M_g|_{L_\infty}$ extends by continuity to L_2 . Thus, $T \in \mathcal{M}(X, \mu)$ and $\mathcal{M}(X, \mu)$ is maximal abelian from (1). Since $M_g^* = M_{\bar{g}}$, $\mathcal{M}(X, \mu)$ is self-adjoint. The general σ -finite case follows the usual ‘‘piecing together’’ argument. Write $X = \cup_{j=1}^\infty X_j$, $\mu(X_j) < \infty$, X_j disjoint. If $T \in (\mathcal{M}(X, \mu))'$, T commutes with $M_{1_{X_j}}$, so T leaves the subspace $\{f \in L_2(X) : f|_{X_j^c} = 0\}$ ($= L_2(X_j)$) invariant, etc.

Proof of (4): If $P \in \mathcal{A}'$, $x \in K$ then $Ax = APx = PAx \in K$. Conversely, if $AK \subseteq K$ with $x \in H$, then $APx \in K$, so $APx = PAPx$. Since $A^* \in \mathcal{A}$, we similarly have $A^*P = PA^*P$. So, $PA = P^*A = (A^*P)^* = (PA^*P)^* = PAP = AP$, so $P \in \mathcal{A}'$. (A projection satisfies $P = P^*$ by Thm. 10.1(b)).

Proof of (5): Let $x \in H$, and let $\overline{\mathcal{A}x}$ be the smallest closed subspace containing $\mathcal{A}x$. Note that $id \in \mathcal{A}$ so $x \in \mathcal{A}x$. Since $\mathcal{A}x$ is invariant under \mathcal{A} , so is $\overline{\mathcal{A}x}$ (via a limiting argument). Suppose $y \perp \mathcal{A}x$. Since $(Ay, Bx) = (y, A^*Bx) = 0$, $\mathcal{A}y \perp \mathcal{A}x$. Let $E = \{x_\alpha\}$ be an orthonormal set such that $\mathcal{A}x_\alpha \perp \mathcal{A}x_\beta$ for $\alpha \neq \beta$. These sets exists (consider e.g. singletons), so Zorn’s lemma gives a maximal such set, E . Note that $H = \text{closed span}_\alpha \{\mathcal{A}x_\alpha\}$ (for if not, this would contradict maximality of E). Since H is separable, $E = \{x_\alpha\}_{\alpha \in \mathbb{N}}$ is countable. Set $z = \sum_{n \geq 1} 2^{-n}x_n$. This is our desired cyclic vector. To see this, let P_n project onto $\overline{\mathcal{A}x_n}$. From (4), $P_n \in \mathcal{A}'$, so $P_n \in \mathcal{A} = \mathcal{A}'$ by (1), so $\mathcal{A}z \supseteq \mathcal{A}P_nz = \mathcal{A}2^{-n}x_n = \mathcal{A}x_n$, so $\overline{\mathcal{A}z} \supseteq \text{closed span}_n \{\mathcal{A}x_n\} = H$. \square

Theorem 10.29. (Unitary Diagonalization of m.a.s.a algebra) *Let \mathcal{A} be a m.a.s.a. algebra on a separable Hilbert space H . Then \exists a finite measure space (X, μ) and a unitary operator $U : H \rightarrow L^2(X, \mu)$ such that $UAU^{-1} = \mathcal{M}(X, \mu)$*

Proof. Idea: Take ‘‘Rayleigh quotient’’ of (unit) cyclic vector (separating vector), use inverse Gelfand map. Using the Gelfand map again, define $U_0Az = \widehat{A}$.

Let z be a unit cyclic vector for \mathcal{A} (using Thm. 10.28(5)). Then z is also a **separating vector** for \mathcal{A} (i.e. if $A \in \mathcal{A}$ and $Az = 0$ then $A = 0$) since if $Az = 0$ then $\forall B \in \mathcal{A}$, $ABz = BAz = 0$, so $AAz = 0$, but $\mathcal{A}z$ is dense, so $A = 0$. From Thm. 10.28(2), \mathcal{A} is a B^* algebra (with identity). Let $X = \text{spectrum}(\mathcal{A}) = \widetilde{\mathcal{A}}$. From Thm. 10.27, the Gelfand map $A \mapsto \widehat{A}$ is an isometric isomorphism $\mathcal{A} \rightarrow C(X)$.

Define $\Lambda: C(X) \rightarrow \mathbb{C}$ by

$$\Lambda(\widehat{A}) := (Az, z)$$

Λ is a bounded linear functional, as it is the composition of bounded linear functionals. In fact, $|\Lambda(\widehat{A})| \leq \|A\| = \|\widehat{A}\|_\infty$ (using that the Gelfand map is an isometry). Λ is positive since

$$\Lambda(\overline{\widehat{A}}\widehat{A}) = \Lambda(\widehat{A^*A}) = \Lambda(\widehat{A^*A}) = (A^*Az, z) = \|Az\|^2 \geq 0$$

(using Claim 2 of Thm. 10.27, Thm.10.26(10)(iii), and Thm. 10.25). So, by Riesz's Representation Theorem and recalling that X is compact, \exists a unique regular Borel measure μ on X such that $\Lambda(\widehat{A}) = \int \widehat{A}d\mu$. Using Thm. 10.26(2), note that $\mu(X) = \int 1d\mu = \Lambda(1) = \|z\|^2 = 1$, so μ is a probability measure. Define $U_0: \mathcal{A}z \rightarrow L_2(X, \mu)$ by

$$U_0Az := \widehat{A}$$

U_0 is well defined since $Az = 0$ implies $A = 0$. U_0 is also linear and densely defined by Thm. 10.28(5). Moreover,

$$\|U_0Az\|^2 = \int \overline{\widehat{A}}\widehat{A}d\mu = \Lambda(\widehat{A^*A}) = (Az, Az) = \|Az\|^2$$

So, U_0 is an isometry, which extends by continuity to $U: H \rightarrow L_2(X, \mu)$, with $\|Ux\| = \|x\| \forall x \in H$. Since U is an isometry and H is a Hilbert space, the range of U is a complete (hence closed) subspace of $L_2(\mu)$ (which contains $C(X)$ by definition of the Gelfand map), so the range of U is all of $L_2(\mu)$, by the proof of Thm. 10.10(b). Thus, U is unitary.

Now, if $A, B \in \mathcal{A}$, $UAU^{-1}\widehat{B} = UABz = \widehat{AB} = M_{\widehat{A}}B$, so $UAU^{-1} = M_{\widehat{A}}$ on a dense set $\{\widehat{B}: B \in \mathcal{A}\} \subseteq L_2(X, \mu)$, so $UAU^{-1} = M_{\widehat{A}}$ on all of $L_2(X, \mu)$. Let $\mathcal{N} := UAU^{-1}$, and let $\mathcal{M} := \mathcal{M}(X, \mu)$. We just showed $\mathcal{N} \subseteq \mathcal{M}$. Now, if $T \in \mathcal{M}$, then $T \in \mathcal{N}'$, so $U^{-1}TU \in \mathcal{A}'$. But $\mathcal{A}' = \mathcal{A}$ (Thm. 10.28(1)), so $U^{-1}TU \in \mathcal{A}$, i.e. $T \in \mathcal{N}$, therefore $\mathcal{M} = \mathcal{N}$. \square

Theorem 10.30. (Spectral theorem, Multiplication Operator Form) Let $\{A_\alpha\}$ be a family of bounded normal operators on a complex separable Hilbert space H . Assume the family is commuting: $A_\alpha A_\beta = A_\beta A_\alpha \forall \alpha, \beta$, and $A_\alpha A_\beta^* = A_\beta^* A_\alpha \forall \alpha, \beta$. Then \exists a finite measure space (X, μ) and a unitary operator $U: H \rightarrow L_2(X, \mu)$ and for each $\alpha \exists$ a function $f_\alpha \in L_\infty$ such that $UA_\alpha U^{-1} = M_{f_\alpha}$

Proof. Let \mathcal{A}_0 be the algebra generated by the $\{A_\alpha, A_\alpha^*\}$. Note that \mathcal{A}_0 is a commutative $*$ algebra. Using Zorn's lemma, \exists a largest such $*$ algebra \mathcal{A} . We then claim that $\mathcal{A} = \mathcal{A}'$. If $B \in \mathcal{A}'$, then $B^* \in \mathcal{A}'$, so $C := B + B^* \in \mathcal{A}'$. Then the algebra generated by \mathcal{A} and C is commutative and self-adjoint, so $C \in \mathcal{A}$. Similarly, $i(B - B^*) \in \mathcal{A}$, so $B \in \mathcal{A}$, so $\mathcal{A}' = \mathcal{A}$, so \mathcal{A} is m.a.s.a., so we may apply Thm. 10.29. \square

Theorem 10.31. (Spectral theorem, Projection Valued Measure Form) Let A be a bounded normal operator on a separable Hilbert space H . There exists a unique projection valued Borel measure E on \mathbb{C} with compact support such that

$$A = \int_{\mathbb{C}} zdE$$

Moreover, if D is any bounded operator on H , then D commutes with A and A^* if and only if D commutes with $E(B)$ for all Borel sets B .

Proof idea: From the previous theorem, $UAU^{-1} = M_f$. Now, $G(B) := M_{1_{f^{-1}(B)}}$ is a projection valued measure. Define $E(B) := U^{-1}G(B)U$.

Theorem 10.32. (Spectral theorem for a bounded Hermitian operator) *If A is a bounded Hermitian operator on a separable Hilbert space H , then \exists a unique projection-valued Borel measure $E(\cdot)$ on the line with compact support such that*

$$A = \int_{-\infty}^{\infty} \lambda dE(\lambda)$$

And for all real Borel sets B , $E(B) \subseteq \sigma(A)$

Theorem 10.33. (Spectral theorem for a unitary operator) *Let U be a unitary operator on a separable Hilbert space. Then \exists a unique projection-valued Borel measure $E(\cdot)$ on $[0, 2\pi)$ such that*

$$U = \int_0^{2\pi} e^{i\theta} dE(\theta)$$

And $E(B) \subseteq \sigma(U)$ for all Borel sets B .

Theorem 10.34. (von Neumann's criteria for self adjointness) *Let T be a symmetric operator on a Hilbert space H . Then the following three statements are equivalent:*

- (a) T is self-adjoint
- (b) T is closed and $\ker(T^* + i) = \ker(T^* - i) = \{0\}$
- (c) $\text{Range}(T + i) = \text{Range}(T - i) = H$

Theorem 10.35. (Spectral Theorem, Unbounded case, Multiplication operator form) *Let T be a self-adjoint operator on a separable Hilbert space H . Then \exists a finite measure space (X, μ) , a unitary operator $U: H \rightarrow L_2(X, \mu)$ and a real valued measurable function f on X such that*

$$UTU^{-1} = M_f$$

Proof. (Sketch) From Thm. 10.34, $T + i$ is bijective from \mathcal{D}_T to H , so its inverse $(T + i)^{-1}$ exists. One can also check that $\|(T + i)\phi\|^2 = ((T + i)\phi, (T + i)\phi) = \|T\phi\|^2 + \|\phi\|^2 \geq \|\phi\|^2$, so $(T + i)^{-1}$ is bounded, and we can apply the bounded spectral theorem and write $U(T + i)^{-1}U = M_g$. Since $(T + i)^{-1}$ is bijective, so is M_g . In particular, g can be zero only on a set of measure zero. Therefore $f := (1/g) - i$ is well-defined. This is our desired f . \square

Theorem 10.36. (Spectral Theorem, Unbounded case, Projection valued measure form) *Let T be a self-adjoint operator on a separable complex Hilbert space H . Then \exists a projection valued measure $E(\cdot)$ on the Borel sets of the line such that*

$$T = \int_{-\infty}^{\infty} \lambda dE(\lambda)$$

Theorem 10.37. (Facts about Compact Operators, etc.) *Let $A: H \rightarrow K$ be a linear map between two Banach spaces.*

- (1) *If A is compact then A is bounded.*
- (2) *Let A be compact and let $B(0, n) = \{x: \|x\| \leq n\}$. Then $\overline{A(B(0, n))}$ is compact and therefore separable. Since the range of A is $\cup_{n \geq 1} \overline{A(B(0, n))}$, the range of A is separable.*

- (3) Let $\{A_n\}$ be a sequence of compact operators such that $\|A_n - A\| \rightarrow 0$. Then A is compact.
- (4) Define $\alpha_n(A) := \inf\{\|A - A_n\| : A_n : H \rightarrow K, \text{rank}(A_n) < n\}$, where $n = 1, 2, \dots$. For $n, m \geq 1$ and for all operators A, B :

$$\begin{aligned}\alpha_{n+m-1}(A + B) &\leq \alpha_n(A) + \alpha_m(B) \\ \alpha_{n+m-1}(A \circ B) &\leq \alpha_n(A) \cdot \alpha_m(B)\end{aligned}$$

- (5) For $0 < p < \infty$ define $A^p(X, Y)$ as the set of all operators $A : X \rightarrow Y$ with $a_p(A) := (\sum_{n=1}^{\infty} \alpha_n(A)^p)^{\frac{1}{p}} < \infty$. The quantity $a_p(A)$ is a quasi-norm that makes $A^p(X, Y)$ a quasi-Banach operator ideal.
- (6) If $A \in A^p(X, Y)$ and $B \in A^q(Y, Z)$ then $BA \in A^s(X, Z)$ with $1/s = 1/p + 1/q$.
- (7) If $A : X \rightarrow Y$ is compact then $\alpha_n(A) = \alpha_n(A^*)$. In particular, A^* is also compact, by (3).

Proof of (3): Let $\{x_n\}$ with $\|x_n\| \leq 1$. A diagonalization argument gives a subsequence $\{y_n\} \subseteq \{x_n\}$ such that for all j , $A_j y_n$ converges as $n \rightarrow \infty$. Then

$$\begin{aligned}\|Ay_k - Ay_\ell\| &\leq \|Ay_k - A_i y_k\| + \|A_i y_k - A_i y_\ell\| + \|A_i y_\ell + Ay_\ell\| \\ &\leq 2\|A - A_i\| + \|A_i y_k - A_i y_\ell\|\end{aligned}$$

So $\limsup_{k, \ell \rightarrow \infty} \|Ay_k - Ay_\ell\| \leq 2\|A - A_i\|$ which can be made arbitrarily small. Therefore, $A(B(0, 1))$ is a compact set, as desired.

Proof of (4): Given A_n, B_m of rank less than n, m respectively, consider $(A + B) - (A_n + B_m)$. The first inequality follows. For the second inequality, note that

$$\|AB - (A_n B + AB_m - A_n B_m)\| = \|(A - A_n)(B - B_m)\| \leq \|A - A_n\| \|B - B_m\|$$

So, using that $\text{rank}(A_n B + AB_m - A_n B_m) \leq \text{rank}(A_n(B - B_m)) + \text{rank}(AB_m) < n + m - 1$, we get

$$\alpha_{n+m-1}(AB) \leq \inf_{A_n, B_m} \|A - A_n\| \|B - B_m\| = \alpha_n(A) \alpha_n(B)$$

Proof of (5): Using that the α_n 's are decreasing in n , (4), and the L_p (quasi)-triangle inequality,

$$\begin{aligned}a_p(A + B) &= \left(\sum_{n=1}^{\infty} \alpha_n(A + B)^p \right)^{\frac{1}{p}} \leq \left(2 \sum_{n=1}^{\infty} \alpha_{2n-1}(A + B)^p \right)^{\frac{1}{p}} \\ &\leq 2^{\frac{1}{p}} \left(\sum_{n=1}^{\infty} (\alpha_n(A) + \alpha_n(B))^p \right)^{\frac{1}{p}} \leq C_p (a_p(A) + a_p(B))\end{aligned}$$

Since also $a_p(\lambda A) = |\lambda| a_p(A)$, we see that $a_p(A)$ is a quasi-norm and $A^p(X, Y)$ is a linear subspace of $L(X, Y)$. Since $\alpha_n(A) = 0$ if $\text{rank}(A) < n$, $A^p(X, Y)$ contains finite rank operators.

Since $\alpha_1(A) = \|A\|$ and the sequence $\{\alpha_n(A)\}$ is decreasing in n , (4) shows $\alpha_n(CBA) \leq \alpha_n(C) \alpha_n(B) \alpha_n(A) \leq \|C\| \alpha_n(B) \|A\|$, so $a_p(CBA) \leq \|C\| a_p(B) \|A\|$. In particular, if $B \in A^p(X, Y)$, $A \in A^p(Z, X)$ and $C \in A^p(Y, V)$, then $CBA \in A^p(Z, V)$. For any rank-one operator $A : X \rightarrow Y$ of the form $A(x) = x^*(x) \cdot y$, $a_p(A) = \alpha_1(A) = \|A\| = \|x^*\| \|y\|$.

It remains to show that $(A^p(X, Y), a_p(\cdot))$ is complete. This procedure is mostly routine. Starting with a Cauchy sequence, take a rapidly convergent subsequence, use $\alpha_1(\cdot) = \|\cdot\|$ to get a limit, use the (quasi)-triangle inequality for ℓ_p , and so on.

Proof of (6): From (4) and Hölder's inequality,

$$\begin{aligned} a_s(BA) &= \left(\sum_{n=1}^{\infty} \alpha_n(BA)^s \right)^{\frac{1}{s}} \leq 2^{\frac{1}{s}} \left(\sum_{n=1}^{\infty} \alpha_{2n-1}(BA)^s \right)^{\frac{1}{s}} \\ &\leq 2^{\frac{1}{s}} \left(\sum_{n=1}^{\infty} \alpha_n(A)^s \cdot \alpha_n(B)^s \right)^{\frac{1}{s}} \leq 2^{\frac{1}{s}} a_p(A) a_q(B) \end{aligned}$$

Proof of (7): If A_n has $\text{rank}(A_n) < n$ with $\|A - A_n\| < \alpha_n(A) + \varepsilon$ then A_n^* has $\text{rank}(A_n^*) < n$ (if S is a finite independent spanning set for $A_n(X)$, then $A_n^*(y^*)$ is determined by evaluating y^* on S , so apply Thm. 10.17(c)). So, $\alpha_n(A^*) \leq \|A^* - A_n^*\| = \|A - A_n\| < \alpha_n(A) + \varepsilon$, so $\alpha_n(A^*) \leq \alpha_n(A)$. It remains to show that $\alpha_n(A^*) \geq \alpha_n(A)$. It suffices to show $\alpha_n(A^{**}) \geq \alpha_n(A)$, since combining this with our first inequality gives $\alpha_n(A^*) \geq \alpha_n(A^{**}) \geq \alpha_n(A)$.

Fix $n, \varepsilon > 0$, and V a finite ε -net in $T(B_X)$. Since T is compact, Goldstine, Thm. 10.18, and Hahn-Banach, Thm. 10.4(d) show that $\iota(T(B_X))$ is norm-dense in $T^{**}(B_{X^{**}})$, so $\iota(V)$ is also an ε -net in $T^{**}(B_{X^{**}})$. Let $T_n: X^{**} \rightarrow Y^{**}$ with $\text{rank}(T_n) < n$ and $\|T^{**} - T_n\| \leq \alpha_n(T^{**}) + \varepsilon$. Let $E = \text{span}\{\iota(V) \cup T_n(X^{**})\}$. From the principle of local reflexivity, Thm. 10.21, there exists $\phi: E \rightarrow Y$ with $\|\phi\| \leq 1 + \varepsilon$ and $\phi|_{E \cap \iota(Y)} = id$, so $\phi|_{\iota(V)} = id$. For $x \in X$ with $\|x\| \leq 1$ fix $v \in V$ with $\|Tx - v\| \leq \varepsilon$. Then

$$\begin{aligned} \|Tx - \phi T_n \iota x\| &\leq \varepsilon + \|v - \phi T_n \iota x\| \leq \varepsilon + (1 + \varepsilon) \|\iota(v) - T_n \iota x\| \\ &\leq \varepsilon + (1 + \varepsilon)(\|\iota(v) - \iota(Tx)\| + \|\iota(Tx) - T_n \iota x\|) \\ &\leq \varepsilon + (1 + \varepsilon)(\varepsilon + \|T^{**} \iota x - T_n \iota x\|) \\ &\leq \varepsilon + (1 + \varepsilon)(\varepsilon + \varepsilon + \alpha_n(T^{**})) \end{aligned}$$

using that ι is a linear isometry, that $\iota(Tx) = T^{**} \iota(x)$, and the definition of T_n . Since $\phi T_n \iota$ has rank less than n and ε is arbitrary, we conclude that $\alpha_n(T) \leq \alpha_n(T^{**})$. \square

Theorem 10.38. (Riesz Theory for Compact Operators) Let $C: H \rightarrow H$ be a compact operator on a Banach space H . Let $B = id - C$ and $B^* = id - C^*$.

- (1) $\ker(B)$ is finite dimensional.
- (2) If B is injective then B is surjective.
- (3) $\dim \ker(B) = \dim \ker(B^*)$.
- (4) Every nonzero point λ of the spectrum of C is an eigenvalue of finite multiplicity. (That is $\dim \ker(\lambda - C)$ is finite.) Moreover, the multiplicity of λ for C is the same as for C^* . The only possible cluster point of the spectrum of C is zero.
- (5) If C has an infinite number of eigenvalues then 0 is a cluster point of eigenvalues. So, the eigenvalues can be arranged in a sequence converging to zero. (Just apply (4).)
- (6) If C is a compact normal operator on a separable complex Hilbert space H , then there is a finite or infinite sequence P_n of mutually orthogonal finite dimensional projections

such that

$$C = \sum_{n=1}^{\infty} \lambda_n P_n \quad (\text{or } C = \sum_{n=1}^k \lambda_n P_n)$$

where $\{\lambda_n\}$ are the nonzero eigenvalues of C and the series converges in the operator norm. Moreover, H has an orthonormal basis consisting of eigenvectors of C .

Proof of (1): We begin with a few claims.

Claim 1: A finite dimensional subspace $F \subseteq H$ is closed in H .

Claim 2: Let $H_0 \subseteq H$ be a closed proper subspace. For any $\varepsilon > 0$, $\exists x_0 \in H$ with $\|x_0\| = 1$ and $\|x - x_0\| \geq 1 - \varepsilon$ for all $x \in H_0$.

Claim 3: Any locally compact Banach space is finite dimensional.

To prove Claim 1, recall that any two norms on a finite dimensional linear space are equivalent, so the norm on F is equivalent to any Euclidean norm on F . Therefore F is complete in its own norm. Let $\{x_n\}_{n=1}^{\infty} \subseteq F$ be such that $x_n \rightarrow x \in H$. Then $\{x_n\}$ is Cauchy in F , but F is complete, so $x_n \rightarrow y \in F$. By uniqueness of limits in H , $x = y$. So $x \in F$ and F is closed in H .

To prove Claim 2, let $\varepsilon < 1$, $z_0 \notin H_0$, $d = \inf_{x \in H_0} \|x - z_0\|$. For any $\delta > 0$, $\exists z \in H_0$ with $\|z - z_0\| \leq d + \delta$. Let $\delta := \varepsilon d / (1 - \varepsilon)$, so that z is determined from the previous sentence, and let $x_0 := (z - z_0) / \|z - z_0\|$. Then $\|x_0\| = 1$ and if $x \in H_0$, then $(\|z - z_0\|)x - z \in H_0$, so

$$\|x - x_0\| = \frac{\|(\|z - z_0\|)x - z + z_0\|}{\|z - z_0\|} \geq \frac{d}{\|z - z_0\|} \geq \frac{d}{d + \delta} = 1 - \varepsilon$$

To prove Claim 3, we instead prove its contrapositive. Assume H is infinite dimensional. We construct a sequence $x_1, x_2, \dots \subseteq H$ with $\|x_n\| = 1$, $\|x_i - x_j\| \geq 1/2$, $i \neq j$. Let x_1 with $\|x_1\| = 1$. Given x_1, \dots, x_n , let $H_0 = \text{span}\{x_1, \dots, x_n\}$. By Claim 1, H_0 is closed. By Claim 2, $\exists x_{n+1}$ such that $\|x_i - x_{n+1}\| \geq 1/2$ for $i = 1, \dots, n$. The sequence $\{x_n\}$ has no Cauchy subsequence, so the closed unit ball of H is not compact. By scaling, the closed ball of radius $r > 0$ is also not compact.

We can now finally prove (1). Let $H_0 = \ker(B)$. Then $x \in H_0$ if and only if $Cx = x$, i.e. $C|_{H_0} = id$. But C maps the unit ball of H_0 into a totally bounded set, i.e. the unit ball of H_0 is locally compact. Therefore H_0 is finite dimensional by Claim 3.

Proof of (2): We first claim that $B = id - C$ satisfies $\|Bx\| \geq m \|x\|$ for all $x \in H$. If not, then there exists a sequence $\{y_n\}$ with $\|y_n\| = 1$ and $By_n \rightarrow 0$. Since C is compact, we may take a subsequence $\{y_{n_j}\}$ so that Cy_{n_j} converges. Since $By_{n_j} = (id - C)y_{n_j}$ converges also, $(B + C)y_{n_j} = y_{n_j}$ converges. Let $y = \lim y_{n_j}$. Since $\|y_n\| = 1$, $\|y\| = 1$, but $By = \lim By_n = 0$ so B is not injective. Since we have arrived at a contradiction, the claim follows.

Let $M \subseteq H$ be a closed subspace. Then $B(M)$ is closed. To see this, suppose $Bx_n \rightarrow y$. Then $\{x_n\}$ is bounded since $\|Bx_n\| \geq m \|x_n\|$. We may take a subsequence so that Cx_{n_j} converges, so $(B + C)x_{n_j} = x_{n_j}$ converges. Let $x = \lim x_{n_j}$, so that $Bx = \lim Bx_{n_j} = y$, i.e. $y \in B(M)$ as desired.

Suppose for the sake of contradiction that $\text{Range}(B) \neq H$. Let $H_0 = H$, $H_1 = BH_0$, $H_2 = BH_1$, and so on. Then H_{m+1} is a closed and proper subspace of H_m . (If $E, F \subseteq H$ are closed with $F = BH$, $E = BF = B^2H$ and $f \in F$ then $Bf =: e \notin F$ implies $B^2f = Be$ with $B^2f \in E$ and $Be \notin E$ by injectivity, a contradiction. So $E \subseteq F$.) By Claim 2 of (1), \exists

$x_n \in H_n$ with $\|x_n\| = 1$ and $d(x_n, H_{n+1}) \geq 1/2$. If $n > m$,

$$Cx_m - Cx_n = x_m - Bx_m - x_n + Bx_n =: x_m - x, \text{ with } x \in H_{m+1}$$

So $\|Cx_m - Cx_n\| \geq d(x_m, H_{m+1}) \geq 1/2$. So the sequence $\{Cx_n\}$ contains no Cauchy subsequence, contradicting compactness of C . Therefore $\text{Range}(B) = H$.

Proof of (3): Using (1) and Thm. 10.37(7), let x_1, \dots, x_n be a basis for $\ker(B)$ and let η_1, \dots, η_ν be a basis for $\ker(B^*)$. By Hahn-Banach (Thm. 10.4(a)) $\exists \xi_j \in H^*$ with $\xi_j(x_i) = \delta_{ij}$, $i, j = 1, \dots, n$. Let $K = \ker B^* \subseteq H^*$ and let $R: H \rightarrow H^{**}$ be such that $R(x) = \iota(x)|_K$. We claim that $R: H \rightarrow K^*$ is surjective. If not, $\exists u \in K$, $u \neq 0$ such that, for all $x \in H$, $u(x) = R(x)(u) = 0$. But then $u = 0$, a contradiction. Therefore $R: H \rightarrow K^*$ is surjective. So using Hahn-Banach again and the surjectivity of R , $\exists y_1, \dots, y_\nu$ in H such that $\eta_j(y_i) = \delta_{ij}$, $i, j = 1, \dots, \nu$.

Suppose for the sake of contradiction that $n < \nu$ and define

$$C'x := Cx + \sum_{j=1}^n \xi_j(x)y_j$$

Since C' is the sum of two compact operators, C' is compact. Let $B' = id - C'$. We claim that B' is injective. To see this, suppose $B'x_0 = 0$. Since $B - B' = (id - C) - (id - C') = C' - C$, we see that $Bx_0 = \sum_{j=1}^n \xi_j(x_0)y_j$. Using $\eta_i \in \ker B^*$, the previous sentence, and the choice of the y_j ,

$$0 = B^*\eta_i(x_0) = \eta_i(Bx_0) = \sum_{j=1}^n \xi_j(x_0)\eta_i(y_j) = \xi_i(x_0)$$

Therefore $Bx_0 = 0$, so $x_0 = \sum_{j=1}^n \alpha_j x_j$ by definition of the x_j . But then $0 = \xi_i(x_0) = \alpha_i$ by definition of the ξ_i , so $x_0 = 0$. We conclude that $\ker B' = 0$.

By (2), B' is also surjective. So $\exists x \in H$ with $y_{n+1} = B'x$. Since $B - B' = C' - C$,

$$\begin{aligned} 1 &= \eta_{n+1}(y_{n+1}) = \eta_{n+1}(B'x) = \eta_{n+1}(Bx) - \eta_{n+1}\left(\sum_{j=1}^n \xi_j(x)y_j\right) \\ &= B^*(\eta_{n+1})(x) - \sum_{j=1}^n \xi_j(x)\eta_{n+1}(y_j) = 0 - 0 = 0 \end{aligned}$$

Since $\eta_{n+1} \in \ker B^*$ and using the definition of η_{n+1} . Since we have achieved a contradiction, we conclude that $n \geq \nu$.

Since $n \geq \nu$, $\dim \ker B \geq \dim \ker B^*$. Since C^* is compact by Thm. 10.37(7), we conclude that $\dim \ker B^* \geq \dim \ker B^{**}$. Since $B^{**}\iota x = \iota(Bx)$, $\dim \ker B^{**} \geq \dim \ker B$. Combining these inequalities concludes the proof.

Proof of (4): Suppose $\lambda \neq 0$ is in $\sigma(C)$. Then $id - \lambda^{-1}C$ is not invertible. Since $\lambda^{-1}C$ is compact, $1 - \lambda^{-1}C$ may not be invertible due to either a lack of injectivity (in which case λ is an eigenvalue, of finite multiplicity by (1)), or due to a lack of surjectivity (in which case $1 - \lambda^{-1}C$ is also not injective by (2), so the first case applies). If $1 - \lambda^{-1}C$ is bijective, then it is also invertible by the Open Mapping Theorem, Thm. 10.7 (or by the proof of (2)). By (3), $\dim \ker(\lambda - C) = \dim \ker(\lambda - C^*)$ for $\lambda \neq 0$, using $\ker(\lambda - C) = \ker(1 - \lambda^{-1}C)$.

We now prove by contradiction that zero is the only possible cluster point of $\sigma(C)$. Let $\lambda_n \in \sigma(C)$ with $\lambda_n \rightarrow \lambda \neq 0$. Without loss of generality, $\lambda_n \neq \lambda_m$ for $n \neq m$ and $\exists \varepsilon > 0$ with

$|\lambda_n| \geq \varepsilon$ for all n . By the part of the theorem already proven, $\exists x_n \neq 0$ with $Cx_n = \lambda_n x_n$. We claim that $\{x_n\}$ is linearly independent. If not, let N be the smallest integer such that a linear relation holds: $x_N = \sum_{j=1}^{N-1} \alpha_j x_j$ with some $\alpha_j \neq 0$. Then $\lambda_N x_N = Cx_N = \sum_{j=1}^{N-1} \alpha_j \lambda_j x_j$, i.e. $\lambda_N \sum_{j=1}^{N-1} \alpha_j x_j = \sum_{j=1}^{N-1} \alpha_j \lambda_j x_j$, i.e. $\sum_{j=1}^{N-1} \alpha_j (\lambda_N - \lambda_j) x_j = 0$. By minimality of N and the distinctness of the λ_j , $\alpha_j = 0$ for $j = 1, \dots, n-1$, a contradiction. Therefore, $\{x_n\}$ is linearly independent.

Let $H_n := \text{span}\{x_1, \dots, x_n\}$. Then H_n is a strictly increasing sequence of subspaces. By Claim 2 of (1), $\exists y_n \in H_n$ with $\|y_n\| = 1$ and $\|y_n - x\| \geq 1/2$ for all $x \in H_{n-1}$. Let $y \in H_n$ so that $y = \sum_{j=1}^n \alpha_j x_j$ and $Cy - \lambda_n y = \sum_{j=1}^n \alpha_j (\lambda_j - \lambda_n) x_j \in H_{n-1}$. Let $n > m$. Then

$$\begin{aligned} \|Cy_n - Cy_m\| &= \|(Cy_n - \lambda_n y_n) + \lambda_n y_n - \lambda_m y_m + (\lambda_m y_m - Cy_m)\| \\ &=: \|\lambda_n y_n - z\|, \text{ with } z \in H_{n-1} \\ &\geq \frac{|\lambda_n|}{2} \geq \frac{\varepsilon}{2} \end{aligned}$$

using the choice of y_n and the definition of ε . Therefore, $\{Cy_m\}$ contains no Cauchy subsequence, a contradiction. We conclude that $\lambda = 0$, as desired.

Proof of (6), sketch: From the Spectral Theorem, Thm. 10.31, we may write

$$C = \int_{\sigma(C)} \lambda dE(\lambda)$$

Let $\lambda_1, \lambda_2, \dots$ be the nonzero eigenvalues of C . Let $P_n := E(\{\lambda_n\})$. □

Theorem 10.39. (Facts about Semigroups of Operators, etc.) Let $T_t: H \rightarrow H$ be a strongly continuous contraction semigroup. Define H_0 as the set of $f \in H$ such that $\lim_{t \downarrow 0} T_t f = f$.

- (1) H_0 is closed
- (2) T_t leaves the subspace H_0 invariant. For $f \in H_0$, $T_t f$ is strongly continuous on $t \geq 0$ ($\lim_{s \rightarrow 0} T_{t+s} f = T_t f$). For $f \in \mathcal{D}_A$, $Af \in H_0$.
- (3) $H_0 = \overline{\mathcal{D}_A}$
- (4) If $f \in \mathcal{D}_A$, then $T_t f$ is strongly differentiable, and $\frac{dT_t f}{dt} = AT_t f = T_t A f$. Also, $T_t f - f = \int_0^t T_s A f ds$
- (5) The operator A is closed. (If $f_n \in \mathcal{D}_A$, if $f_n \rightarrow f$ and if $A f_n \rightarrow g$ then $f \in \mathcal{D}_A$ and $A f = g$).
- (6) Let $g \in H_0$, $\lambda > 0$. The equation $\lambda f - A f = g$ has exactly one solution $f \in \mathcal{D}_A$. Moreover, f satisfies $f = R_\lambda(g) := \int_0^\infty e^{-\lambda t} T_t g dt$. The operator R_λ is linear and $\|R_\lambda g\| \leq \frac{1}{\lambda} \|g\|$. (So, $(\lambda - A): \mathcal{D}_A \rightarrow H_0$ is bijective, and $(\lambda - A)^{-1} = R_\lambda$. R_λ is called the **resolvent** of A . Note also that $R_\lambda(H_0) = \mathcal{D}_A$. Moreover, the proof shows that $R_\lambda(H) \subseteq H_0$.)
- (7) Let U be a linear operator that extends A . Suppose \mathcal{D} satisfies: (i) $\mathcal{D}_A \subseteq \mathcal{D} \subseteq H_0$, (ii) $\mathcal{D} \subseteq \mathcal{D}_U$ and $U(\mathcal{D}) \subseteq H_0$, (iii) If $U f = f$ and $f \in \mathcal{D}$ then $f = 0$. Given such \mathcal{D} , we conclude that $\mathcal{D} = \mathcal{D}_A$.

Proof of (1): Let $f_n \in H_0$ with $f_n \rightarrow f$. Then

$$\|T_t f - f\| \leq \|T_t f - T_t f_n\| + \|T_t f_n - f_n\| + \|f_n - f\| \leq 2 \|f_n - f\| + \|T_t f_n - f_n\|$$

Now choose n so that $\|f_n - f\|$ is small, and then choose t so that $\|T_t f_n - f_n\|$ is small.

Proof of (2): Let $f \in H_0$. If $h \downarrow 0$, then $\|T_{t+h}f - T_t f\| = \|T_t(T_h f - f)\| \leq \|T_h f - f\| \rightarrow 0$. If $h \uparrow 0$, $\|T_{t+h} - T_t f\| = \|T_{t+h}(f - T_{-h}f)\| \leq \|f - T_{-h}f\| \rightarrow 0$. Therefore, $\lim_{h \rightarrow 0} T_{t+h}f = T_t f$. The first two assertions follow. Now, let $f \in H$. If $\lim_{t \downarrow 0} T_t f$ is not equal to f , or this limit does not exist, then Af is undefined. Therefore, $\mathcal{D}_A \subseteq H_0$. Let $f \in \mathcal{D}_A$ so that $f \in H_0$. Then the third assertion follows by the definition $A(f) := \lim_{h \downarrow 0} \frac{T_h f - f}{h}$ and from (1).

Proof of (3): As shown in (2), $\mathcal{D}_A \subseteq H_0$. It is therefore sufficient to show that any $f \in H_0$ can be approximated by elements of \mathcal{D}_A . If $f \in \mathcal{D}_A$, then $T_t f$ is strongly continuous for $t \geq 0$ by (2), so $T_t f$ is strongly integrable on any finite interval in $[0, \infty)$. Let $g_a := \int_0^a T_t f dt$. Observe

$$T_h g_a = \int_0^a T_h T_t f dt = \int_0^a T_{h+t} f dt = \int_h^{h+a} T_t f dt = g_a + \int_a^{a+h} T_t f dt - \int_0^h T_t f dt$$

Therefore

$$\lim_{h \downarrow 0} \frac{T_h g_a - g_a}{h} = \lim_{h \downarrow 0} \left(\frac{1}{h} \int_a^{a+h} T_t f dt - \frac{1}{h} \int_0^h T_t f dt \right) = T_a f - f$$

So, $g_a \in \mathcal{D}_a$. Finally, by definition of g_a , $\lim_{h \downarrow 0} g_h/h = f$, as desired.

Proof of (4): For $t, h \geq 0$,

$$T_t \left(\frac{T_h f - f}{h} \right) = \frac{T_t T_h f - T_t f}{h} = \frac{T_{t+h} f - T_t f}{h} = \frac{T_h T_t f - T_t f}{h} \quad (*)$$

If $f \in \mathcal{D}_A$, then $\lim_{h \downarrow 0} \frac{T_h f - f}{h} = Af$, so $\lim_{h \downarrow 0} T_t \left(\frac{T_h f - f}{h} \right) = T_t Af$. So, using (*), we conclude that $\lim_{h \downarrow 0} \frac{T_{t+h} f - T_t f}{h} =: \frac{d^+ T_t f}{dt}$, with $T_t f \in \mathcal{D}_A$, and with $\frac{d^+ T_t f}{dt} = AT_t f = T_t Af$.

We now need to show that $\lim_{h \downarrow 0} \frac{T_t f - T_{t-h} f}{h} = \frac{d^+ T_t f}{dt}$. For $t > h > 0$, we have

$$\begin{aligned} \left\| \frac{T_t f - T_{t-h} f}{h} - T_t Af \right\| &\leq \left\| T_{t-h} \left(\frac{T_h f - f}{h} - Af \right) \right\| + \|T_{t-h}(Af - T_h Af)\| \\ &\leq \left\| \frac{T_h f - f}{h} - Af \right\| + \|Af - T_h Af\| \end{aligned}$$

By (2), $Af \in H_0$, so, $\lim_{h \downarrow 0} \frac{T_t f - T_{t-h} f}{h} = T_t Af = \frac{d^+ T_t f}{dt}$, as desired. Finally, $\int_0^t T_s Af ds = \int_0^t \frac{dT_s f}{ds} ds = T_t f - f$, as desired.

Proof of (5): From (4), $T_t f_n - f_n = \int_0^t T_s A f_n ds$. Letting $n \rightarrow \infty$, (noting that $\|\int_0^t (T_s A f_n - T_s g) ds\| \leq \int_0^t \|T_s\| \|A f_n - g\| ds$), we get $T_t f - f = \int_0^t T_s g ds$. So $\lim_{t \rightarrow 0} \frac{T_t f - f}{t} = \lim_{t \downarrow 0} \frac{1}{t} \int_0^t T_s g ds = g$. So $f \in \mathcal{D}_A$ and $Af = g$, as desired.

Proof of (6): First note that $\int_0^\infty e^{-\lambda t} T_t g$ is well-defined since the integrand is strongly continuous by (2), with norm bounded by $e^{-\lambda t} \|g\|$. Linearity of R_λ follows by its definition. Note that $\|\int_0^\infty e^{-\lambda t} T_t g dt\| \leq \int_0^\infty e^{-\lambda t} \|g\| dt = \frac{1}{\lambda} \|g\|$. Also,

$$\begin{aligned} T_h f &= \int_0^\infty e^{-\lambda t} T_{t+h} g dt = \int_h^\infty e^{-\lambda(t-h)} T_t g dt \\ &= e^{\lambda h} \int_h^\infty e^{-\lambda t} T_t g dt = e^{\lambda h} \left(f - \int_0^h e^{-\lambda t} T_t g dt \right) \end{aligned}$$

by definition of f . So, $\frac{T_h f - f}{h} = \frac{e^{\lambda h} - 1}{h} f - \frac{1}{h} e^{\lambda h} \int_0^h e^{-\lambda t} T_t g dt$. Letting $h \downarrow 0$, the right side of this equality goes to $\lambda f - g$. Therefore, $f \in \mathcal{D}_A$, with $Af = \lambda f - g$, as desired.

We now treat uniqueness. Let g such that the equation $\lambda f - Af = g$ has two solutions. By subtracting these solutions, $\exists \phi \in \mathcal{D}_A$ with $\lambda\phi - A\phi = 0$. Since $d(T_t\phi)/dt = AT_t\phi = T_tA\phi = \lambda T_t\phi$ by (4), we conclude that $d(e^{-\lambda t}T_t\phi)/dt = e^{-\lambda t}\lambda T_t\phi + T_t\phi(-\lambda)e^{-\lambda t} = 0$. Integrating this in s , we get $\int_0^t \frac{d(e^{-\lambda s}T_s\phi)}{ds} ds = e^{-\lambda t}T_t\phi - \phi = 0$. So, $\|\phi\| \leq e^{-\lambda t} \|\phi\|$ for all $t > 0$, i.e. $\phi = 0$.

Proof of (7): Let $f \in \mathcal{D}$. Then $f - Uf =: g \in H_0$. By (6), $\exists \tilde{f} \in \mathcal{D}_A$ with $\tilde{f} - A\tilde{f} = g$. Since $A \subseteq U$, we conclude that $\tilde{f} - U\tilde{f} = g$ as well. Therefore, $h := f - \tilde{f} \in \mathcal{D}$ satisfies $Uh = h$, so $h = 0$, i.e. $f = \tilde{f} \in \mathcal{D}_A$. \square

Theorem 10.40. (Characterization of Semigroups) *Let $T_t: H \rightarrow H$ be a strongly continuous semigroup of bounded linear operators (so $H = H_0$). Then its infinitesimal generator A is a closed densely defined linear function. T_t is uniquely determined by A in the sense that distinct semigroups have distinct infinitesimal generators. Moreover, if $f \in \mathcal{D}_A$ then $u(t) := T_t f$ solves the differential equation*

$$\frac{du}{dt}(t) = Au(t), \quad t \geq 0, u(0) = f$$

Proof. Closedness of A follows from Thm. 10.39(5), and the density of \mathcal{D}_A in H_0 follows from Thm. 10.39(3). We now prove two claims

Claim 1: Let $\phi(t)$ be a Borel function such that for all $\lambda > 0$, $\int_0^\infty e^{-\lambda t}\phi(t)dt = 0$. Then $\phi(t) = 0$ for almost all t .

Claim 2: Let $u_t: [0, \infty) \rightarrow H$ be a strongly continuous function such that for all $\lambda > 0$, $\int_0^\infty e^{-\lambda t}u_t dt = 0$. Then $u_t = 0$ for all t .

Proof of Claim 1: Letting $\lambda = n = 1, 2, \dots$, and substituting $z = e^{-t}$ gives $\int_0^1 z^{n-1}\psi(z)dz = 0$ with $\psi(z) := \phi(-\log(z))$. (In particular, $\psi \in L_1$.) Applying Weierstrass's Approximation Theorem, $\int_0^1 f(z)\psi(z)dz = 0$ for all continuous functions f . Arguing as in Thm. 10.10, $\int_0^1 f(z)\psi(z)dz = 0$ for all bounded measurable functions f . In particular, letting $m > 0$ and $f(z) = \psi(z)1_{|\psi| \leq m}$ gives $\int_0^1 \psi(z)^2 1_{|\psi| \leq m}(z)dz = 0$. Letting $m \rightarrow \infty$ by the Monotone Convergence Theorem, $\int_0^1 \psi(z)^2 dz = 0$, i.e. $\psi = 0$, i.e. $\phi = 0$ as desired.

Proof of Claim 2: Let ξ be a linear functional. Then $\int_0^\infty e^{-\lambda t}\xi(u_t)dt = 0$, so $\xi(u_t) = 0$ for almost all t by Claim 1. Since $\xi(u_t)$ is continuous, $\xi(u_t) = 0$ for all t . By Hausdorffness of the weak topology, $u_t = 0$.

We now prove our uniqueness statement. Suppose two semigroups T_t, T'_t have the same infinitesimal generator A . We will show that $T_t f = T'_t f$ for all $f \in H_0$ (with H_0 identical for both T_t and T'_t). First, recall from Thm. 10.39(3) that H_0 is the strong closure of the domain of the infinitesimal generator. Therefore, H_0 is in fact the same for both T_t and T'_t . By Thm. 10.39(6), R_λ is identical for T_t, T'_t . As a result, $\forall f \in H_0 \forall \lambda > 0$, $\int_0^\infty e^{-\lambda t}T_t f dt = \int_0^\infty e^{-\lambda t}T'_t f dt$. By Claim 2, $T_t f - T'_t f = 0 \forall f \in H_0 \forall \lambda > 0$, giving our desired uniqueness.

Let $f \in \mathcal{D}_A$. We now show that $u_t := T_t f$ is the unique solution of the equation $du_t/dt = Au_t$ such that (i) u_t is strongly differentiable and its derivative is strongly continuous for $t \geq 0$, (ii) $\|u_t\|$ is bounded, and (iii) $u_0 = \lim_{t \downarrow 0} u_t = f$.

Thm. 10.39(4) shows that u_t satisfies $du_t/dt = Au_t$ and (i) holds (using also $Af \in H_0$ by Thm. 10.39(2)). Then (ii) follows from (i) since $\|T_t f\| \leq \|f\|$. Also, using Thm. 10.39(3), $f \in \mathcal{D}_A \subseteq H_0$, so (iii) holds. We now prove uniqueness. Suppose u_t is a solution of du_t/dt satisfying (i),(ii) and such that $u_0 = \lim_{t \downarrow 0} u_t = 0$. We will show that $u_t = 0$. Define

$v_t := e^{-\lambda t} u_t$. By definition of u_t , $dv_t/dt = -\lambda e^{-\lambda t} u_t + e^{-\lambda t} du_t/dt = -\lambda v_t + A v_t$. So, by Thm. 10.39(6), $v_t = -R_\lambda(dv_t/dt)$.

Now, dv_t/dt is strongly continuous, and R_λ is bounded by Thm. 10.39(6), so $\int_0^t v_s ds = -R_\lambda \int_0^t (dv_s/ds) ds = -R_\lambda(v_t - v_0) = -R_\lambda v_t$. By (ii), $\|v_t\| \rightarrow 0$ as $t \rightarrow \infty$, so $\|R_\lambda v_t\| \leq \|R_\lambda\| \|v_t\| \rightarrow 0$ as $t \rightarrow \infty$. So, for all $\lambda > 0$, $\int_0^\infty v_s ds = \int_0^\infty e^{-\lambda s} u_s ds = 0$. Applying Claim 2 shows $u_t = 0$ for all $t > 0$, as desired. \square

11. APPENDIX: ISOPERIMETRIC INEQUALITIES

Theorem 11.1. (Prékopa-Leindler Inequality) *Let $f, g, m: \mathbb{R}^n \rightarrow [0, \infty)$ be integrable. Suppose $\forall x, y \in \mathbb{R}^n$ $m(\lambda x + (1 - \lambda)y) \geq f(x)^\lambda g(y)^{1-\lambda}$. Then*

$$\int_{\mathbb{R}^n} m(x) dx \geq \left(\int_{\mathbb{R}^n} f(x) dx \right)^\lambda \left(\int_{\mathbb{R}^n} g(x) dx \right)^{1-\lambda}$$

Proof. We induct on n . Let μ_n denote Lebesgue measure. Note that $\mu_1(A + B) \geq \mu_1(A) + \mu_1(B)$ for $A, B \subseteq \mathbb{R}$. (To prove this, by approximation we may assume compactness, and then translation with $\max(A) = 0 = \min(B)$ shows $A + B \supseteq A \cup B$, etc.) By scaling, we may assume $\|f\|_\infty = \|g\|_\infty = 1$. Let $A := \{x: f(x) \geq t\}$, $B := \{x: g(x) \geq t\}$, $C := \{x: m(x) \geq t\}$. Let $x \in A, y \in B$. Then $m(\lambda x + (1 - \lambda)y) \geq f(x)^\lambda g(y)^{1-\lambda} \geq t$. Thus, $\lambda A + (1 - \lambda)B \subseteq C$. Applying our claim, we have $\mu_1(C) \geq \mu_1(\lambda A + (1 - \lambda)B) \geq \lambda \mu_1(A) + (1 - \lambda) \mu_1(B)$. So, applying the definitions of A, B, C , integrating with respect to t , and applying AMGM gives

$$\int_{\mathbb{R}} m(x) dx \geq \lambda \int_{\mathbb{R}} f(x) dx + (1 - \lambda) \int_{\mathbb{R}} g(x) dx \geq \left(\int_{\mathbb{R}} f(x) dx \right)^\lambda \left(\int_{\mathbb{R}} g(x) dx \right)^{1-\lambda}$$

Now, let $(t, s) \in \mathbb{R} \times \mathbb{R}^{n-1}$. We can check that $f_{t_0}(s) = f(t_0, s)$, $g_{t_1}(s) = g(t_1, s)$, $m_t(s) = m(t, s)$ satisfy the hypotheses of the theorem, for $t = \lambda t_0 + (1 - \lambda)t_1$. So, we can apply the inductive hypothesis (for the case $n - 1$). But the result is the $n = 1$ hypotheses, viewed as functions of the t 's. Thus applying the $n = 1$ case of the theorem gives the desired result. \square

Theorem 11.2. (Brunn-Minkowski Inequality) *Let $A, B \subseteq \mathbb{R}^n$ measurable, nonempty.*

- (1) $\forall \lambda \in [0, 1]$ $\text{vol}(\lambda A + (1 - \lambda)B) \geq \text{vol}(A)^\lambda \text{vol}(B)^{1-\lambda}$
- (2) $\text{vol}(A + B)^{1/n} \geq \text{vol}(A)^{1/n} + \text{vol}(B)^{1/n}$

Proof. To prove (1), let $m = 1_{(\lambda A + (1-\lambda)B)}$, $f = 1_A$, $g = 1_B$, and apply Prékopa-Leindler (Thm. 11.1). To prove (2), let $\tilde{A} = \frac{A}{\text{vol}(A)^{1/n}}$, $\tilde{B} = \frac{B}{\text{vol}(B)^{1/n}}$, $\lambda = \frac{\text{vol}(A)^{1/n}}{\text{vol}(A)^{1/n} + \text{vol}(B)^{1/n}}$, and apply (1). \square

Theorem 11.3. (The Isoperimetric Inequality) *Let $A \subseteq \mathbb{R}^n$ as above. Let $B = rB_2^n$ be the Euclidean ball, scaled so that $\text{vol}(A) = \text{vol}(B) = r^n \text{vol}(B_2^n)$. Then $\text{vol}_{n-1}(\partial B) \leq \text{vol}_{n-1}(\partial A)$.*

Proof. We show that $\text{vol}(A_\varepsilon) \geq \text{vol}(B_\varepsilon)$. Observe that $r = \left(\frac{\text{vol}(A)}{\text{vol}(B_2^n)}\right)^{\frac{1}{n}}$. So,

$$\begin{aligned} \text{vol}(A_\varepsilon)^{\frac{1}{n}} &= \text{vol}(A + \varepsilon B_2^n)^{\frac{1}{n}} \quad , \text{ definition of } A_\varepsilon \\ &\geq \text{vol}(A)^{\frac{1}{n}} + \text{vol}(\varepsilon B_2^n)^{\frac{1}{n}} \quad , \text{ Brunn-Minkowski, Thm. 11.2} \\ &= r \text{vol}(B_2^n)^{\frac{1}{n}} + \varepsilon \text{vol}(B_2^n)^{\frac{1}{n}} \quad , \text{ definition of } r \\ &= \text{vol}((r + \varepsilon)B_2^n)^{\frac{1}{n}} \\ &= \text{vol}(B_\varepsilon)^{\frac{1}{n}} \end{aligned}$$

as desired. □

Theorem 11.4 (Spherical Isoperimetric Inequality). *Among all domains of fixed volume on the sphere, one with minimal boundary volume is the geodesic ball.*

Proof due to Figiel-Lindenstrauss-Milman-1977. Idea: if we start with an optimal set that is not a geodesic ball, we can apply a finite number of symmetrizations to it so that its interior is squished into a smaller region. This gives a contradiction, so we had a ball at the beginning. The technical device of outer radius allows the argument to proceed rigorously.

Let $S^{n-1} \subseteq \mathbb{R}^n$ be the unit sphere centered at the origin, and let $A \subseteq S^{n-1}$ be closed. Given two antipodal points $a, b \in S^{n-1}$, let $\gamma \subseteq S^{n-1}$ be a geodesic joining a and b . We define the symmetrization $\sigma_\gamma(A)$ as follows. For each $y \in \gamma$, let Π_y be the plane containing y , such that Π_y is perpendicular to the line in \mathbb{R}^n connecting a and b . Note that $\Pi_y \cap S^{n-1}$ is a dilation and translation of S^{n-2} , so let $\mu_{n-2,y}$ be the normalized Haar measure on $\Pi_y \cap S^{n-1}$. We let $\sigma_\gamma(A) \cap \Pi_y$ be a geodesic ball in S^{n-2} with center y , such that $\mu_{n-2,y}(\sigma_\gamma(A) \cap \Pi_y) = \mu_{n-2,y}(A \cap \Pi_y)$. From Fubini's Theorem, we have $\mu_{n-1}(B) = \mu_{n-1}(A)$, where μ_{n-1} is normalized Haar measure on S^{n-1} .

We say $\sigma_\gamma(A)$ is the **symmetrization** of A with respect to γ . Let $r(A) := \min\{r > 0: \exists x \in S^{n-1}, A \subseteq \overline{B(x, r)}\}$ be the (outer) **radius** of A . Here $B(x, r) := \{y \in S^{n-1}: d(x, y) < r\}$ is the open ball of radius r centered at y on S^{n-1} , and d is the usual metric on S^{n-1} , so that $d(x, y) = \cos^{-1}(\langle x, y \rangle)$ for all $x, y \in S^{n-1}$. The minimum in the definition of $r(A)$ exists by closedness of A and $\overline{B(x, r)}$. We claim that $\sigma_\gamma(A)$ is closed.

To show this, we use the Hausdorff distance on closed sets in S^{n-1} , $\delta(A, B) := \min\{r > 0: A_r \supseteq B, B_r \supseteq A\}$. (Here $A_r := \{x \in S^{n-1}: d(x, A) < r\}$.) Let $B := \sigma_\gamma(A)$. Recall that the set of closed sets is a complete metric space with respect to the metric δ . Note that the function $y \mapsto \mu_{n-2,y}(A) = \mu_{n-2,y}(B)$ is upper semicontinuous in $y \in \gamma$, i.e. $\mu_{n-2,y_0}(A) \geq \limsup_{y \rightarrow y_0} \mu_{n-2,y}(A)$ when $y, y_0 \in \gamma$. This follows by the definition of the product topology and by the closedness of A . Now, writing S^{n-1} as $S^{n-2} \times [-1, 1] / \sim$ where $(x, 1) \sim (x', 1)$ and $(x, -1) \sim (x', -1) \forall x, x' \in S^{n-2}$, we can treat $A \subseteq S^{n-1}$ as a closed set in the product topology of $S^{n-2} \times [-1, 1]$. Given $(x, y) \in B^c \subseteq S^{n-1} \times [-1, 1]$, we wish to find a box $F \times G \subseteq S^{n-2} \times [-1, 1]$ with F, G open, so that $(x, y) \in F \times G$ and $F \times G$ is disjoint from B . Since $B \cap \Pi_y$ is a geodesic ball (which is not all of S^{n-2}), we can find $F \times G$ as required, by the upper semicontinuity of $y \mapsto \mu_{n-2,y}(B)$. (Specifically, our inability to find such a box $F \times G$ would violate this upper semicontinuity.)

Below we also use that $\mu_{n-1}(\cdot)$ is upper semi-continuous with respect to δ , that is if $A^{(1)}, A^{(2)}, \dots \subseteq S^{n-1}$ satisfy $\lim_{k \rightarrow \infty} \delta(A^{(k)}, A) = 0$, then $\mu_{n-1}(A) \geq \limsup_{k \rightarrow \infty} \mu_{n-1}(A^{(k)})$.

To see this, let $x_k \in A^{(k)}$ for any $k \geq 1$. Since $d(x_k, A) \leq \delta(A^{(k)}, A) \rightarrow 0$ as $k \rightarrow \infty$, any limit point of the set $\{x_k\}_{k=1}^\infty$ must be contained in A . Therefore, for any fixed $\varepsilon > 0$, there exists $K > 0$ such that $k \geq K$ implies $A^{(k)} \subseteq A_\varepsilon$. Let $\lambda > \mu_{n-1}(A)$. Since μ_{n-1} is a Borel measure, there exists an open set U such that $A \subseteq U$ and $\mu_{n-1}(U) < \lambda$. Since A is compact, $d(A, U^c) > 0$, and there exists $\varepsilon > 0$ such that $A_\varepsilon \subseteq U$. Combining these observations, $\limsup_{k \rightarrow \infty} \mu_{n-1}(A^{(k)}) \leq \mu_{n-1}(A_\varepsilon) \leq \mu_{n-1}(U) < \lambda$. Therefore, $\limsup_{k \rightarrow \infty} \mu_{n-1}(A^{(k)}) \leq \mu_{n-1}(A)$, as desired.

We are now ready to proceed by inducting on n . For the case $n = 1$, the theorem is clear. We require the following claims, which are proven by induction.

Claim 1: Let $A \subseteq S^{n-1}$ be closed, and define

$$M(A) := \{C \subseteq S^{n-1} : C \text{ is closed,} \\ \mu_{n-1}(C) = \mu_{n-1}(A), \mu_{n-1}(C_\varepsilon) \leq \mu_{n-1}(A_\varepsilon) \forall \varepsilon > 0\}$$

Then there is a $B \in M(A)$ with minimal radius, i.e. $\min\{r(C) : C \in M(A)\}$ exists.

Claim 2: Let $A \subseteq S^{n-1}$ be closed. Then for every half circle γ , $\sigma_\gamma(A) \in M(A)$.

Claim 3: Let $B \subseteq S^{n-1}$ be a closed set that is not a geodesic ball. There exists a finite family of half circles $\{\gamma_i\}_{i=1}^n \subseteq S^{n-1}$ so that $r(\sigma_{\gamma_n}(\sigma_{\gamma_{n-1}}(\cdots \sigma_{\gamma_1}(B)\cdots))) < r(B)$.

We prove the theorem assuming these claims. By definition of $M(A)$, $B \in M(A)$ and $C \in M(B)$ implies $C \in M(A)$. So, using Claim 2, $B \in M(A)$ and $\sigma_{\gamma_1}(B) \in M(B)$ implies $\sigma_{\gamma_1}(B) \in M(A)$, $\sigma_{\gamma_2}(\sigma_{\gamma_1}(B)) \in M(A)$, etc. Using Claim 3, we therefore see that an element of minimal (outer) radius in $M(A)$ must be a geodesic ball. Claim 1 says that this minimal element must exist, so $M(A)$ must contain a geodesic ball. The theorem is therefore proven. We now prove the claims.

Proof of Claim 1: $B \mapsto r(B)$ is continuous (with respect to the Hausdorff metric for $B \subseteq S^{n-1}$), so it suffices to show that $M(A)$ is a closed subset in the space of closed subsets of S^{n-1} (since the latter space is compact with respect to δ). Let $B^{(1)}, B^{(2)}, \dots \in M(A)$ with $\lim_{k \rightarrow \infty} \delta(B^{(k)}, B) = 0$ for some $B \subseteq S^{n-1}$, and let $\varepsilon \geq 0$. We will show $B \in M(A)$. For any fixed $\eta > 0$, there exists $K > 0$ such that, if $k \geq K$, then $B \subseteq B_\eta^{(k)}$, so $B_\varepsilon \subseteq B_{\varepsilon+\eta}^{(k)}$. So, for all $k \geq K$, $\mu_{n-1}(B_\varepsilon) \leq \mu_{n-1}(B_{\varepsilon+\eta}^{(k)}) \leq \mu_{n-1}(A_{\varepsilon+\eta})$, since $B^{(1)}, B^{(2)}, \dots \in M(A)$. Therefore,

$$\mu_{n-1}(B_\varepsilon) \leq \inf_{\eta > 0} \mu_{n-1}(A_{\varepsilon+\eta}) = \mu_{n-1}(\bigcap_{\eta > 0} A_{\varepsilon+\eta}) = \mu_{n-1}(A_\varepsilon)$$

So, letting $\varepsilon = 0$, we get $\mu_{n-1}(B) \leq \mu_{n-1}(A)$. Moreover, $\mu_{n-1}(B) \geq \limsup_{k \rightarrow \infty} \mu_{n-1}(B^{(k)}) = \mu_{n-1}(A)$, using the upper semicontinuity of $\mu_{n-1}(\cdot)$ mentioned above, and the definition of $B^{(1)}, B^{(2)}, \dots \in M(A)$. So $B \in M(A)$, as desired.

Proof of Claim 2: Let $A \subseteq S^{n-1}$ be closed and let γ be a half circle on S^{n-1} joining $z \in S^{n-1}$ with $-z$. Let u be the midpoint of γ . As usual, identify $S^{n-2,u} := S^{n-1} \cap \Pi_u$ with S^{n-2} . For any $y \in \gamma, y \neq \pm z$, define a map $\tau_y : S^{n-2,y} \rightarrow S^{n-2,u}$ by letting $\tau_y(x) := \gamma \cap S^{n-2,u}$ for any $x \in S^{n-2,y}$. (Note that this intersection is a single point). By applying polar coordinates, we see that there exists a function f such that, if $y_1, y_2 \in \gamma$ and if $x_1 \in S^{n-2,y_1}, x_2 \in S^{n-2,y_2}$, we have

$$d(x_1, x_2) = f(y_1, y_2, d(\tau_{y_1}(x_1), \tau_{y_2}(x_2))).$$

Moreover, for y_1, y_2 fixed, f is monotonically increasing with respect to its third argument, $d(\tau_{y_1}(x_1), \tau_{y_2}(x_2)) \leq \pi$.

For every $y_1, y_2 \in \gamma, \varepsilon > 0$ (with $d(y_1, y_2) < \varepsilon$) there is an $\eta(y_1, y_2, \varepsilon)$ so that, for every $C \subseteq S^{n-2, y_1}$, we have

$$C_\varepsilon \cap S^{n-2, y_2} = \tau_{y_2}^{-1}((\tau_{y_1} C)_{\eta(y_1, y_2, \varepsilon)}) \quad (*)$$

To see this, it suffices to consider the case that $C = \{x_1\}$. Then

$$\begin{aligned} C_\varepsilon \cap S^{n-2, y_2} &= \{x_2 \in S^{n-2, y_2} : d(x_1, x_2) < \varepsilon\} \\ &= \{x_2 \in S^{n-2, y_2} : f(y_1, y_2, d(\tau_{y_1}(x_1), \tau_{y_2}(x_2))) < \varepsilon\} \\ &= \{x_2 \in S^{n-2, y_2} : d(\tau_{y_1}(x_1), \tau_{y_2}(x_2)) < \eta\} \end{aligned}$$

Here η is determined by the existence and monotonicity of f . (If $d(y_1, y_2) \geq \varepsilon$, then $C_\varepsilon \cap S^{n-2, y_2} = \emptyset$.) Note that the subscript ε on the left of $(*)$ denotes an ε neighborhood in S^{n-1} , whereas the subscript η on the right of $(*)$ denotes an η neighborhood in S^{n-2} . Let $A^y := A \cap S^{n-2, y}$. By fixing $y_2 = y$ and varying $y_1 = z$ in $(*)$, we have

$$\tau_y((A_\varepsilon)^y) = \cup_{\{z \in \gamma : d(z, y) < \varepsilon\}} (\tau_z(A^z))_{\eta(z, y, \varepsilon)} \quad (**)$$

Substituting $B := \sigma_\gamma(A)$ gives

$$\tau_y((B_\varepsilon)^y) = \cup_{\{z \in \gamma : d(z, y) < \varepsilon\}} (\tau_z(B^z))_{\eta(z, y, \varepsilon)} \quad (\dagger)$$

By definition of B , $\tau_z(B^z)$ is a geodesic ball in $S^{n-2, u} \forall z \in \gamma$, and $\mu_{n-2, u}(\tau_z(B^z)) = \mu_{n-2, u}(\tau_z(A^z))$. So, the induction hypothesis (i.e. the full theorem) says

$$\mu_{n-2, u}((\tau_z(B^z))_{\eta(z, y, \varepsilon)}) \leq \mu_{n-2, u}((\tau_z(A^z))_{\eta(z, y, \varepsilon)}) \quad (\ddagger)$$

for admissible y, z, ε . Since the sets on the right side of (\ddagger) are all $(n-2)$ -dimensional geodesic balls with the same center, we have

$$\begin{aligned} \mu_{n-2, u}(\tau_y((B_\varepsilon)^y)) &= \sup_{z \in \gamma : d(z, y) \leq \varepsilon} \mu_{n-2, u}((\tau_z(B^z))_{\eta(z, y, \varepsilon)}) \\ &\leq \sup_{z \in \gamma : d(z, y) \leq \varepsilon} \mu_{n-2, u}((\tau_z(A^z))_{\eta(z, y, \varepsilon)}) \quad , \text{ from } (\ddagger) \\ &\leq \mu_{n-2, u}(\tau_y((A_\varepsilon)^y)) \quad , \text{ from } (**). \end{aligned}$$

Re-writing this inequality, we see that for every $y \in \gamma, y \neq \pm x$ we have

$$\mu_{n-2, y}((B_\varepsilon)^y) \leq \mu_{n-2, y}((A_\varepsilon)^y)$$

So by Fubini's Theorem, we can integrate this inequality to get $\mu_{n-1}(B_\varepsilon) \leq \mu_{n-1}(A_\varepsilon)$, so that $B \in M(A)$ as desired.

Proof of Claim 3: Let $B \subseteq S^{n-1}$ be closed, and suppose B is not a geodesic ball. Let $r = r(B)$ as above, and let $u \in S^{n-1}$ be such that $B \subseteq \overline{B(u, r)}$. Let γ be a half circle with midpoint u , so that we will symmetrize with respect to γ , leaving $\overline{B(u, r)}$ fixed. Since B is not a geodesic ball, $E := B^c \cap \partial B(u, r) \neq \emptyset$.

We need two observations. First, any symmetrization σ_γ does not decrease the set E . That is, $E \subseteq (\sigma_\gamma(B))^c \cap \partial B(u, r)$. Second, we can find symmetrizations that increase E . To see the second claim, let $G \subseteq \partial B(u, r)$ be a relatively open set. Given any $x \in \partial B(u, r) \setminus G$, there exists a relatively open set $G_x \subseteq \partial B(u, r)$ and γ_x such that $x \in G_x$, and $G_x \cap \sigma_{\gamma_x}(B) = \emptyset$. To construct γ_x , consider the straight line ℓ (in \mathbb{R}^n) between x and some point $y \in B^c \cap \partial B(u, r)$ (which exists since B is not a ball). Let P reflect $\partial B(u, r)$ across a hyperplane perpendicular to ℓ and intersecting ℓ at its midpoint. Then, let G_x be a small ball (in $\partial B(u, r)$) around x disjoint from G , such that $PG_x \subseteq B^c \cap \partial B(u, r)$ (which is possible since B is closed).

Observe that G_x does what we claimed above. Also note that G_x, γ_x depend on x and G , but not on B .

Now, apply the above observations to B and $G := B^c \cap \partial B(u, r)$ to produce γ_1, G_{x_1} . Then, apply these same observations to $\sigma_{\gamma_1}(B)$ and $G := \sigma_{\gamma_1}(B)^c \cap \partial B(u, r)$ to produce γ_2 and G_{x_2} , and so on. By compactness of S^{n-1} (using a cover by $\{G_{x_i}\}_{i \geq 1}$), after a finite number of symmetrizations we have $\sigma_{\gamma_n}(\cdots \sigma_{\gamma_1}(B) \cdots)$ disjoint from $\partial B(u, r)$. Therefore, $r(\sigma_{\gamma_n}(\cdots \sigma_{\gamma_1}(B) \cdots)) < r(B)$. \square

As an application, we prove the following concentration of measure result. Note that the exponential dependence on n implies that almost all of a high dimensional sphere is close to any given set of Haar measure $1/2$. Put another way, a high dimensional sphere has a ‘‘large waist.’’

Theorem 11.5 (Concentration of measure on the sphere). *Let μ be the normalized Haar measure on S^{n+1} . Let $A \subseteq S^{n+1}$, let $\varepsilon > 0$, and define $A_\varepsilon := \{x \in S^{n+1} : \exists y \in S^{n+1} \text{ with } d_{S^{n+1}}(x, y) \leq \varepsilon\}$. If $\mu(A) \geq 1/2$ then $\mu(A_\varepsilon) \geq 1 - \sqrt{\frac{\pi}{8}} e^{-\varepsilon^2 n/2}$.*

Proof. By Theorem 11.4, it suffices to prove this claim for geodesic balls, i.e. it suffices to analyze the quantity

$$\mu(B(\pi/2 + \varepsilon)) = \frac{\int_{-\pi/2}^{\varepsilon} \cos^n(t) dt}{\int_{-\pi/2}^{\pi/2} \cos^n(t) dt}.$$

For any $n \geq 1$, let $I_n := \int_0^{\pi/2} \cos^n(t) dt$. Changing variables and using $\cos(t) \leq e^{-t^2/2}$, valid for any $0 \leq t \leq \pi/2$ (which follows since $f(t) := \log \cos t$ satisfies $f''(t) = -1/\cos^2(t) \leq -1$ for all $0 \leq t \leq \pi/2$),

$$\begin{aligned} 1 - \mu(B(\pi/2 + \varepsilon)) &= \int_{\varepsilon}^{\pi/2} \cos^n(t) \frac{dt}{2I_n} = \frac{1}{\sqrt{n}} \int_{\varepsilon\sqrt{n}}^{(\pi/2)\sqrt{n}} \cos^n(t/\sqrt{n}) \frac{dt}{2I_n} \\ &\leq \frac{1}{\sqrt{n}} \int_{\varepsilon\sqrt{n}}^{(\pi/2)\sqrt{n}} e^{-t^2/2} \frac{dt}{2I_n} \leq \frac{1}{\sqrt{n}} e^{-\varepsilon^2 n/2} \int_0^{(\pi/2-\varepsilon)\sqrt{n}} e^{-t^2/2} \frac{dt}{2I_n} \\ &\leq \frac{1}{\sqrt{n}} e^{-\varepsilon^2 n/2} \int_0^{\infty} e^{-t^2/2} \frac{dt}{2I_n} = \frac{1}{2\sqrt{n}I_n} e^{-\varepsilon^2 n/2} \sqrt{\pi/2}. \end{aligned}$$

Integration by parts shows that $I_n = \frac{n-1}{n} I_{n-2}$. Since $(n-1)/\sqrt{n(n-2)} \geq 1$ for any $n \geq 3$, we get $\sqrt{n}I_n \geq \sqrt{n-2}I_{n-2}$ for any $n \geq 3$, so that

$$\sqrt{n}I_n \geq \min(I_1, \sqrt{2}I_2) = \min(1, \sqrt{2}\pi/4) = 1, \quad \forall n \geq 1$$

In summary, $1 - \mu(B(\pi/2 + \varepsilon)) \leq e^{-\varepsilon^2 n/2} \sqrt{\pi/8}$. \square

Theorem 11.5 implies a corresponding statement for Lipschitz functions. That is, Lipschitz functions on high-dimensional spheres are typically close to their average value.

For any $x = (x_1, \dots, x_n) \in \mathbb{R}^n$, we denote $\|x\| := (x_1^2 + \cdots + x_n^2)^{1/2}$.

Theorem 11.6 (Concentration of measure, Lipschitz function form). *Let $f: S^{n+1} \rightarrow \mathbb{R}$. Suppose that for all $x, y \in S^{n+1}$, $|f(x) - f(y)| \leq \|x - y\|$, so that f is 1-Lipschitz. Let μ denote normalized Haar measure on S^{n+1} . Then for all $\varepsilon > 0$,*

$$\mu \left(x \in S^{n+1} : \left| f(x) - \int_{S^{n+1}} f(y) d\mu(y) \right| \geq \varepsilon \right) \leq \sqrt{\frac{\pi}{2}} e^{-n\varepsilon^2/4}.$$

Proof. Let $m \in \mathbb{R}$ such that $\mu(x \in S^{n+1}: f(x) \leq m) \geq 1/2$ and $\mu(x \in S^{n+1}: f(x) \geq m) \geq 1/2$. Let $C := \{x \in S^{n+1}: f(x) \leq m\}$. Then $x \in C_\varepsilon$ if and only if $\exists y \in C$ with $\|x - y\|_2 \leq \varepsilon$. Since f is 1-Lipschitz, $|f(x) - f(y)| \leq \varepsilon$, so that $f(x) \leq m + \varepsilon$. Taking the contrapositive,

$$\{x \in S^{n+1}: f(x) > m + \varepsilon\} \subseteq S^{n+1} \setminus A_\varepsilon$$

So, from Thm. 11.5, since $\mu(C) \geq 1/2$, we have

$$\mu(x \in S^{n+1}: f(x) > m + \varepsilon) \leq \sqrt{\pi/8} e^{-n\varepsilon^2/2}.$$

Similarly, $\mu(x \in S^{n+1}: f(x) < m - \varepsilon) \leq \sqrt{\pi/8} e^{-n\varepsilon^2/2}$. In conclusion,

$$\mu(x \in S^{n+1}: |f(x) - m| > \varepsilon) \leq 2\sqrt{\pi/8} e^{-n\varepsilon^2/2}. \quad (*)$$

It remains to replace m with $\int_{S^{n+1}} f(y) d\mu(y)$. Consider $\mu \times \mu$ on $S^{n+1} \times S^{n+1}$. Observe

$$\begin{aligned} & (\mu \times \mu) ((x, y) \in S^{n+1} \times S^{n+1}: |f(x) - f(y)| \geq \varepsilon) \\ & \leq (\mu \times \mu) (\{|f(x) - m| \geq \varepsilon/2\} \cup \{|f(y) - m| \geq \varepsilon/2\}) \\ & \leq 2\mu(|f(x) - m| \geq \varepsilon/2) \leq 4\sqrt{\pi/8} e^{-n\varepsilon^2/2}, \text{ from } (*) \end{aligned}$$

Let $\lambda > 0$. Then from Theorem 11.5, if $\lambda^2 := n/4$,

$$\begin{aligned} & \int_{S^{n+1} \times S^{n+1}} e^{\lambda^2(f(x)-f(y))^2} d\mu(x) d\mu(y) \\ & = \int_0^\infty 2\lambda^2 t e^{\lambda^2 t^2} (\mu \times \mu) ((x, y) \in S^{n+1} \times S^{n+1}: |f(x) - f(y)| \geq t) dt \\ & \leq 4\sqrt{\pi/8} \int_0^\infty \lambda^2 t e^{\lambda^2 t^2} e^{-nt^2/2} dt = \sqrt{\pi/8} \int_0^\infty t n e^{-nt^2/4} dt = 2\sqrt{\pi/8} = \sqrt{\pi/2}. \end{aligned}$$

So, for this λ , Jensen's inequality in y implies that

$$\sqrt{\pi/2} \geq \int_{S^{n+1} \times S^{n+1}} e^{\lambda^2(f(x)-f(y))^2} d\mu(x) d\mu(y) \geq \int_{S^{n+1}} e^{\lambda^2(f(x)-\int_{S^{n+1}} f(y) d\mu(y))^2} d\mu(x).$$

Finally, by Chebyshev's inequality,

$$\begin{aligned} \mu(x \in S^{n+1}: |f(x) - \int_{S^{n+1}} f d\mu| \geq \varepsilon) & = \mu(x \in S^{n+1}: e^{\lambda^2|f(x)-\int_{S^{n+1}} f(y) d\mu(y)|^2} \geq e^{\lambda^2\varepsilon^2}) \\ & \leq e^{-\lambda^2\varepsilon^2} \int_{S^{n+1}} e^{\lambda^2|f(x)-\int_{S^{n+1}} f(y) d\mu(y)|^2} d\mu(x) \leq \sqrt{\pi/2} e^{-\lambda^2\varepsilon^2}. \end{aligned}$$

□

Theorem 11.7. (Spheres and Gaussians, Strong Version) Let γ_n be the standard Gaussian measure on \mathbb{R}^n , and let $A \subseteq \mathbb{R}^n$ be a Borel set. Let $\sigma_{\sqrt{N}}^N$ denote the normalized Haar measure on $\sqrt{N} \cdot S^N$. Let $P_{N+1,n}$ (for $N \geq n$) be the standard linear projection from \mathbb{R}^{N+1} onto \mathbb{R}^n , i.e. $P_{N+1,n}(x_1, \dots, x_{N+1}) = (x_1, \dots, x_n)$. Then

$$\lim_{N \rightarrow \infty} \sigma_{\sqrt{N}}^N(P_{N+1,n}^{-1}(A) \cap (\sqrt{N} \cdot S^N)) = \gamma_n(A)$$

Proof. Let $\{g_i\}$ be iid standard Gaussian random variables, and define $R_N^2 := g_1^2 + \dots + g_N^2$. By rotation invariance, $\frac{\sqrt{N}}{R_{N+1}}(g_1, \dots, g_{N+1})$ has the same distribution as $\sigma_{\sqrt{N}}^N$. So, by projecting both sides, $\frac{\sqrt{N}}{R_{N+1}}(g_1, \dots, g_n)$ has the same distribution as $P_{N+1,n}(\sigma_{\sqrt{N}}^N) = \sigma_{\sqrt{N}}^N(P_{N+1,n}^{-1}(\cdot))$ (for $N \geq n$). In the previous sentence, we calculated both mentioned measures as pushforward measures, under the map $P_{N+1,n}: \mathbb{R}^{N+1} \rightarrow \mathbb{R}^n$.

$R_N^2/N \rightarrow 1$ a.s. by the Strong Law of Large Numbers, giving a weak convergence result. However, we need to be more precise. Note that $R_n^2, R_{N+1}^2 - R_n^2$ and $(g_1, \dots, g_n)/R_n$ are independent of each other. The first two are independent by definition, and the third is independent from the others by, say, invoking polar coordinates. (The third term is the ‘‘angular’’ part of the Gaussian vector, and the first is the ‘‘radial’’ part.) Therefore, R_n^2/R_{N+1}^2 is independent of $(g_1, \dots, g_n)/R_n$.

Now, write $R_n^2/R_{N+1}^2 = 1/(1 + (g_{n+1}^2 + \dots + g_{N+1}^2)/(g_1^2 + \dots + g_n^2))$ and observe that the ratio of g_i 's is a ratio of two independent chi squared distributions. A computation then shows that R_n^2/R_{N+1}^2 has beta distribution with parameters $n/2$ and $(N + 1 - n)/2$. The corresponding distribution function is therefore

$$\frac{\Gamma\left(\frac{N+1}{2}\right)}{\Gamma\left(\frac{n}{2}\right)\Gamma\left(\frac{N-n+1}{2}\right)} x^{\frac{n}{2}-1}(1-x)^{\frac{N-n+1}{2}-1}$$

Combining our results and using the definition of the pushforward, we have

$$\begin{aligned} \sigma_{\sqrt{N}}^N(P_{N+1,n}^{-1}(A) \cap (\sqrt{N} \cdot S^N)) &= \mathbb{P}(P_{N+1,n}(\sigma_{\sqrt{N}}^N) \in A) \\ &= \mathbb{P}\left(\frac{\sqrt{N}}{R_{N+1}}(g_1, \dots, g_n) \in A\right) = \mathbb{P}\left(\left(N \frac{R_n^2}{R_{N+1}^2}\right)^{1/2} \cdot \frac{1}{R_n}(g_1, \dots, g_n) \in A\right) \\ &= \frac{\Gamma\left(\frac{N+1}{2}\right)}{\Gamma\left(\frac{n}{2}\right)\Gamma\left(\frac{N-n+1}{2}\right)} \int_{S^{n-1}} \int_0^1 1_A(\sqrt{N}tx) t^{\frac{n}{2}-1} (1-t)^{\frac{N-n+1}{2}-1} dt d\sigma^{n-1}(x) \\ &= \frac{\Gamma\left(\frac{N+1}{2}\right)}{\Gamma\left(\frac{n}{2}\right)\Gamma\left(\frac{N-n+1}{2}\right)} \frac{2}{N^{n/2}} \int_{S^{n-1}} \int_0^{\sqrt{N}} 1_A(ux) u^{n-1} \left(1 - \frac{u^2}{N}\right)^{\frac{N-n+1}{2}-1} dud\sigma^{n-1}(x) \end{aligned}$$

In the last line we used the change of variables $u = \sqrt{N}t$. Applying the Dominated Convergence Theorem and letting $N \rightarrow \infty$, the last quantity converges to

$$\frac{1}{2^{\frac{n}{2}-1}\Gamma\left(\frac{n}{2}\right)} \int_{S^{n-1}} \int_0^\infty 1_A(ux) u^{n-1} e^{-u^2/2} dud\sigma^{n-1}(x)$$

which is $\gamma_n(A)$ in polar coordinates, as desired. (We used the formula $\Gamma(n+1) = n\Gamma(n)$ to get the correct constant in front of the integral.) \square

We can now verify Gaussian isoperimetry via Thms. 11.7 and 11.4.

Theorem 11.8. (Gaussian Isoperimetric inequality) For $A \subseteq \mathbb{R}^n$ define $d\gamma_n(x) := e^{-\|x\|_2^2/2} dx / (2\pi)^{n/2}$. Define $\gamma(\partial A) := \liminf_{\varepsilon \rightarrow 0^+} (\gamma(A_\varepsilon) - \gamma(A)) / \varepsilon$. Then among all sets with the same (Gaussian) volume, one with minimal (Gaussian) boundary measure is the half space.

Proof. [Led96] Define $\Phi(t) := \frac{1}{\sqrt{2\pi}} \int_{-\infty}^t e^{-x^2/2} dx$, $t \in \mathbb{R}$. We may assume that $\gamma_n(A) = \Phi(a)$ for some $a \in \mathbb{R}$. It suffices to show that

$$\gamma_n(A_r) \geq \Phi(a + r)$$

That is, we want $\Phi^{-1}(\gamma_n(A_r)) \geq \Phi^{-1}(\gamma_n(A)) + r$. Now, let $b \in \mathbb{R}$, $b < a$. Since $\Phi(a) = \gamma_n(A)$, we have

$$\gamma_n(A) = \Phi(a) > \Phi(b) = \gamma_1((-\infty, b])$$

So, for N large enough ($N \geq n$), Thm. 11.7 gives

$$\sigma_{\sqrt{N}}^N(P_{N+1,n}^{-1}(A) \cap (\sqrt{N})S^N) > \sigma_{\sqrt{N}}^N(P_{N+1,1}^{-1}((-\infty, b]) \cap (\sqrt{N})S^N) \quad (*)$$

Note that $P_{N+1,n}^{-1}(A_r) \cap (\sqrt{N})S^N \supseteq (P_{N+1,n}^{-1}(A) \cap (\sqrt{N})S^N)_r$, where the subscript on the left denotes a neighborhood in \mathbb{R}^n , and the subscript on the right denotes a neighborhood on $(\sqrt{N})S^N$ with respect to geodesic distance. This set containment follows since the map $P_{n+1,n}^{-1}$ increases distances. Note also that $P_{N+1,1}^{-1}((-\infty, b]) \cap (\sqrt{N})S^N$ is a geodesic ball on $(\sqrt{N})S^N$. This gives us two things. First, we can apply the spherical isoperimetric inequality. Second, we have $(P_{N+1,1}^{-1}((-\infty, b]) \cap (\sqrt{N})S^N)_r = P_{N+1,1}^{-1}((-\infty, b + r(N)])$, where

$$r(N) = -\sqrt{N} \cos(\cos^{-1}(b/\sqrt{N}) + r/\sqrt{N}) + b, \quad b < 0$$

with a similar equality for $b > 0$. In either case, the cosine addition formula shows that $\lim_{N \rightarrow \infty} r(N) = r$. Combining all of these observations, we have

$$\begin{aligned} & \sigma_{\sqrt{N}}^N(P_{N+1,n}^{-1}(A_r) \cap (\sqrt{N})S^N) \\ & \geq \sigma_{\sqrt{N}}^N((P_{N+1,n}^{-1}(A) \cap (\sqrt{N})S^N)_r) \quad , \text{ by set containment} \\ & \geq \sigma_{\sqrt{N}}^N((P_{N+1,n}^{-1}((-\infty, b]) \cap (\sqrt{N})S^N)_r) \quad , \text{ by Thm. 11.4, and } (*) \\ & = \sigma_{\sqrt{N}}^N(P_{N+1,n}^{-1}((-\infty, b + r(N)]) \cap (\sqrt{N})S^N) \end{aligned}$$

So, letting $N \rightarrow \infty$ and applying Thm. 11.7, we get

$$\gamma_n(A_r) \geq \Phi(b + r)$$

And since $b < a$ is arbitrary, we get $\gamma_n(A_r) \geq \Phi(a + r)$ as desired. \square

Proof. [Led96] Instead of using the isoperimetric inequality on the sphere, we argue more directly, as in Thm. 11.3, and we use Brunn-Minkowski (for Gaussian space, Thm. 11.9). Let $\lambda \in (0, 1)$ and let $B \subseteq \mathbb{R}^n$ be the Euclidean ball with radius $r/(1-\lambda)$, so $B = (r/(1-\lambda))B_2^n$. From Gaussian Brunn-Minkowski (Thm. 11.9) we have

$$\begin{aligned} \Phi^{-1}(\gamma_n(\lambda A)_r) &= \Phi^{-1}(\gamma_n(\lambda A + rB_2^n)) = \Phi^{-1}(\gamma_n(\lambda A + (1-\lambda)B)) \\ &\geq \lambda \Phi^{-1}(\gamma_n(A)) + (1-\lambda) \Phi^{-1}(\gamma_n(B)) \\ &= \lambda \Phi^{-1}(\gamma_n(A)) + (1-\lambda) \varepsilon(\lambda, r) \Phi^{-1}(\gamma_1(-\infty, r/(1-\lambda)]) \\ &= \lambda \Phi^{-1}(\gamma_n(A)) + \varepsilon(\lambda, r)r \end{aligned}$$

where $\varepsilon(\lambda, r) \rightarrow 1$ as $\lambda \rightarrow 1$, for fixed r . So, letting $\lambda \rightarrow 1$ gives $\Phi^{-1}(\gamma_n(A_r)) \geq \Phi^{-1}(\gamma_n(A)) + r$, as desired. \square

Proof. [Bob97] We begin with a functional form of the isoperimetric inequality. We prove this via a “two-point” inequality, which is leveraged to the whole Gaussian space, via the central limit theorem. With Φ as above, and $\phi(t) := \frac{1}{\sqrt{2\pi}}e^{-t^2/2}$, define $I := \phi \circ \Phi^{-1}$. We want to prove

$$\gamma(\partial A) \geq I(\gamma(A)) \quad (*)$$

This inequality is equivalent to $\gamma(A_r) \geq \Phi(\Phi^{-1}(\gamma(A)) + r)$. To see this, differentiate the latter inequality to get (*). Conversely, if $\gamma(\partial A) \geq I(\gamma(A))$, for all A , we can conclude that $\gamma(A_r) \geq \Phi(\Phi^{-1}(\gamma(A)) + r) - \varepsilon$, for all small $\varepsilon > 0$, with ε independent of A , but with r small and dependent on A . Now, let $p \in [0, 1]$, $r \in \mathbb{R}$, and consider $R_r(p) := \Phi(\Phi^{-1}(p) + r)$. Note that $R_{r_1+r_2}(p) = R_{r_1}(R_{r_2}(p))$ and $A_{r_1+r_2} = (A_{r_1})_{r_2}$. Thus, we can improve the inequality to $\gamma(A_r) \geq \Phi(\Phi^{-1}(\gamma(A)) + r) - \varepsilon$ with r, ε both independent of A , since r small implies (via addition), the same inequality for r large. Letting $\varepsilon \rightarrow 0$ then gives our result.

Now, we will deduce (*) from the following functional inequality

$$I\left(\int f d\gamma\right) \leq \int \sqrt{I(f)^2 + |\nabla f|^2} d\gamma \quad (**)$$

where $f: \mathbb{R}^n \rightarrow [0, 1]$. Surprisingly, we can deduce (**) from the following inequality (which we will not prove)

$$I\left(\frac{a+b}{2}\right) \leq \frac{1}{2}\sqrt{(I(a))^2 + \left|\frac{a-b}{2}\right|^2} + \frac{1}{2}\sqrt{(I(b))^2 + \left|\frac{a-b}{2}\right|^2}$$

where $0 \leq a, b \leq 1$. For $f: \{-1, 1\} \rightarrow [0, 1]$, write $f(-1) = a$ and $f(1) = b$, so that we re-write this inequality as

$$I(\mathbf{E}f) \leq \mathbf{E}\sqrt{(I(f))^2 + |\nabla f|^2} \quad (\dagger)$$

Here the expected value is taken with respect to the Bernoulli measure $\mu := (1/2)\delta_{-1} + (1/2)\delta_1$, and $|\nabla f| = |f(1) - f(-1)|/2$. To get (**) from (\dagger), we apply the Central Limit Theorem.

We first show that (\dagger) tensorizes. That is, (\dagger) holds for $f: \{-1, 1\}^n \rightarrow [0, 1]$ with measure μ^n on $\{-1, 1\}^n$, and $|\nabla f(x)|^2 := \frac{1}{4} \sum_{j=1}^n |f(x) - f(s_j(x))|^2$, with $s_j(x) := (x_1, \dots, x_{j-1}, -x_j, x_{j+1}, \dots, x_n)$. Let $x \in \{-1, 1\}^n$, and now let $f: \{-1, 1\}^{n+1} \rightarrow [0, 1]$. Let $f_0(x) := f(x, -1)$, $f_1(x) := f(x, 1)$, $p_0 := \mu(\{-1\})$, $p_1 := \mu(\{1\})$, $a_0 := \int f_0 d\mu^n$, $a_1 := \int f_1 d\mu^n$. Then $\int f d\mu^{n+1} = p_0 a_0 + p_1 a_1$. Also,

$$\begin{aligned} |\nabla f(x, -1)|^2 &= |\nabla f_0(x)|^2 + (1/4) |f_0(x) - f_1(x)|^2 \\ |\nabla f(x, 1)|^2 &= |\nabla f_1(x)|^2 + (1/4) |f_0(x) - f_1(x)|^2 \\ \int \sqrt{I(f)^2 + |\nabla f|^2} d\mu^{n+1} &= p_0 \int \sqrt{I(f_0)^2 + |\nabla f_0|^2 + \frac{1}{4} |f_0 - f_1|^2} d\mu^n \\ &\quad + p_1 \int \sqrt{I(f_1)^2 + |\nabla f_1|^2 + \frac{1}{4} |f_0 - f_1|^2} d\mu^n \quad (\ddagger) \end{aligned}$$

By applying (the proof of) Jensen’s inequality twice, we have

$$\sqrt{\left(\int u\right)^2 + \left(\int v\right)^2} \leq \int \sqrt{u^2 + v^2} \leq \int \sqrt{u^2 + v^2} \quad (***)$$

Specifically, we use that $\phi(x, y) := \sqrt{x^2 + y^2}$ is convex in one argument, when the other argument is fixed. Now, let $\ell_v(u) = au + b$ be a linear function such that $\ell_v(u) \leq \phi(u, \int v)$ and $\ell_v(\int u) = \phi(\int u, \int v)$. Then

$$\phi(\int u, \int v) = \ell_v(\int u) = a \int u + b = \int \ell_v(u) \leq \int \phi(u, \int v)$$

Now, let $\ell_u(v) = av + b$ be a linear function such that $\ell_u(v) \leq \phi(u, v)$ and $\ell_u(\int v) = \phi(u, \int v)$. Then

$$\phi(u, \int v) = \ell_u(\int v) = a \int v + b = \int \ell_u(v) \leq \int \phi(u, v)$$

We now use $(***)$ twice, with $u_0 = \sqrt{I(f_0)^2 + |\nabla f_0|^2}$, $v = (f_0 - f_1)/2$, and $u_1 = \sqrt{I(f_1)^2 + |\nabla f_1|^2}$. Then (\dagger) implies

$$\begin{aligned} \int \sqrt{I(f)^2 + |\nabla f|^2} d\mu^{n+1} &\geq p_0 \sqrt{\left(\int u_0 d\mu^n\right)^2 + \left(\int v d\mu^n\right)^2} \\ &\quad + p_1 \sqrt{\left(\int u_1 d\mu^n\right)^2 + \left(\int v d\mu^n\right)^2} \end{aligned}$$

By the inductive assumption of (\dagger) , $\int u_0 d\mu^n \geq I(a_0)$, $\int u_1 d\mu^n \geq I(a_1)$. By definition of v , $\int v d\mu^n = (a_0 - a_1)/2$. Therefore,

$$\begin{aligned} \int \sqrt{I(f)^2 + |\nabla f|^2} d\mu^{n+1} &\geq p_0 \sqrt{I(a_0)^2 + (1/4)(a_0 - a_1)^2} \\ &\quad + p_1 \sqrt{I(a_1)^2 + (1/4)(a_0 - a_1)^2} \end{aligned}$$

By the $n = 1$ case of (\dagger) ,

$$\begin{aligned} p_0 \sqrt{I(a_0)^2 + (1/4)(a_0 - a_1)^2} + p_1 \sqrt{I(a_1)^2 + (1/4)(a_0 - a_1)^2} \\ \geq I(p_0 a_0 + p_1 a_1) = I\left(\int f d\mu^{n+1}\right) \end{aligned}$$

So, the inductive step is complete, and (\dagger) is proven.

We now show how to get Gaussian isoperimetry from (\dagger) . Let $f: \mathbb{R}^n \rightarrow [0, 1]$ be a smooth function. Let $x_1, \dots, x_k \in \{-1, 1\}^n$, and define $f_k(x_1, \dots, x_k) := f((x_1 + \dots + x_k)/\sqrt{k})$. By the Central Limit Theorem, as $k \rightarrow \infty$,

$$\int_{\{-1, 1\}^{nk}} f_k d\mu^{nk} \rightarrow \int_{\mathbb{R}^n} f d\gamma_n$$

Also,

$$\begin{aligned}
|\nabla f_k(x_1, \dots, x_k)|^2 &= \frac{1}{4} \sum_{i=1}^n \sum_{j=1}^k |f_k(x_1, \dots, x_k) - f_k(x_1, \dots, s_i(x_j), \dots, x_k)|^2 \\
&= \frac{1}{4} \sum_{i=1}^n \sum_{j=1}^k \left| f\left(\frac{x_1 + \dots + x_k}{\sqrt{k}}\right) - f\left(\frac{x_1 + \dots + s_i(x_j) + \dots + x_k}{\sqrt{k}}\right) \right|^2 \\
&= \|\nabla f\left(\frac{x_1 + \dots + x_k}{\sqrt{k}}\right)\|_2^2 + O(1/\sqrt{k})
\end{aligned}$$

Since f is smooth, the error $O(1/\sqrt{k})$ is uniform over all (x_1, \dots, x_k) . So, using the Central Limit Theorem again, as $k \rightarrow \infty$,

$$\int_{\{-1,1\}^{nk}} \sqrt{I(f_k)^2 + |\nabla f_k|^2} d\mu^{nk} \rightarrow \int_{\mathbb{R}^n} \sqrt{I(f)^2 + |\nabla f|^2} d\gamma_n$$

So, we have the analogue of (†) for the Gaussian measure. To finally get Gaussian isoperimetry, let f approximate 1_A for a set A , and note that $I(0) = I(1) = 0$. \square

Theorem 11.9. (Gaussian Brunn-Minkowski, [Bor03]) *As above, let $\Phi(t) := \frac{1}{\sqrt{2\pi}} \int_{-\infty}^t e^{-x^2/2} dx$, let $\gamma = \gamma_n$ be the standard Gaussian measure on \mathbb{R}^n , let $A, B \subseteq \mathbb{R}^n$ be Borel sets, and let $\lambda \in (0, 1)$. Then*

$$\Phi^{-1}(\gamma_n(\lambda A + (1 - \lambda)B)) \geq \lambda \Phi^{-1}(\gamma_n(A)) + (1 - \lambda) \Phi^{-1}(\gamma_n(B))$$

(Compare to Euclidean Brunn-Minkowski, Thm. 11.2.)

Proof. We use a method of heat flows, together with the maximum principle. Let $u = T_t f$ be a solution of the heat equation $\partial u / \partial t = \frac{1}{2} \Delta u$. Suppose u has range $[0, 1]$. Note that $d\Phi/dt = \phi := e^{-t^2/2} / \sqrt{2\pi}$. Define $U := \Phi^{-1}(u)$. Straightforward calculations show that

$$\begin{aligned}
\frac{\partial u}{\partial t} &= \phi(U) \frac{\partial U}{\partial t}, & \nabla u &= \Phi(U) \nabla U, & \Delta u &= \phi(U) (\nabla U - U |\nabla U|^2) \\
\frac{\partial U}{\partial t} &= \frac{1}{2} \Delta U - \frac{1}{2} U |\nabla U|^2 & (*) & &
\end{aligned}$$

Let $A, B \subseteq \mathbb{R}^n$ be compact subsets, and let $\varepsilon \in (0, 1)$, $\delta \in (0, \varepsilon)$. Let F be a smooth bump function adapted to an epsilon neighborhood of A . Let $f := \delta + (1 - \varepsilon)F$, so that $\delta \leq f \leq \alpha := \delta + 1 - \varepsilon < 1$. Similarly define g for B . Define

$$\kappa := \max\{\Phi(\lambda \Phi^{-1}(\alpha) + (1 - \lambda) \Phi^{-1}(\delta)), \Phi(\lambda \Phi^{-1}(\delta) + (1 - \lambda) \Phi^{-1}(\alpha))\}$$

Crucially, $\kappa \rightarrow 0$ as $\delta \rightarrow 0$ (since $\Phi^{-1}(0) = -\infty$). Now, let h such that $\kappa \leq h \leq \alpha$, such that $h = \alpha$ on $\lambda A_\varepsilon + (1 - \lambda) B_\varepsilon$ and $h = \kappa$ off an ε -neighborhood of $\lambda A_\varepsilon + (1 - \lambda) B_\varepsilon$. Observe that the following inequality holds pointwise for $x, y \in \mathbb{R}^n$.

$$\Phi^{-1}(h(\lambda x + (1 - \lambda)y)) \geq \lambda \Phi^{-1}(f(x)) + (1 - \lambda) \Phi^{-1}(g(y)) \quad (**)$$

This inequality is clear for $x \in A, y \in B$ and for $x \notin A, y \notin B$, since we reduce to $\alpha \geq \alpha$ or $\kappa \geq \delta$. For $x \in A, y \notin B$, the inequality holds by definition of κ . (And similarly for $x \notin A, y \in B$).

Now, by letting $\delta \rightarrow 0$ and then $\varepsilon \rightarrow 0$ in the following inequality, we will complete the theorem.

$$\Phi^{-1}\left(\int hd\gamma\right) \geq \lambda\Phi^{-1}\left(\int fd\gamma\right) + (1-\lambda)\Phi^{-1}\left(\int gd\gamma\right) \quad (\ddagger)$$

We prove (\ddagger) via the more general inequality

$$\Phi^{-1}(T_t h(\lambda x + (1-\lambda)y)) \geq \lambda\Phi^{-1}(T_t f(x)) + (1-\lambda)\Phi^{-1}(T_t g(y)) \quad (\clubsuit)$$

where T_t is the heat semigroup operator.

Surprisingly, the inequality (\clubsuit) is almost a symbolic exercise. Define

$$C(t, x, y) := \Phi^{-1}(T_t h(\lambda x + (1-\lambda)y)) - \lambda\Phi^{-1}(T_t f(x)) + (1-\lambda)\Phi^{-1}(T_t g(y))$$

Then $(**)$ expresses $C(0, x, y) \geq 0$ and our desired inequality (\ddagger) corresponds to $C(1, 0, 0) \geq 0$. Moreover, by construction, C vanishes as $|x|, |y| \rightarrow \infty$. Since we know $C(0, x, y) \geq 0$, in order to prove (\clubsuit) it suffices to show via a maximum principle that $C(t, x, y) \geq 0$ for all $x, y \in \mathbb{R}^n, t \geq 0$.

Let $\xi := (t, x)$, $\eta := (t, y)$ and $\zeta := (t, \lambda x + (1-\lambda)y)$. Then

$$\begin{aligned} \nabla_x C &= \lambda[(\nabla(\Phi^{-1}(T_t h)))(\zeta) - (\nabla(\Phi^{-1}(T_t f)))(\xi)] \\ \nabla_y C &= (1-\lambda)[(\nabla(\Phi^{-1}(T_t h)))(\zeta) - (\nabla(\Phi^{-1}(T_t g)))(\eta)] \\ \Delta_x C &= \lambda^2(\Delta(\Phi^{-1}(T_t h)))(\zeta) - \lambda(\Delta\Phi^{-1}(T_t f))(\xi) \\ \Delta_y C &= (1-\lambda)^2(\Delta(\Phi^{-1}(T_t h)))(\zeta) - (1-\lambda)(\Delta\Phi^{-1}(T_t g))(\eta) \end{aligned}$$

Also,

$$\sum_{1 \leq i \leq n} \frac{\partial^2 C}{\partial x_i \partial y_i} = \lambda(1-\lambda)(\Delta(\Phi^{-1}(T_t h)))(\zeta)$$

Let $\mathcal{E} := \frac{1}{2} \left[\Delta_x + 2 \sum_{1 \leq i \leq n} \frac{\partial^2}{\partial x_i \partial y_i} + \Delta_y \right]$. Then, using $(*)$,

$$\begin{aligned} \mathcal{E}C &= \frac{1}{2} [(\Delta(\Phi^{-1}(T_t h)))(\zeta) - \lambda(\Delta(\Phi^{-1}(T_t f)))(\xi) - (1-\lambda)(\Delta(\Phi^{-1}(T_t g)))(\eta)] \\ &= \frac{\partial}{\partial t}(\Phi^{-1}(T_t h))(\zeta) + \frac{1}{2}(\Phi^{-1}(T_t h))(\zeta) |(\nabla(\Phi^{-1}(T_t h)))(\zeta)|^2 \\ &\quad - \lambda \frac{\partial}{\partial t}(\Phi^{-1}(T_t f))(\xi) - \frac{\lambda}{2}(\Phi^{-1}(T_t f))(\xi) |(\nabla(\Phi^{-1}(T_t f)))(\xi)|^2 \\ &\quad - (1-\lambda) \frac{\partial}{\partial t}(\Phi^{-1}(T_t g))(\eta) - \frac{1-\lambda}{2}(\Phi^{-1}(T_t g))(\eta) |(\nabla(\Phi^{-1}(T_t g)))(\eta)|^2 \\ &=: \frac{\partial C}{\partial t} + \Psi(t, x, y) \end{aligned}$$

Now, write

$$\begin{aligned} |(\nabla(\Phi^{-1}(T_t f)))(\xi)|^2 &= |(\nabla(\Phi^{-1}(T_t h)))(\zeta)|^2 \\ &+ \sum_{1 \leq i \leq n} \left[\frac{\partial \Phi^{-1}(T_t f)}{\partial x_i}(\xi) + \frac{\partial \Phi^{-1}(T_t h)}{\partial x_i}(\zeta) \right] \left[\frac{\partial \Phi^{-1}(T_t f)}{\partial x_i}(\xi) - \frac{\partial \Phi^{-1}(T_t h)}{\partial x_i}(\zeta) \right] \end{aligned}$$

$$\begin{aligned} |(\nabla(\Phi^{-1}(T_t g)))(\eta)|^2 &= |(\nabla(\Phi^{-1}(T_t h)))(\zeta)|^2 \\ &+ \sum_{1 \leq i \leq n} \left[\frac{\partial \Phi^{-1}(T_t g)}{\partial x_i}(\eta) + \frac{\partial \Phi^{-1}(T_t h)}{\partial x_i}(\zeta) \right] \left[\frac{\partial \Phi^{-1}(T_t g)}{\partial x_i}(\eta) - \frac{\partial \Phi^{-1}(T_t h)}{\partial x_i}(\zeta) \right] \end{aligned}$$

Then by these equalities and the definition of Ψ , we can write $\Psi(t, x, y) =: \frac{1}{2} |(\nabla(\Phi^{-1}(T_t h)))(\zeta)|^2 C - b(t, x, y) \cdot \nabla_{(x,y)} C$. In summary,

$$\mathcal{E}C + b(t, x, y) \cdot \nabla_{(x,y)} C = \frac{\partial C}{\partial t} + \frac{1}{2} |(\nabla(\Phi^{-1}(T_t h)))(\zeta)|^2 C \quad (***)$$

Note that \mathcal{E} is elliptic, since it can be written as $\mathcal{E} = \frac{1}{2}(\nabla_x, \nabla_y)^T A(\nabla_x, \nabla_y)$, where A is a block matrix of the form $\begin{pmatrix} id & id \\ id & id \end{pmatrix}$, and each id is an $n \times n$ identity matrix.

We now argue by contradiction. By the definitions of f, g, h , we know that $\inf_{0 \leq t \leq T} C(t, x, y)$ is non-negative as $|x| + |y| \rightarrow \infty$. We make the contrary assumption that $C(t, x, y) < 0$ for some $(t, x, y) \in [0, T] \times \mathbb{R}^n \times \mathbb{R}^n$. Then there exists $\varepsilon > 0$ such that $\varepsilon t + C(t, x, y)$ has a strictly negative minimum in $[0, T] \times \mathbb{R}^n \times \mathbb{R}^n$. Suppose this minimum occurs at $P = (t_0, x_0, y_0)$ with $t_0 > 0$. Note that $t_0 > 0$, since $t_0 = 0$ cannot occur, since (\clubsuit) reduces to $(**)$ at $t = 0$. Then, using that P is a minimum and \mathcal{E} is an elliptic operator,

$$C(P) < 0, \quad \frac{\partial C}{\partial t}(P) + \varepsilon \leq 0, \quad \nabla_{(x,y)} C(P) = 0, \quad \mathcal{E}C(P) \geq 0$$

However, these inequalities contradict $(***)$. We conclude that $C(t, x, y) \geq 0$, so (\clubsuit) holds, as desired. \square

We begin with Gross' Log-Sobolev inequality. Let μ be a probability measure on some measure space Ω . Given a non-negative measurable function f on Ω with $\int f \log(1+f) < \infty$, define the **entropy** of f as

$$\text{Ent}_\mu(f) := \int f \log f d\mu - \left(\int f d\mu \right) \cdot \log \left(\int f d\mu \right) = \int f \log \left(\frac{f}{\int f d\mu} \right) d\mu$$

Note that $\text{Ent}_\mu(f) \geq 0$ by Jensen's inequality (since $g(x) = x \log x$, $x \geq 0$ is convex). Also, a short calculation shows that $\text{Ent}_\mu(\lambda f) = \lambda \text{Ent}_\mu(f)$, i.e. the entropy is homogeneous of order 1.

Theorem 11.10. (Log-Sobolev Inequality, Gross, [Led01]) *Let $f: \mathbb{R}^n \rightarrow \mathbb{R}$ be smooth. As usual, let $\gamma := (1/(2\pi)^{n/2})e^{-\|x\|_2^2/2} dx$ be the standard Gaussian measure on \mathbb{R}^n , and let $|\nabla f| := \|\nabla f\|_2$. Then*

$$\text{Ent}_\gamma(f^2) \leq 2 \int |\nabla f|^2 d\gamma$$

Proof. Recall that the operator $L := \Delta - x \cdot \nabla$ on \mathbb{R}^n has the associated semigroup $\{P_t\}_{t \geq 0}$ with representation

$$P_t f(x) = \int f(e^{-t}x + (1 - e^{-2t})^{1/2}y) d\gamma(y), \quad t \geq 0, x \in \mathbb{R}^n$$

This semigroup is known as the Ornstein-Uhlenbeck semigroup. Now, let f be smooth and non-negative with $\varepsilon \leq f \leq 1/\varepsilon$ for some $\varepsilon > 0$. Via abstract semigroup theory or direct verification with the above formula for P_t , we have: $P_t = e^{tL}$, $(\partial/\partial t)P_t f = LP_t f$,

$P_0 f = f$, $P_\infty f := \lim_{t \rightarrow \infty} P_t f = \int f d\gamma$. Also, recall the Gaussian integration by parts formula $\int f(-Lg)d\gamma = \int \nabla f \cdot \nabla g d\gamma$, via real variable integration by parts.

Lastly, we need the formula $\int P_t f(x) d\gamma(x) = \int f d\gamma$. This follows immediately from the fact that, if X, Y are independent standard Gaussians, then $\alpha X + \sqrt{1 - \alpha^2} Y$ is also a standard Gaussian. Thus,

$$\int P_t f(x) d\gamma(x) = \mathbf{E} f(e^{-t} X + \sqrt{1 - e^{-2t}} Y) = \mathbf{E} f(X') = \int f d\gamma$$

where X' is another standard Gaussian. Alternately,

$$\frac{\partial}{\partial t} \int P_t f = \int \frac{\partial}{\partial t} P_t f = \int L P_t f = - \int P_t f \cdot (\nabla 1) = 0$$

In either case, we are ready to begin the proof.

$$\begin{aligned} \text{Ent}_\gamma(f) &= \int f \log f - \left(\int f \right) \log \left(\int f \right) \\ &= \int (P_0 f \log P_0 f) - P_\infty f \log P_\infty f \\ &= \int_{\mathbb{R}^n} \left(- \int_0^\infty \frac{\partial}{\partial t} (P_t \log P_t) dt \right) d\gamma \quad , \text{integrating by parts} \\ &= - \int_0^\infty \int_{\mathbb{R}^n} \frac{\partial}{\partial t} P_t \log P_t d\gamma dt \\ &= - \int_0^\infty \left(\int L P_t f \log P_t f d\gamma + \int L P_t f d\gamma \right) dt \quad , \text{chain rule} \\ &= - \int_0^\infty \left(- \int \nabla P_t f \cdot \nabla \log P_t f - \int \nabla P_t f \cdot (\nabla 1) \right) dt \quad , L \text{ int. by parts} \\ &= \int_0^\infty \int \frac{|\nabla P_t f|^2}{P_t f} d\gamma dt \end{aligned}$$

Now, from the formula for P_t , we see that $\nabla P_t f = e^{-t} P_t (\nabla f)$. So, using the formula for P_t and Minkowski's inequality we have $|\nabla P_t f| \leq e^{-t} P_t (|\nabla f|)$. Note that we may apply Cauchy-Schwarz to $P_t(fg)(x)$ for fixed x on $L_2(\gamma)$, so that $P_t(fg) \leq \sqrt{P_t(f^2)} \sqrt{P_t(g^2)}$. Therefore

$$|\nabla P_t f| \leq e^{-t} P_t (|\nabla f|) = e^{-t} P_t (|\nabla f| \sqrt{f}/\sqrt{f}) \leq e^{-t} \sqrt{P_t f} \sqrt{P_t (|\nabla f|^2/f)}$$

that is,

$$\frac{|\nabla P_t f|^2}{P_t f} \leq e^{-2t} P_t \left(\frac{|\nabla f|^2}{f} \right)$$

So, combining this with our above equality for entropy gives

$$\text{Ent}_\gamma(f) \leq \int_0^\infty e^{-2t} \int P_t \left(\frac{|\nabla f|^2}{f} \right) d\gamma dt = \frac{1}{2} \int \frac{|\nabla f|^2}{f} d\gamma$$

using that $\int P_t g d\gamma = \int g d\gamma$, and then integrating in t . Letting $f = g^2$ and applying the chain rule completes the proof. \square

Remark 11.11. The above applies equally well to measures $d\mu = e^{-U} dx$ where $\text{Hess}U(x) \geq cId$. The result is then

$$\text{Ent}_\mu(f^2) \leq \frac{2}{c} \int |\nabla f|^2 d\mu$$

12. APPENDIX: NOTATION

Let n, m be a positive integers. Let A, B be sets contained in a universal set Ω .

$\mathbb{N} = \{1, 2, \dots\}$ denotes the set of natural numbers

$\mathbb{Z} = \{\dots, -2, -1, 0, 1, 2, \dots\}$ denotes the set of integers

$\mathbb{Q} = \{a/b: a, b, \in \mathbb{Z}, b \neq 0\}$ denotes the set of rational numbers

\mathbb{R} denotes the set of real numbers

$\mathbb{C} = \{a + b\sqrt{-1}: a, b \in \mathbb{R}\}$ denotes the set of complex numbers

\in means “is an element of.” For example, $2 \in \mathbb{R}$ is read as “2 is an element of \mathbb{R} .”

\forall means “for all”

\exists means “there exists”

$\mathbb{R}^n = \{(x_1, x_2, \dots, x_n): x_i \in \mathbb{R} \forall 1 \leq i \leq n\}$

$f: A \rightarrow B$ means f is a function with domain A and range B . For example,

$f: \mathbb{R}^2 \rightarrow \mathbb{R}$ means that f is a function with domain \mathbb{R}^2 and range \mathbb{R}

\emptyset denotes the empty set

$A \subseteq B$ means $\forall a \in A$, we have $a \in B$, so A is contained in B

$A \setminus B := \{a \in A: a \notin B\}$

$A^c := \Omega \setminus A$, the complement of A in Ω

$A \cap B$ denotes the intersection of A and B

$A \cup B$ denotes the union of A and B

$A \Delta B := (A \setminus B) \cup (B \setminus A)$

\mathbf{P} denotes a probability law on Ω

Let $n \geq m \geq 0$ be integers. We define

$$\binom{n}{m} := \frac{n!}{(n-m)!m!} = \frac{n(n-1)\cdots(n-m+1)}{m(m-1)\cdots(2)(1)}.$$

Let a_1, \dots, a_n be real numbers. Let n be a positive integer.

$$\sum_{i=1}^n a_i = a_1 + a_2 + \cdots + a_{n-1} + a_n.$$

$$\prod_{i=1}^n a_i = a_1 \cdot a_2 \cdots a_{n-1} \cdot a_n.$$

$\min(a_1, a_2)$ denotes the minimum of a_1 and a_2 .

$\max(a_1, a_2)$ denotes the maximum of a_1 and a_2 .

The min of a set of nonnegative real numbers is the smallest element of that set. We also define $\min(\emptyset) := \infty$.

Let $A \subseteq \mathbb{R}$.

$\sup A$ denotes the supremum of A , i.e. the least upper bound of A .
 $\inf A$ denotes the infimum of A , i.e. the greatest lower bound of A .

Let $X: \Omega \rightarrow \mathbb{R}$ be a random variable on a probability space $(\Omega, \mathcal{F}, \mu)$.

$\mathbf{E}(X)$ denotes the expected value of X

$\|X\|_p := (\mathbf{E}|X|^p)^{1/p}$, denotes the L_p -norm of X when $1 \leq p < \infty$

$\|X\|_\infty := \inf\{c > 0: \mathbf{P}(|X| \leq c) = 1\}$, denotes the L_∞ -norm of X

$\text{var}(X) = \mathbf{E}(X - \mathbf{E}(X))^2$, the variance of X

$\sigma_X = \sqrt{\text{var}(X)}$, the standard deviation of X

Let $A \subseteq \Omega$.

$\mathbf{E}(X|A) := \mathbf{E}(X1_A)/\mathbf{P}(A)$ denotes the expected value of X conditioned on the event A .

$1_A: \Omega \rightarrow \{0, 1\}$, denotes the indicator function of A , so that

$$1_A(\omega) = \begin{cases} 1 & , \text{if } \omega \in A \\ 0 & , \text{otherwise.} \end{cases}$$

Let $\langle \cdot, \cdot \rangle: \mathbb{R}^n \times \mathbb{R}^n \rightarrow \mathbb{R}$ be the standard inner product on \mathbb{R}^n , so that for any $x = (x_1, \dots, x_n), y = (y_1, \dots, y_n) \in \mathbb{R}^n$, we have $\langle x, y \rangle := \sum_{i=1}^n x_i y_i$. We also denote $\|x\| := (\sum_{i=1}^n x_i^2)^{1/2}$ as the standard norm on \mathbb{R}^n .

Let X be a random variable on a sample space Ω , so that $X: \Omega \rightarrow \mathbb{R}$. Let \mathbf{P} be a probability law on Ω . Let $x, t \in \mathbb{R}$.

$$F_X(x) = \mathbf{P}(X \leq x) = \mathbf{P}(\{\omega \in \Omega: X(\omega) \leq x\})$$

the Cumulative Distribution Function of X .

$$M_X(t) = \mathbf{E}e^{tX} \text{ denotes the Moment Generating Function of } X \text{ at } t \in \mathbb{R}$$

Let $g, h: \mathbb{R} \rightarrow \mathbb{R}$. Let $t \in \mathbb{R}$.

$$(g * h)(t) = \int_{-\infty}^{\infty} g(x)h(t-x)dx \text{ denotes the convolution of } g \text{ and } h \text{ at } t \in \mathbb{R}$$

Let $f, g: \mathbb{R} \rightarrow \mathbb{C}$. We use the notation $f(t) = o(g(t)), \forall t \in \mathbb{R}$ to denote $\lim_{t \rightarrow \infty} \left| \frac{f(t)}{g(t)} \right| = 0$. We use the notation $f(t) = O(g(t))$ to denote that $\exists c > 0$ such that $|f(t)| \leq c|g(t)|$ for all $t \in \mathbb{R}$. We write $f(t) = \Omega(g(t))$ when $\exists c > 0$ such that $|f(t)| \geq c|g(t)|$ for all $t \in \mathbb{R}$. We write $f(t) = \Theta(g(t))$ when $f(t) = O(g(t))$ and $g(t) = O(f(t))$.

REFERENCES

- [ABSS97] Sanjeev Arora, László Babai, Jacques Stern, and Z. Sweedyk, *The hardness of approximate optima in lattices, codes, and systems of linear equations*, J. Comput. Syst. Sci. **54** (1997), no. 2, 317–331.
- [ACKS15] Pranjal Awasthi, Moses Charikar, Ravishankar Krishnaswamy, and Ali Kemal Sinop, *The hardness of approximation of euclidean k -means*, Preprint, [arXiv:1502.03316](https://arxiv.org/abs/1502.03316), 2015.
- [BFKV98] Avrim Blum, Alan M. Frieze, Ravi Kannan, and Santosh Vempala, *A polynomial-time algorithm for learning noisy linear threshold functions*, Algorithmica **22** (1998), no. 1/2, 35–52.
- [Bob97] S. G. Bobkov, *An isoperimetric inequality on the discrete cube, and an elementary proof of the isoperimetric inequality in gauss space*, Ann. Probab. **25** (1997), no. 1, 206–214.
- [Bor03] Christer Borell, *The Ehrhard inequality*, C. R. Math. Acad. Sci. Paris **337** (2003), no. 10, 663–666. MR 2030108 (2004k:60102)
- [CEM⁺15] Michael B. Cohen, Sam Elder, Cameron Musco, Christopher Musco, and Madalina Persu, *Dimensionality reduction for k -means clustering and low rank approximation*, Proceedings of the Forty-seventh Annual ACM Symposium on Theory of Computing (New York, NY, USA), STOC '15, ACM, 2015, pp. 163–172.
- [Cha05] Sourav Chatterjee, *An error bound in the Sudakov-Fernique inequality*, Preprint, [arXiv.0510424](https://arxiv.org/abs/0510424), 2005.
- [Che09] Ke Chen, *On coresets for k -median and k -means clustering in metric and euclidean spaces and their applications*, SIAM Journal on Computing **39** (2009), no. 3, 923–947.
- [CW09] Kenneth L. Clarkson and David P. Woodruff, *Numerical linear algebra in the streaming model*, Proceedings of the Forty-first Annual ACM Symposium on Theory of Computing (New York, NY, USA), STOC '09, ACM, 2009, pp. 205–214.
- [DKS18] Ilias Diakonikolas, Daniel M. Kane, and Alistair Stewart, *Learning geometric concepts with nasty noise*, Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2018, Los Angeles, CA, USA, June 25-29, 2018, 2018, pp. 1061–1073.
- [FGKP06] Vitaly Feldman, Parikshit Gopalan, Subhash Khot, and Ashok Kumar Ponnuswami, *New results for learning noisy parities and halfspaces*, 47th Annual IEEE Symposium on Foundations of Computer Science (FOCS 2006), 21-24 October 2006, Berkeley, California, USA, Proceedings, 2006, pp. 563–574.
- [FGRW12] Vitaly Feldman, Venkatesan Guruswami, Prasad Raghavendra, and Yi Wu, *Agnostic learning of monomials by halfspaces is hard*, SIAM J. Comput. **41** (2012), no. 6, 1558–1590. MR 3029261
- [FMS07] Dan Feldman, Morteza Monemizadeh, and Christian Sohler, *A ptas for k -means clustering based on weak coresets*, Proceedings of the Twenty-third Annual Symposium on Computational Geometry (New York, NY, USA), SCG '07, ACM, 2007, pp. 11–18.
- [FS97] Yoav Freund and Robert E Schapire, *A decision-theoretic generalization of on-line learning and an application to boosting*, J. Comput. Syst. Sci. **55** (1997), no. 1, 119–139.
- [Gal14] François Le Gall, *Powers of tensors and fast matrix multiplication*, Preprint, [arXiv:1401.7714](https://arxiv.org/abs/1401.7714). ISAAC 2014., 2014.
- [HMT11] Nathan Halko, Per-Gunnar Martinsson, and Joel A. Tropp, *Finding structure with randomness: Probabilistic algorithms for constructing approximate matrix decompositions*, SIAM Review **53** (2011), no. 2, 217–288.
- [HPM04] Sarel Har-Peled and Soham Mazumdar, *On coresets for k -means and k -median clustering*, Proceedings of the Thirty-sixth Annual ACM Symposium on Theory of Computing (New York, NY, USA), STOC '04, ACM, 2004, pp. 291–300.
- [JP78] David S. Johnson and Franco P. Preparata, *The densest hemisphere problem*, Theor. Comput. Sci. **6** (1978), 93–107.
- [KKMS08] Adam Tauman Kalai, Adam R. Klivans, Yishay Mansour, and Rocco A. Servedio, *Agnostically learning halfspaces*, SIAM J. Comput. **37** (2008), no. 6, 1777–1805.
- [KKP17] Daniel Kane, Sushrut Karmalkar, and Eric Price, *Robust polynomial regression up to the information theoretic limit*, 58th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2017, Berkeley, CA, USA, October 15-17, 2017, 2017, pp. 391–402.

- [KMN⁺04] Tapas Kanungo, David M. Mount, Nathan S. Netanyahu, Christine D. Piatko, Ruth Silverman, and Angela Y. Wu, *A local search approximation algorithm for k -means clustering*, *Comput. Geom.* **28** (2004), no. 2-3, 89–112. MR 2062789 (2005a:68210)
- [KR08] Subhash Khot and Oded Regev, *Vertex cover might be hard to approximate to within $2-\varepsilon$* , *J. Comput. Syst. Sci.* **74** (2008), no. 3, 335–349.
- [KV94] Michael J. Kearns and Umesh V. Vazirani, *An introduction to computational learning theory*, MIT Press, Cambridge, MA, USA, 1994.
- [KV09] Ravi Kannan and Santosh Vempala, *Spectral algorithms*, *Foundations and Trends in Theoretical Computer Science* **4** (2009), no. 3-4, 157–288.
- [Led96] Michel Ledoux, *Isoperimetry and Gaussian analysis*, *Lectures on probability theory and statistics (Saint-Flour, 1994)*, *Lecture Notes in Math.*, vol. 1648, Springer, Berlin, 1996, pp. 165–294. MR 1600888 (99h:60002)
- [Led01] ———, *The concentration of measure phenomenon*, *Mathematical Surveys and Monographs*, vol. 89, American Mathematical Society, Providence, RI, 2001. MR 1849347 (2003k:28019)
- [LOGT12] James R. Lee, Shayan Oveis Gharan, and Luca Trevisan, *Multi-way spectral partitioning and higher-order Cheeger inequalities*, *STOC'12—Proceedings of the 2012 ACM Symposium on Theory of Computing*, ACM, New York, 2012, pp. 1117–1130. MR 2961569
- [Mah11] Michael W. Mahoney, *Randomized algorithms for matrices and data*, *Found. Trends Mach. Learn.* **3** (2011), no. 2, 123–224.
- [Mat00] J. Matoušek, *On approximate geometric k -clustering*, *Discrete Comput. Geom.* **24** (2000), no. 1, 61–84. MR 1765234 (2001e:52036)
- [MMR18] Konstantin Makarychev, Yury Makarychev, and Ilya P. Razenshteyn, *Performance of johnson-lindenstrauss transform for k -means and k -medians clustering*, *CoRR* **abs/1811.03195** (2018).
- [Sch90] Robert E. Schapire, *The strength of weak learnability*, *Machine Learning* **5** (1990), no. 2, 197–227.

USC MATHEMATICS, LOS ANGELES, CA
E-mail address: `stevenmheilman@gmail.com`