A Short Course on

Approximation Theory

Math 682 Summer 1998

N. L. Carothers

Department of Mathematics and Statistics

Bowling Green State University

Table of Contents

Preface

These are notes for a six week summer course on approximation theory that I offer occasionally at Bowling Green State University. Our summer classes meet for 90 minutes, five days a week, for a total of 45 hours. But the pace is somewhat leisurely and there is probably not quite enough material here for a "regulation" one semester (45 hour) course. On the other hand, there is more than enough material here for a one quarter (30 hour) course and evidently enough for a five or six week summer course.

I should stress that my presentation here is by no means original: I borrow heavily from a number of well known texts on approximation theory (see the list of references at the end of these notes). I use T. J. Rivlin's book, *An Introduction to the Approximation of Functions*, as a complementary text and thus you will see many references to Rivlin throughout the notes. Also, a few passages here and there are taken from my book, *Real Analysis*. In particular, large portions of these notes are based on copyrighted material. They are offered here solely as an aid to teachers and students of approximation theory and are intended for limited personal use only. I should also point out that I am not an expert in approximation theory and I make no claims that the material presented here is in current fashion among experts in the field.

My interest in approximation theory stems from its beauty, its utility, and its rich history. There are also many connections that can be drawn to questions in both classical and modern analysis. For the purposes of this short introductory course, I focus on a handful of classical topics (with a little bit of modern terminology here and there) and "name" theorems. Indeed, the Weierstrass approximation theorem, along with its various relatives, is the central theme of the course.

In terms of prerequisites, I assume at least a one semester course in advanced calculus or real analysis (compactness, completeness, uniform convergence, uniform continuity, normed spaces, etc.) along with a course in linear algebra. The first chapter, entitled Preliminaries, contains four brief appendices that provide an all too brief review of such topics; they are included in order to make the notes as self-contained as possible. The course is designed for beginning master's students (in both pure and applied mathematics), but should be largely accessible to advanced undergraduates. From my experience, there are plenty of topics here that even advanced PhD students will find entertaining.

Introduction

In 1853, the great Russian mathematician, P. L. Chebyshev [Čebyšev], while working on a problem of *linkages*, devices which translate the linear motion of a steam engine into the circular motion of a wheel, considered the following problem:

> Given a continuous function f defined on a closed interval $[\,a,b\,]$ and a positive integer n, can we "represent" f by a polynomial $p(x) = \sum_{k=0}^{n} a_k x^k$, of degree at most n, in such a way that the maximum error at any point x in $[\,a,b\,]$ is controlled? In particular, is it possible to construct p in such a way that the error $\max\limits_{a\le x\le b} |f(x) - p(x)|$ is minimized?

This problem raises several questions, the first of which Chebyshev himself ignored:

- Why should such a polynomial even *exist*?
- If it does, can we hope to *construct* it?
- If it exists, is it also *unique*?
- What happens if we change the measure of the error to, say, $\int_a^b |f(x) - p(x)|^2\,dx$?

Chebyshev's problem is perhaps best understood by rephrasing it in modern terms. What we have here is a problem of *linear approximation* in a *normed linear space*. Recall that a *norm* on a (real) *vector space* X is a nonnegative function on X satisfying

$$\|x\| \ge 0, \text{ and } \|x\| = 0 \iff x = 0$$

$$\|\alpha x\| = |\alpha|\|x\| \text{ for } \alpha \in \mathbb{R}$$

$$\|x + y\| \le \|x\| + \|y\| \text{ for any } x,\ y \in X.$$

Any norm on X induces a metric or distance function by setting $\text{dist}(x,y) = \|x - y\|$. The abstract version of our problem(s) can now be restated:

- Given a subset (or even a *subspace*) Y of X and a point $x \in X$, is there an element $y \in Y$ which is "nearest" to x; that is, can we find a vector $y \in Y$ such that $\|x - y\| = \inf\limits_{z\in Y} \|x - z\|$? If there is such a "best approximation" to x from elements of Y, is it unique?

Examples

1. In $X = \mathbb{R}^n$ with its usual norm $\|(x_k)_{k=1}^n\|_2 = \left(\sum_{k=1}^n |x_k|^2\right)^{1/2}$, the problem has a complete solution for any subspace (or, indeed, any *closed convex set*) Y. This problem is often considered in Calculus or Linear Algebra where it is called "least-squares approximation." A large part of the current course will be taken up with least-squares approximations, too. For now let's simply note that the problem changes character dramatically if we consider a different norm on $\mathbb{R}^n$.

   Consider $X = \mathbb{R}^2$ under the norm $\|(x,y)\| = \max\{|x|, |y|\}$, and consider the subspace $Y = \{(0,y) : y \in \mathbb{R}\}$ (i.e., the y-axis). It's not hard to see that the point $x = (1,0) \in \mathbb{R}^2$ has infinitely many nearest points in Y; indeed, every point $(0,y)$, $-1 \le y \le 1$, is nearest to x.

2. There are many norms we might consider on $\mathbb{R}^n$. Of particular interest are the ℓ_p-norms; that is, the scale of norms:

$$\|(x_i)_{i=1}^n\|_p = \left(\sum_{k=1}^n |x_k|^p\right)^{1/p}, \qquad 1 \le p < \infty,$$

   and

$$\|(x_i)_{i=1}^n\|_\infty = \max_{1 \le i \le n} |x_i|.$$

   It's easy to see that $\|\cdot\|_1$ and $\|\cdot\|_\infty$ define norms. The other cases take a bit more work; we'll supply full details later.

3. Our original problem concerns $X = C[\,a,b\,]$, the space of all continuous functions $f : [\,a,b\,] \to \mathbb{R}$ under the *uniform norm* $\|f\| = \max_{a \le x \le b} |f(x)|$. The word "uniform" is used because convergence in this norm is the same as uniform convergence on $[\,a,b\,]$:

$$\|f_n - f\| \to 0 \iff f_n \rightrightarrows f \text{ on } [\,a,b\,].$$

   In this case we're interested in approximations by elements of $Y = \mathcal{P}_n$, the *subspace* of all polynomials of degree at most n in $C[\,a,b\,]$. It's not hard to see that $\mathcal{P}_n$ is a finite-dimensional subspace of $C[\,a,b\,]$ of dimension exactly $n+1$. (Why?)

   If we consider the subspace $Y = \mathcal{P}$ consisting of *all* polynomials in $X = C[\,a,b\,]$, we readily see that the existence of best approximations can be problematic. It follows from the Weierstrass theorem, for example, that each $f \in C[\,a,b\,]$ has distance 0 from $\mathcal{P}$ but, since not every $f \in C[\,a,b\,]$ is a polynomial (why?), we can't hope for a best approximating polynomial to exist in every case. For example, the function $f(x) = x\sin(1/x)$ is continuous on $[\,0,1\,]$ but can't possibly agree with any polynomial on $[\,0,1\,]$. (Why?)

The key to the problem of polynomial approximation is the fact that each $\mathcal{P}_n$ is *finite-dimensional*. To see this, it will be most efficient to consider the abstract setting of finite-dimensional subspaces of arbitrary normed spaces.

"Soft" Approximation

Lemma. *Let V be a finite-dimensional vector space. Then, all norms on V are equivalent. That is, if $\|\cdot\|$ and $|||\cdot|||$ are norms on V, then there exist constants $0 < A, B < \infty$ such that*

$$A\,\|x\| \le |||\,x\,||| \le B\,\|x\|$$

for all vectors $x \in V$.

PROOF. Suppose that V is n-dimensional and that $\|\cdot\|$ is a norm on V. Fix a basis $e_1, \ldots, e_n$ for V and consider the norm

$$\left\|\sum_{i=1}^n a_i e_i\right\|_1 = \sum_{i=1}^n |a_i| = \|(a_i)_{i=1}^n\|_1$$

for $x = \sum_{i=1}^n a_i e_i \in V$. Since $e_1, \ldots, e_n$ is a basis for V, it's not hard to see that $\|\cdot\|_1$ is, indeed, a norm on V. It now suffices to show that $\|\cdot\|$ and $\|\cdot\|_1$ are equivalent. (Why?)

One inequality is easy to show; indeed, notice that

$$\left\|\sum_{i=1}^n a_i e_i\right\| \le \sum_{i=1}^n |a_i|\,\|e_i\| \le \left(\max_{1\le i\le n}\|e_i\|\right)\sum_{i=1}^n |a_i| = B\left\|\sum_{i=1}^n a_i e_i\right\|_1.$$

The real work comes in establishing the other inequality.

To begin, notice that we've actually set-up a correspondence between $\mathbb{R}^n$ and V; specifically, the map $(a_i)_{i=1}^n \mapsto \sum_{i=1}^n a_i e_i$ is obviously both one-to-one and onto. Moreover, this correspondence is an *isometry* between $(\mathbb{R}^n, \|\cdot\|_1)$ and $(V, \|\cdot\|_1)$.

Now the inequality we've just established shows that the function $x \mapsto \|x\|$ is *continuous* on the space $(V, \|\cdot\|_1)$ since

$$\big|\, \|x\| - \|y\| \,\big| \le \|x - y\| \le B\, \|x - y\|_1$$

for any x, $y \in V$. Thus, $\|\cdot\|$ assumes a *minimum* value on the *compact* set

$$S = \{x \in V : \|x\|_1 = 1\}.$$

(Why is S compact?) In particular, there is some $A > 0$ such that $\|x\| \ge A$ whenever $\|x\|_1 = 1$. (Why can we assume that $A > 0$?) The inequality we need now follows from the homogeneity of the norm:

$$\left\| \frac{x}{\|x\|_1} \right\| \ge A \Longrightarrow \|x\| \ge A\, \|x\|_1. \qquad \square$$

Corollary. *Given $a < b$ (fixed) and a positive integer n, there exist $0 < A_n$, $B_n < \infty$ (constants which may depend on n) such that*

$$A_n \sum_{k=0}^{n} |a_k| \le \max_{a \le x \le b} \left| \sum_{k=0}^{n} a_k x^k \right| \le B_n \sum_{k=0}^{n} |a_k|.$$

Exercise

Find explicit "formulas" for A_n and B_n, above. (This can be done without any fancy theorems.) If it helps, you may consider the case $[\,a, b\,] = [\,0, 1\,]$.

Corollary. *Let Y be a finite-dimensional normed space and let $M > 0$. Then, any closed ball $\{y \in Y : \|y\| \le M\}$ is compact.*

PROOF. Again suppose that Y is n-dimensional and that $e_1, \ldots, e_n$ is a basis for Y. From our previous lemma we know that there is some constant $A > 0$ such that

$$A \sum_{i=1}^{n} |a_i| \le \left\| \sum_{i=1}^{n} a_i e_i \right\|$$

for all $x = \sum_{i=1}^{n} a_i e_i \in Y$. In particular,

$$A\, |a_i| \le \left\| \sum_{i=1}^{n} a_i e_i \right\| \le M \implies |a_i| \le \frac{M}{A} \text{ for } i = 1, \ldots, n.$$

Thus, $\{y \in Y : \|y\| \le M\}$ is a *closed* subset (why?) of the *compact* set

$$\left\{ x = \sum_{i=1}^{n} a_i e_i : |a_i| \le \frac{M}{A},\ i = 1, \ldots, n \right\}. \qquad \square$$

Corollary. *Every finite-dimensional normed space is complete. In particular, if Y is a finite-dimensional subspace of a normed linear space X, then Y is a closed subset of X.*

Theorem. *Let Y be a finite-dimensional subspace of a normed linear space X, and let $x \in X$. Then, there exists a (not necessarily unique) $y^* \in Y$ such that*

$$\|x - y^*\| = \min_{y \in Y} \|x - y\|$$

for all $y \in Y$. That is, there is a best approximation to x by elements of Y.

Proof. First notice that since $0 \in Y$, we know that a nearest point y^* will satisfy $\|x - y^*\| \le \|x\| = \|x - 0\|$. Thus, it suffices to look for y^* among the vectors $y \in Y$ satisfying $\|x - y\| \le \|x\|$. It will be convenient to use a slightly larger set of vectors, though. By the triangle inequality,

$$\|x - y\| \le \|x\| \Longrightarrow \|y\| \le \|x - y\| + \|x\| \le 2\|x\|.$$

Thus, we may restrict our attention to those y's in the *compact* set

$$K = \{y \in Y : \|y\| \le 2\|x\|\}.$$

To finish the proof, we need only notice that the function $f(y) = \|x - y\|$ is *continuous*:

$$|f(y) - f(z)| = \big|\, \|x - y\| - \|x - z\| \,\big| \le \|y - z\|,$$

hence attains a minimum value at some point $y^* \in K$. $\square$

Corollary. *For each $f \in C[\,a, b\,]$, and each positive integer n, there is a (not necessarily unique) polynomial $p_n^* \in \mathcal{P}_n$ such that*

$$\|f - p_n^*\| = \min_{p \in \mathcal{P}_n} \|f - p\|.$$

Corollary. *Given $f \in C[a,b]$ and a (fixed) positive integer n, there exists a constant $R < \infty$ such that if*

$$\left\| f - \sum_{k=0}^{n} a_k x^k \right\| \le \|f\|,$$

then $\max_{0 \le k \le n} |a_k| \le R$.

Examples

Nothing in our Corollary says that p_n^* will be a polynomial of degree *exactly* n—rather, a polynomial of degree *at most* n. For example, the best approximation to $f(x) = x$ by a polynomial of degree at most 3 is, of course, $p(x) = x$. Even examples of non-polynomial functions are easy to come by; for instance, the best linear approximation to $f(x) = |x|$ on $[-1, 1]$ is actually the constant function $p(x) = 1/2$, and this makes for an entertaining **exercise**.

Before we leave these "soft" arguments behind, let's discuss the problem of *uniqueness* of best approximations. First, let's see why we want best approximations to be unique:

Lemma. *Let Y be a finite-dimensional subspace of a normed linear space X, and suppose that each $x \in X$ has a unique nearest point $y_x \in Y$. Then, the nearest point map $x \mapsto y_x$ is continuous.*

Proof. Let's write $P(x) = y_x$ for the nearest point map, and let's suppose that $x_n \to x$ in X. We want to show that $P(x_n) \to P(x)$, and for this it's enough to show that there is a subsequence of $(P(x_n))$ which converges to $P(x)$. (Why?)

Since the sequence (x_n) is bounded in X, say $\|x_n\| \le M$ for all n, we have

$$\|P(x_n)\| \le \|P(x_n) - x_n\| + \|x_n\| \le 2\|x_n\| \le 2M.$$

Thus, $(P(x_n))$ is a bounded sequence in Y, a finite-dimensional space. As such, by passing to a subsequence, we may suppose that $(P(x_n))$ converges to some element $P_0 \in Y$. (How?) Now we need to show that $P_0 = P(x)$. But

$$\|P(x_n) - x_n\| \le \|P(x) - x_n\| \qquad \text{(why?)},$$

for any n, and hence, letting $n \to \infty$,

$$\|P_0 - x\| \le \|P(x) - x\|.$$

Since nearest points in Y are unique, we must have $P_0 = P(x)$. $\square$

Exercise

Let X be a normed linear space and let $P : X \to X$. Show that P is continuous at $x \in X$ if and only if, whenever $x_n \to x$ in X, some subsequence of $(P(x_n))$ converges to $P(x)$. [Hint: The forward direction is easy; for the backward implication, suppose that $(P(x_n))$ fails to converge to $P(x)$ and work toward a contradiction.]

It should be pointed out that the nearest point map is, in general, *nonlinear* and, as such, can be very difficult to work with. Later we'll see at least one case in which nearest point maps always turn out to be linear.

We next observe that the set of best approximations is always pretty reasonable:

Theorem. *Let Y be a subspace of a normed linear space X, and let $x \in X$. The set Y_x, consisting of all best approximations to x out of Y, is a bounded, convex set.*

PROOF. As we've seen, the set Y_x is a subset of $\{y \in X : \|y\| \le 2\|x\|\}$ and, hence, is bounded.

Now recall that a subset K of a vector space V is said to be *convex* if K contains the line segment joining any pair of its points. Specifically, K is convex if

$$x, y \in K,\ 0 \le \lambda \le 1 \Longrightarrow \lambda x + (1-\lambda)y \in K.$$

Now, y_1, $y_2 \in Y_x$ means that

$$\|x - y_1\| = \|x - y_2\| = \min_{y \in Y} \|x - y\|.$$

Next, given $0 \le \lambda \le 1$, set $y^* = \lambda y_1 + (1-\lambda)y_2$. We want to show that $y^* \in Y_x$, but notice that we at least have $y^* \in Y$. Finally, we estimate:

$$\begin{aligned}\|x - y^*\| &= \|x - (\lambda y_1 + (1-\lambda)y_2)\| \\ &= \|\lambda(x - y_1) + (1-\lambda)(x - y_2)\| \\ &\le \lambda\|x - y_1\| + (1-\lambda)\|x - y_2\| \\ &= \min_{y \in Y} \|x - y\|.\end{aligned}$$

Hence, $\|x - y^*\| = \min_{y \in Y} \|x - y\|$; that is, $y^* \in Y_x$. $\square$

Exercise

If, in addition, Y is finite-dimensional, show that Y_x is *closed* (hence *compact*).

If Y_x contains more than one point, then, in fact, it contains an entire line segment. Thus, Y_x is either empty, contains exactly one point, or contains infinitely many points. This observation gives us a *sufficient condition* for uniqueness of nearest points: If our normed space X contains no line segments on any sphere $\{x \in X : \|x\| = r\}$, then any best approximation (out of any set) will be unique.

A norm $\|\cdot\|$ on a vector space X is said to be *strictly convex* if, for any $x \neq y \in X$ with $\|x\| = r = \|y\|$, we always have $\|\lambda x + (1 - \lambda) y\| < r$ for any $0 < \lambda < 1$. That is, the open line segment between any pair of points on the surface of the ball of radius r in X lies entirely inside the ball. We often simply say that the space X is strictly convex, with the understanding that a property of the norm in X is implied. Here's an immediate corollary to our last result:

Corollary. *If X has a strictly convex norm, then, for any subspace Y of X and any point $x \in X$, there can be at most one best approximation to x out of Y. That is, Y_x is either empty or consists of a single point.*

In order to arrive at a condition that's somewhat easier to check, let's translate our original definition into a statement about the triangle inequality in X.

Lemma. *X has a strictly convex norm if and only if the triangle inequality is strict on non-parallel vectors; that is, if and only if*

$$x \neq \alpha y,\ y \neq \alpha x,\ \textit{all } \alpha \in \mathbb{R} \Longrightarrow \|x + y\| < \|x\| + \|y\|.$$

PROOF. First suppose that X is strictly convex, and let x and y be non-parallel vectors in X. Then, in particular, the vectors $x/\|x\|$ and $y/\|y\|$ must be different. (Why?) Hence,

$$\left\| \left(\frac{\|x\|}{\|x\| + \|y\|} \right) \frac{x}{\|x\|} + \left(\frac{\|y\|}{\|x\| + \|y\|} \right) \frac{y}{\|y\|} \right\| < 1.$$

That is, $\|x + y\| < \|x\| + \|y\|$.

Next suppose that the triangle inequality is strict on non-parallel vectors, and let $x \neq y \in X$ with $\|x\| = r = \|y\|$. If x and y are parallel, then we must have $y = -x$. (Why?) In this case,

$$\|\lambda x + (1-\lambda)\, y\| = |2\lambda - 1|\, \|x\| < r,$$

since $|2\lambda - 1| < 1$ whenever $0 < \lambda < 1$. Otherwise, x and y are non-parallel. In this case, for any $0 < \lambda < 1$, the vectors λx and $(1-\lambda)\, y$ are likewise non-parallel. Thus,

$$\|\lambda x + (1-\lambda)\, y\| < \lambda\|x\| + (1-\lambda)\|y\| = r. \qquad \square$$

Examples

1. The usual norm on $C[\,a,b\,]$ is *not* strictly convex (and so the problem of uniqueness of best approximations is all the more interesting to tackle). For example, if $f(x) = x$ and $g(x) = x^2$ in $C[\,0,1\,]$, then $\|f\| = 1 = \|g\|$, $f \neq g$, while $\|f+g\| = 2$. (Why?)
2. The usual norm on $\mathbb{R}^n$ is strictly convex, as is any one of the norms $\|\cdot\|_p$, $1 < p < \infty$. (We'll prove these facts shortly.) The norms $\|\cdot\|_1$ and $\|\cdot\|_\infty$, on the other hand, are *not* strictly convex. (Why?)

Appendix A

For completeness, we supply a few of the missing details concerning the ℓ_p-norms. We begin with a handful of classical inequalities of independent interest. First recall that we have defined a scale of "norms" on $\mathbb{R}^n$ by setting:

$$\|x\|_p = \left(\sum_{i=1}^n |x_i|^p\right)^{1/p}, \qquad 1 \le p < \infty,$$

and

$$\|x\|_\infty = \max_{1\le i\le n} |x_i|,$$

where $x = (x_i)_{i=1}^n \in \mathbb{R}^n$. Please note that the case $p = 2$ gives the usual Euclidean norm on $\mathbb{R}^n$ and that the cases $p = 1$ and $p = \infty$ clearly give rise to legitimate norms on $\mathbb{R}^n$.

Common parlance is to refer to these expressions as ℓ_p-norms and to refer to the space $(\mathbb{R}^n, \|\cdot\|_p)$ as ℓ_p^n. The space of all infinite sequences $x = (x_n)_{n=1}^\infty$ for which the analogous infinite sum (or supremum) $\|x\|_p$ is finite is referred to as ℓ_p. What's more, there is a "continuous" analogue of this scale: We might also consider the norms

$$\|f\|_p = \left(\int_a^b |f(x)|^p \, dx \right)^{1/p}, \qquad 1 \le p < \infty,$$

and

$$\|f\|_\infty = \sup_{a \le x \le b} |f(x)|,$$

where f is in $C[\,a, b\,]$ (or is simply Lebesgue integrable). The subsequent discussion actually covers all of these cases, but we will settle for writing our proofs in the $\mathbb{R}^n$ setting only.

Lemma. (Young's inequality): *Let $1 < p < \infty$, and let $1 < q < \infty$ be defined by $\frac{1}{p} + \frac{1}{q} = 1$; that is, $q = \frac{p}{p-1}$. Then, for any $a,\, b \ge 0$, we have*

$$ab \le \frac{1}{p}\, a^p + \frac{1}{q}\, b^q.$$

Moreover, equality can only occur if $a^p = b^q$. (We refer to p and q as *conjugate* exponents; note that p satisfies $p = \frac{q}{q-1}$. Please note that the case $p = q = 2$ yields the familiar *arithmetic-geometric mean inequality.*)

PROOF. A quick calculation before we begin:

$$q - 1 = \frac{p}{p-1} - 1 = \frac{p - (p-1)}{p-1} = \frac{1}{p-1}.$$

Now we just estimate areas; for this you might find it helpful to draw the graph of $y = x^{p-1}$ (or, equivalently, the graph of $x = y^{q-1}$). Comparing areas we get:

$$ab \le \int_0^a x^{p-1}\, dx + \int_0^b y^{q-1}\, dy = \frac{1}{p}\, a^p + \frac{1}{q}\, b^q.$$

The case for equality also follows easily from the graph of $y = x^{p-1}$ (or $x = y^{q-1}$), since $b = a^{p-1} = a^{p/q}$ means that $a^p = b^q$. $\square$

Corollary. (Hölder's inequality): *Let $1 < p < \infty$, and let $1 < q < \infty$ be defined by $\frac{1}{p} + \frac{1}{q} = 1$. Then, for any $a_1, \ldots, a_n$ and $b_1, \ldots, b_n$ in $\mathbb{R}$ we have:*

$$\sum_{i=1}^n |a_i b_i| \le \left(\sum_{i=1}^n |a_i|^p \right)^{1/p} \left(\sum_{i=1}^n |b_i|^q \right)^{1/q}.$$

(Please note that the case $p = q = 2$ yields the familiar *Cauchy-Schwarz inequality.*)

Moreover, equality in Hölder's inequality can only occur if there exist nonnegative scalars α and β such that $\alpha\, |a_i|^p = \beta\, |b_i|^q$ for all $i = 1, \ldots, n$.

PROOF. Let $A = \left(\sum_{i=1}^n |a_i|^p\right)^{1/p}$ and let $B = \left(\sum_{i=1}^n |b_i|^q\right)^{1/q}$. We may clearly assume that A, $B \neq 0$ (why?), and hence we may divide (and appeal to Young's inequality):

$$\frac{|a_i b_i|}{AB} \leq \frac{|a_i|^p}{pA^p} + \frac{|b_i|^q}{qB^q}.$$

Adding, we get:

$$\frac{1}{AB}\sum_{i=1}^n |a_i b_i| \leq \frac{1}{pA^p}\sum_{i=1}^n |a_i|^p + \frac{1}{qB^q}\sum_{i=1}^n |b_i|^q = \frac{1}{p} + \frac{1}{q} = 1.$$

That is, $\sum_{i=1}^n |a_i b_i| \leq AB$.

The case for equality in Hölder's inequality follows from what we know about Young's inequality: Equality in Hölder's inequality means that either $A = 0$, or $B = 0$, or else $|a_i|^p/pA^p = |b_i|^q/qB^q$ for all $i = 1, \ldots, n$. In short, there must exist nonnegative scalars α and β such that $\alpha\,|a_i|^p = \beta\,|b_i|^q$ for all $i = 1, \ldots, n$. $\square$

Notice, too, that the case $p = 1$ ($q = \infty$) works, and is easy:

$$\sum_{i=1}^n |a_i b_i| \leq \left(\sum_{i=1}^n |a_i|\right)\left(\max_{1\leq i\leq n} |b_i|\right).$$

Exercise

When does equality occur in the case $p = 1$ ($q = \infty$)?

Finally, an application of Hölder's inequality leads to an easy proof that $\|\cdot\|_p$ is actually a norm. It will help matters here if we first make a simple observation: If $1 < p < \infty$ and if $q = \frac{p}{p-1}$, notice that

$$\left\|\,(\,|a_i|^{p-1})_{i=1}^n\right\|_q = \left(\sum_{i=1}^n |a_i|^p\right)^{(p-1)/p} = \|a\|_p^{p-1}.$$

Lemma. (Minkowski's inequality): *Let $1 < p < \infty$ and let $a = (a_i)_{i=1}^n$, $b = (b_i)_{i=1}^n \in \mathbb{R}^n$. Then, $\|a + b\|_p \leq \|a\|_p + \|b\|_p$.*

PROOF. In order to prove the triangle inequality, we once again let q be defined by $\frac{1}{p} + \frac{1}{q} = 1$, and now we use Hölder's inequality to estimate:

$$\sum_{i=1}^n |a_i + b_i|^p = \sum_{i=1}^n |a_i + b_i| \cdot |a_i + b_i|^{p-1}$$

$$\leq \sum_{i=1}^{n} |a_i| \cdot |a_i + b_i|^{p-1} + \sum_{i=1}^{n} |b_i| \cdot |a_i + b_i|^{p-1}$$
$$\leq \|a\|_p \cdot \| (|a_i + b_i|^{p-1})_{i=1}^{n} \|_q + \|y\|_p \cdot \| (|a_i + b_i|^{p-1})_{i=1}^{n} \|_q$$
$$= \|a + b\|_p^{p-1} (\|a\|_p + \|b\|_p) .$$

That is, $\|a + b\|_p^p \leq \|a + b\|_p^{p-1} (\|a\|_p + \|b\|_p)$, and the triangle inequality follows. $\square$

If $1 < p < \infty$, then equality in Minkowski's inequality can only occur if a and b are parallel; that is, the ℓ_p-norm is *strictly convex* for $1 < p < \infty$. Indeed, if $\|a + b\|_p = \|a\|_p + \|b\|_p$, then either $a = 0$, or $b = 0$, or else $a, b \neq 0$ and we have equality at each stage of our proof. Now equality in the first inequality means that $|a_i + b_i| = |a_i| + |b_i|$, which easily implies that a_i and b_i have the same sign. Next, equality in our application of Hölder's inequality implies that there are nonnegative scalars C and D such that $|a_i|^p = C\,|a_i + b_i|^p$ and $|b_i|^p = D\,|a_i + b_i|^p$ for all $i = 1, \ldots, n$. Thus, $a_i = E\,b_i$ for some scalar E and all $i = 1, \ldots, n$.

Of course, the triangle inequality also holds in either of the cases $p = 1$ or $p = \infty$ (with much simpler proofs).

Exercises

When does equality occur in the triangle inequality in the cases $p = 1$ or $p = \infty$? In particular, show that neither of the norms $\| \cdot \|_1$ or $\| \cdot \|_\infty$ is strictly convex.

Appendix B

Next, we provide a brief review of completeness and compactness. Such review is doomed to inadequacy; the reader unfamiliar with these concepts would be well served to consult a text on advanced calculus such as *Analysis in Euclidean Spaces* by K. Hoffman, or *Principles of Mathematical Analysis* by W. Rudin.

To begin, we recall that a subset A of normed space X (such as $\mathbb{R}$ or $\mathbb{R}^n$) is said to be *closed* if A is closed under the taking of sequential limits. That is, A is closed if, whenever (a_n) is a sequence from A converging to some point $x \in X$, we always have $x \in A$. It's not hard to see that any closed interval, such as $[\,a, b\,]$ or $[\,a, \infty)$, is, indeed, a closed subset of $\mathbb{R}$ in this sense. There are, however, much more complicated examples of closed sets in $\mathbb{R}$.

A normed space X is said to be *complete* if every Cauchy sequence from X converges (to a point in X). It is a familiar fact from Calculus that $\mathbb{R}$ is complete, as is $\mathbb{R}^n$. In fact, the completeness of $\mathbb{R}$ is often assumed as an axiom (in the form of the least upper bound axiom). There are, however, many examples of normed spaces which are not complete; that is, there are examples of normed spaces in which Cauchy sequences need not converge.

We say that a subset A of a normed space X is complete if every Cauchy sequence from A converges *to a point in A*. Please note here that we require not only that Cauchy sequences from A converge, but also that the limit be back in A. As you might imagine, the completeness of A depends on properties of both A and the containing space X.

First note that a complete subset is necessarily also closed. Indeed, since every convergent sequence is also Cauchy, it follows that a complete subset is closed.

Exercise

If A is a complete subset of a normed space X, show that A is also closed.

If the containing space X is itself complete, then it's easy to tell which of its subsets are complete. Indeed, since every Cauchy sequence in X converges (somewhere), all we need to know is whether the subset is closed.

Exercise

Let A be a subset of a complete normed space X. Show that A is complete if and only if A is a closed subset of X. In particular, please note that every closed subset of $\mathbb{R}$ (or $\mathbb{R}^n$) is complete.

Finally, we recall that a subset A of a normed space X is said to be *compact* if every sequence from A has a subsequence which converges *to a point in A*. Again, since we have insisted that certain limits remain in A, it's not hard to see that compact sets are necessarily also closed.

Exercise

If A is a compact subset of a normed space X, show that A is also closed.

Moreover, since a Cauchy sequence with a convergent subsequence must itself converge (why?), we actually have that every compact set is necessarily complete.

Exercise

If A is a compact subset of a normed space X, show that A is also complete.

Since the compactness of a subset A has something to do with *every* sequence in A, it's not hard to believe that it is a more stringent property than the others we've considered so far. In particular, it's not hard to see that a compact set must be bounded.

Exercise

If A is a compact subset of a normed space X, show that A is also bounded. [Hint: If not, then A would contain a sequence (a_n) with $\|a_n\| \to \infty$.]

Now it is generally not so easy to describe the compact subsets of a particular normed space X, however, it is quite easy to describe the compact subsets of $\mathbb{R}$ (or $\mathbb{R}^n$). This well-known result goes by many names; we will refer to it as the Heine-Borel theorem.

Theorem. *A subset A of $\mathbb{R}$ (or $\mathbb{R}^n$) is compact if and only if A is both closed and bounded.*

Proof. One direction of the proof is easy: As we've already seen, compact sets in $\mathbb{R}$ are necessarily closed and bounded. For the other direction, notice that if A is a bounded subset of $\mathbb{R}$, then it follows from the Bolzano-Weierstrass theorem that every sequence from A has a subsequence which converges in $\mathbb{R}$. If A is also a closed set, then this limit must, in fact, be back in A. Thus, every sequence in A has a subsequence converging to a point in A. $\square$

Appendix C

We next offer a brief review of pointwise and uniform convergence. We begin with an elementary example:

Example

(a) For each $n = 1, 2, 3, \ldots$, consider the function $f_n(x) = e^x + \frac{x}{n}$ for $x \in \mathbb{R}$. Note that for each (fixed) x the sequence $(f_n(x))_{n=1}^{\infty}$ converges to $f(x) = e^x$ because

$$|f_n(x) - f(x)| = \frac{|x|}{n} \to 0 \quad \text{as} \quad n \to \infty.$$

In this case we say that the sequence of functions (f_n) converges *pointwise* to the function f on $\mathbb{R}$. But notice, too, that the *rate* of convergence depends on x. In particular, in order to get $|f_n(x) - f(x)| < 1/2$ we would need to take $n > 2|x|$. Thus, at $x = 2$, the inequality is satisfied for all $n > 4$, while at $x = 1000$, the inequality is satisfied only for $n > 2000$. In short, the rate of convergence is *not uniform* in x.

(b) Consider the same sequence of functions as above, but now let's suppose that we restrict that values of x to the interval $[-5, 5]$. Of course, we still have that $f_n(x) \to f(x)$ for each (fixed) x in $[-5, 5]$; in other words, we still have that (f_n) converges pointwise to f on $[-5, 5]$. But notice that the rate of convergence is now uniform over x in $[-5, 5]$. To see this, just rewrite the initial calculation:

$$|f_n(x) - f(x)| = \frac{|x|}{n} \leq \frac{5}{n} \quad \text{for} \quad x \in [-5, 5],$$

and notice that the upper bound $5/n$ tends to 0, as $n \to \infty$, *independent of the choice of* x. In this case, we say that (f_n) converges *uniformly* to f on $[-5, 5]$. The point here is that the notion of uniform convergence depends on the underlying domain as well as on the sequence of functions at hand.

With this example in mind, we now offer formal definitions of pointwise and uniform convergence. In both cases we consider a sequence of functions $f_n : X \to \mathbb{R}$, $n = 1, 2, 3, \ldots$, each defined on the same underlying set X, and another function $f : X \to \mathbb{R}$ (the candidate for the limit).

We say that (f_n) *converges pointwise to* f *on* X if, for each $x \in X$, we have $f_n(x) \to f(x)$ as $n \to \infty$; thus, for each $x \in X$ and each $\varepsilon > 0$, we can find an integer N (which depends on ε and which may also depend on x) such that $|f_n(x) - f(x)| < \varepsilon$ whenever $n > N$. A convenient shorthand for pointwise convergence is: $f_n \to f$ on X or, if X is understood, simply $f_n \to f$.

We say that (f_n) *converges uniformly to* f *on* X if, for each $\varepsilon > 0$, we can find an integer N (which depends on ε but *not* on x) such that $|f_n(x) - f(x)| < \varepsilon$ for each $x \in X$, provided that $n > N$. Please notice that the phrase "for each $x \in X$" now occurs well after the phrase "for each $\varepsilon > 0$" and, in particular, that the rate of convergence N does not depend on x. It should be reasonably clear that uniform convergence implies pointwise convergence; in other words, uniform convergence is "stronger" than pointwise convergence. For this reason, we sometimes use the shorthand: $f_n \rightrightarrows f$ on X or, if X is understood, simply $f_n \rightrightarrows f$.

The definition of uniform convergence can be simplified by "hiding" one of the quantifiers under different notation; indeed, note that the phrase "$|f_n(x) - f(x)| < \varepsilon$ for any $x \in X$" is (essentially) equivalent to the phrase "$\sup_{x \in X} |f_n(x) - f(x)| < \varepsilon$." Thus, our definition may be reworded as follows: (f_n) converges uniformly to f on X if, given $\varepsilon > 0$, there is an integer N such that $\sup_{x \in X} |f_n(x) - f(x)| < \varepsilon$ for all $n > N$.

The notion of uniform convergence exists for one very good reason: Continuity is preserved under uniform limits. This fact is well worth stating.

Exercise

Let X be a subset of $\mathbb{R}$, let $f,\ f_n : X \to \mathbb{R}$ for $n = 1, 2, 3, \ldots$, and let $x_0 \in X$. If each f_n is continuous at x_0, and if $f_n \rightrightarrows f$ on X, then f is continuous at x_0. In particular, if each f_n is continuous on all of X, then so is f. Give an example showing that this result may fail if we only assume that $f_n \to f$ on X.

Appendix D

Lastly, we discuss continuity for linear transformations between normed vector spaces. Throughout this section, we consider a linear map $T : V \to W$ between vector spaces V and W; that is we suppose that T satisfies $T(\alpha x + \beta y) = \alpha T(x) + \beta T(y)$ for all $x,\ y \in V$, and all scalars $\alpha,\ \beta$. Please note that every linear map T satisfies $T(0) = 0$. If we further suppose that V is endowed with the norm $\|\cdot\|$, and that W is endowed with the norm $|||\cdot|||$, the we may consider the issue of continuity of the map T.

The key result for our purposes is that, for linear maps, continuity—even at a single point—is equivalent to uniform continuity (and then some!).

Theorem. *Let $(V, \|\cdot\|)$ and $(W, |||\cdot|||)$ be normed vector spaces, and let $T : V \to W$ be a linear map. Then, the following are equivalent:*

(i) *T is Lipschitz;*

(ii) *T is uniformly continuous;*

(iii) *T is continuous (everywhere);*

(iv) *T is continuous at $0 \in V$;*

(v) *there is a constant $C < \infty$ such that $|||\, T(x)\, ||| \le C\|x\|$ for all $x \in V$.*

PROOF. Clearly, (i) $\Longrightarrow$ (ii) $\Longrightarrow$ (iii) $\Longrightarrow$ (iv). We need to show that (iv) $\Longrightarrow$ (v), and that (v) $\Longrightarrow$ (i) (for example). The second of these is easier, so let's start there.

(v) $\Longrightarrow$ (i): If condition (v) holds for a linear map T, then T is Lipschitz (with constant C) since $|||\, T(x) - T(y)\, ||| = |||\, T(x-y)\, ||| \le C\|x - y\|$ for any $x,\ y \in V$.

(iv) $\Longrightarrow$ (v): Suppose that T is continuous at 0. Then we may choose a $\delta > 0$ so that $|||\, T(x)\, ||| = |||\, T(x) - T(0)\, ||| \le 1$ whenever $\|x\| = \|x - 0\| \le \delta$. (How?)

Given $0 \ne x \in V$, we may scale by the factor $\delta/\|x\|$ to get $\big\|\, \delta x/\|x\|\, \big\| = \delta$. Hence, $\big|\big|\big|\, T\big(\delta x/\|x\|\big)\, \big|\big|\big| \le 1$. But $T\big(\delta x/\|x\|\big) = (\delta/\|x\|)\, T(x)$, since T is linear, and so we get $|||\, T(x)\, ||| \le (1/\delta)\|x\|$. That is, $C = 1/\delta$ works in condition (v). (Note that since condition (v) is trivial for $x = 0$, we only care about the case $x \ne 0$.) $\square$

A linear map satisfying condition (v) of the Theorem (i.e., a continuous linear map) is often said to be *bounded*. The meaning in this context is slightly different than usual. Here it means that T maps bounded sets to bounded sets. This follows from the fact that T is Lipschitz. Indeed, if $|||\, T(x)\, ||| \le C\|x\|$ for all $x \in V$, then (as we've seen) $|||\, T(x) - T(y)\, ||| \le C\|x - y\|$ for any $x,\ y \in V$, and hence T maps the ball about x of radius r into the ball about $T(x)$ of radius Cr. In symbols, $T\big(B_r(x)\big) \subset B_{Cr}(\,T(x))$. More generally, T maps a set of diameter d into a set of diameter at most Cd. There's no danger of confusion in our using the word bounded to mean something new here; the ordinary usage of the word (as applied to functions) is uninteresting for linear maps. A nonzero linear map always has an unbounded range. (Why?)

The smallest constant that works in (v) is called the *norm* of the operator T and is usually written $\|T\|$. In symbols,

$$\|T\| \;=\; \sup_{x \ne 0} \frac{|||Tx|||}{\|x\|} \;=\; \sup_{\|x\|=1} |||Tx|||.$$

Thus, T is bounded (continuous) if and only if $\|T\| < \infty$.

The fact that all norms on a finite-dimensional normed space are equivalent provides a final (rather spectacular) corollary.

Corollary. *Let V and W be normed vector spaces with V finite-dimensional. Then, every linear map $T : V \to W$ is continuous.*

Proof. Let $x_1, \ldots, x_n$ be a basis for V and let $\|\sum_{i=1}^n \alpha_i x_i\|_1 = \sum_{i=1}^n |\alpha_i|$, as before. From the Lemma on page 3, we know that there is a constant $B < \infty$ such that $\|x\|_1 \le B\,\|x\|$ for every $x \in V$.

Now if $T : (V, \|\cdot\|\,) \to (W, |||\cdot|||)$ is linear, we get

$$\begin{aligned}
\left|\left|\left| \, T\left(\sum_{i=1}^n \alpha_i x_i\right) \right|\right|\right| &= \left|\left|\left| \sum_{i=1}^n \alpha_i T(x_i) \right|\right|\right| \\
&\le \sum_{i=1}^n |\alpha_i| \, |||\, T(x_i)\, ||| \\
&\le \left(\max_{1 \le j \le n} |||\, T(x_j)\, ||| \right) \sum_{i=1}^n |\alpha_i| \\
&\le B \left(\max_{1 \le j \le n} |||\, T(x_j)\, ||| \right) \left\| \sum_{i=1}^n \alpha_i x_i \right\|.
\end{aligned}$$

That is, $|||\, T(x)\, ||| \le C\|x\|$, where $C = B \max\limits_{1 \le j \le n} |||\, T(x_j)\, |||$ (a constant depending only on T and the choice of basis for V). From our last result, T is continuous (bounded). $\square$

[Problems marked ($\triangleright$) are essential to a full understanding of the course; we will discuss most of these in class. Problems marked ($*$) are of general interest and are offered as a contribution to your personal growth. Unmarked problems are just for fun.]

The most important collection of functions for our purposes is the space $C[\,a, b\,]$, consisting of all continuous functions $f : [\,a, b\,] \to \mathbb{R}$. It's easy to see that $C[\,a, b\,]$ is a vector space under the usual pointwise operations on functions: $(f+g)(x) = f(x)+g(x)$ and $(\alpha f)(x) = \alpha f(x)$ for $\alpha \in \mathbb{R}$. Actually, we will be most interested in the finite-dimensional subspaces $\mathcal{P}_n$ of $C[\,a, b\,]$, consisting of all *algebraic polynomials* of degree at most n.

$\triangleright$ **1.** The subspace $\mathcal{P}_n$ has dimension exactly $n+1$. Why?

Another useful subset of $C[\,a, b\,]$ is the collection $\text{lip}_K\alpha$, consisting of all those f's which satisfy a *Lipschitz condition* of order $\alpha > 0$ with constant $0 < K < \infty$; i.e., those f's for which $|f(x) - f(y)| \le K\,|x-y|^\alpha$ for all x, y in $[\,a, b\,]$. [Some authors would say that f is *Hölder continuous* with exponent α.]

$*$ **2.** (a) Show that $\text{lip}_K\alpha$ is, indeed, a subset of $C[\,a, b\,]$.

(b) If $\alpha > 1$, show that $\text{lip}_K\alpha$ contains only the constant functions.

(c) Show that $\sqrt{x}$ is in $\text{lip}_1(1/2)$ and that $\sin x$ is in $\text{lip}_1 1$ on $[\,0, 1\,]$.

(d) Show that the collection $\text{lip}\,\alpha$, consisting of all those f's which are in $\text{lip}_K\alpha$ for some K, is a subspace of $C[\,a, b\,]$.

(e) Show that $\text{lip}\,1$ contains all the polynomials.

(f) If $f \in \text{lip}\,\alpha$ for some $\alpha > 0$, show that $f \in \text{lip}\,\beta$ for all $0 < \beta < \alpha$.

(g) Given $0 < \alpha < 1$, show that x^α is in $\text{lip}_1\alpha$ on $[\,0, 1\,]$ but *not* in $\text{lip}\,\beta$ for any $\beta > \alpha$.

We will also want to consider a norm on the vector space $C[\,a, b\,]$; we typically use the *uniform* or *sup norm* (Rivlin calls this the *Chebyshev norm*) defined by $\|f\| = \max\limits_{a\le x\le b} |f(x)|$. [Some authors write $\|f\|_u$ or $\|f\|_\infty$.]

$*$ **3.** Show that $\mathcal{P}_n$ and $\text{lip}_K\alpha$ are closed subsets of $C[\,a, b\,]$ (under the sup norm). Is $\text{lip}\,\alpha$ closed? A bit harder: Show that $\text{lip}\,1$ is both first category and dense in $C[\,a, b\,]$.

$\triangleright$ $*$ **4.** Fix n and consider the norm $\|p\|_1 = \sum_{k=0}^n |a_k|$ for $p(x) = a_0 + a_1 x + \cdots + a_n x^n \in \mathcal{P}_n$.

Show that there are constants $0 < A_n, B_n < \infty$ such that $A_n \|p\|_1 \le \|p\| \le B_n \|p\|_1$, where $\|p\| = \max_{a \le x \le b} |p(x)|$. Do A_n and B_n really depend on n?

We will occasionally consider spaces of real-valued functions defined on finite sets; that is, we will consider $\mathbb{R}^n$ under various norms. (Why is this the same?) We define a scale of norms on $\mathbb{R}^n$ by $\|x\|_p = \left(\sum_{i=1}^n |x_i|^p\right)^{1/p}$, where $x = (x_1, \ldots, x_n)$ and $1 \le p < \infty$ (we need $p \ge 1$ in order for this expression to be a legitimate norm, but the expression makes perfect sense for any $p > 0$, and even for $p < 0$ provided no x_i is 0). Notice, please, that the usual norm on $\mathbb{R}^n$ is given by $\|x\|_2$.

5. Show that $\lim_{p\to\infty} \|x\|_p = \max_{1\le i\le n} |x_i|$. For this reason we define $\|x\|_\infty = \max_{1\le i\le n} |x_i|$. Thus $\mathbb{R}^n$ under the norm $\|\cdot\|_\infty$ is the same as $C(\{1, 2, \ldots, n\})$ with its usual norm.

6. Assuming $x_i \ne 0$ for $i = 1, \ldots, n$, compute $\lim_{p\to 0+} \|x\|_p$ and $\lim_{p\to -\infty} \|x\|_p$.

7. Consider $\mathbb{R}^2$ under the norm $\|x\|_p$. Draw the graph of the unit sphere $\{x : \|x\|_p = 1\}$ for various values of p (especially $p = 1, 2, \infty$).

$*$ **8.** (Young's inequality): Let $1 < p < \infty$ and let q satisfy $\frac{1}{p} + \frac{1}{q} = 1$. Show that $ab \le \frac{1}{p}\, a^p + \frac{1}{q}\, b^q$ for all $a, b \ge 0$ with equality if and only if $a^p = b^q$.

$*$ **9.** (Hölder's inequality): Let $1 < p < \infty$ and let q satisfy $\frac{1}{p} + \frac{1}{q} = 1$. Show that

(a) $\sum_{i=1}^n |a_i\, b_i| \le \left(\sum_{i=1}^n |a_i|^p\right)^{1/p} \left(\sum_{i=1}^n |b_i|^q\right)^{1/q}$, and

(b) $\int_a^b |f(x)\, g(x)|\, dx \le \left(\int_a^b |f(x)|^p\, dx\right)^{1/p} \left(\int_a^b |g(x)|^q\, dx\right)^{1/q}$.

Describe the case for equality in each inequality. What happens if $p = 1$ ($q = \infty$)?

$*$ **10.** (Minkowski's inequality): For $1 \le p < \infty$, show that

(a) $\left(\sum_{i=1}^n |a_i + b_i|^p\right)^{1/p} \le \left(\sum_{i=1}^n |a_i|^p\right)^{1/p} + \left(\sum_{i=1}^n |b_i|^p\right)^{1/p}$ and that

(b) $\left(\int_a^b |f(x) + g(x)|^p\, dx\right)^{1/p} \le \left(\int_a^b |f(x)|^p\, dx\right)^{1/p} + \left(\int_a^b |g(x)|^p\, dx\right)^{1/p}$.

Describe the case for equality in each inequality. What happens if $p = \infty$?

Exercise **10** shows that $\|\cdot\|_p$ is indeed a norm for $1 \le p < \infty$. We write $L_p[\,a, b\,]$ to mean the vector space of functions on $[\,a, b\,]$ for which the integral norm is defined and finite, we write ℓ_p^n to mean the vector space of sequences of length n; that is, $\mathbb{R}^n$ supplied with the norm $\|\cdot\|_p$, and we write ℓ_p to mean the vector space of infinite sequences $x = (x_n)_{n=1}^{\infty}$ for which $\|x\|_p < \infty$. In each space, the usual algebraic operations are defined pointwise (or coordinatewise) and the norm is understood to be $\|\cdot\|_p$.

A normed space $(X, \|\cdot\|)$ is said to be *strictly convex* if $\|x + y\| = \|x\| + \|y\|$ always implies that x and y lie in the same direction; that is, either $x = \alpha y$ or $y = \alpha x$ for some nonnegative scalar α. Equivalently, X is strictly convex if the triangle inequality is always strict on nonparallel vectors.

11. Prove that the following are equivalent:

(a) $(X, \|\cdot\|)$ is strictly convex.

(b) If $x, y \in X$ are nonparallel, then $\left\|\dfrac{x+y}{2}\right\| < \dfrac{\|x\| + \|y\|}{2}$.

(c) If $x \neq y \in X$ with $\|x\| = 1 = \|y\|$, then $\left\|\dfrac{x+y}{2}\right\| < 1$.

12. Show that L_p and ℓ_p are strictly convex for $1 < p < \infty$. Show also that this fails in case $p = 1$. [Hint: This is actually a statement about the function $|t|^p$, $1 < p < \infty$.]

Strictly convex spaces are of interest when considering the problem of *nearest points*: Given a nonempty subset K of a normed space X and a point $x \notin K$, we ask whether there is a best approximation to x from elements of K; that is, we want to know if there exist one or more points $y_0 \in K$ satisfying

$$\|x - y_0\| = \inf_{y \in K} \|x - y\| = \operatorname{dist}(x, K).$$

It's not hard to see that a satisfactory answer to this question will require that we take K to be a *closed* set in X (for otherwise the points in $\overline{K} \setminus K$ wouldn't have nearest points). Less easy to see is that we typically also want to assume that K is a *convex* set. Recall that a subset K of a vector space X is said to be convex if it contains the line segment joining any pair of its points; that is, K is convex if

$$x, y \in K, \ 0 \le \lambda \le 1 \Longrightarrow \lambda x + (1 - \lambda) y \in K.$$

Obviously, any subspace of X is a convex set and, for our purposes at least, this is the most important example.

∗ **13.** Let X be a normed space and let $B = \{x \in X : \|x\| \le 1\}$. Show that B is a closed convex set.

14. Consider $\mathbb{R}^2$ under the norm $\|\cdot\|_\infty$. Let $B = \{y \in \mathbb{R}^2 : \|y\|_\infty \le 1\}$ and let $x = (2, 0)$. Show that there are infinitely many points in B nearest to x.

15. (a) Let $K = \{f \in L_1[\,0,1\,] : f \ge 0 \text{ and } \|f\|_1 = 1\}$. Show that K is a closed convex set in $L_1[\,0,1\,]$, that $0 \notin K$, and that every point in K is a nearest point to 0.

(b) Let $K = \{f \in C[\,0,1\,] : f(0) = 0 \text{ and } \int_0^1 f = 1\}$. Again, show that K is a closed convex set in $C[\,0,1\,]$, that $0 \notin K$, but that no point in K is nearest to 0.

16. Let K be a *compact* convex set in a strictly convex space X and let $x \in X$. Show that x has a unique nearest point $y_0 \in K$.

17. Let K be a closed subset of a complete normed space X. Prove that K is convex if and only if K is *midpoint convex*; that is, if and only if $(x + y)/2 \in K$ whenever x, $y \in K$. Is this result true in more general settings? For example, can you prove it without assuming completeness? Or, for that matter, is it true for arbitrary sets in any vector space (i.e., without even assuming the presence of a norm)?

Introduction

Let's begin with some notation. Throughout, we're concerned with the problem of best (uniform) approximation of a given function $f \in C[\,a,b\,]$ by elements from $\mathcal{P}_n$, the subspace of algebraic polynomials of degree at most n in $C[\,a,b\,]$. We know that the problem has a solution (possibly more than one), which we've chosen to write as p_n^*. We set

$$E_n(f) = \min_{p \in \mathcal{P}_n} \|f - p\| = \|f - p_n^*\|.$$

Since $\mathcal{P}_n \subset \mathcal{P}_{n+1}$ for each n, it's clear that $E_n(f) \geq E_{n+1}(f)$ for each n. Our goal in this chapter is to prove that $E_n(f) \to 0$. We'll accomplish this by proving:

Theorem. (The Weierstrass Approximation Theorem, 1885): *Let $f \in C[\,a,b\,]$. Then, for every $\varepsilon > 0$, there is a polynomial p such that $\|f - p\| < \varepsilon$.*

It follows from the Weierstrass theorem that $p_n^* \rightrightarrows f$ for each $f \in C[\,a,b\,]$. (Why?) This is an important first step in determining the exact nature of $E_n(f)$ as a function of f and n. We'll look for much more precise information in later sections.

Now there are many proofs of the Weierstrass theorem (a mere three are outlined in the exercises, but there are hundreds!), and all of them start with one simplification: The underlying interval $[\,a,b\,]$ is of no consequence.

Lemma. *If the Weierstrass theorem holds for $C[\,0,1\,]$, then it also holds for $C[\,a,b\,]$, and conversely. In fact, $C[\,0,1\,]$ and $C[\,a,b\,]$ are, for all practical purposes, identical: They are linearly isometric as normed spaces, order isomorphic as lattices, and isomorphic as algebras (rings).*

PROOF. We'll settle for proving only the first assertion; the second is outlined in the exercises (and uses a similar argument).

Given $f \in C[\,a,b\,]$, notice that the function

$$g(x) = f\big(a + (b-a)x\big), \quad 0 \leq x \leq 1,$$

defines an element of $C[\,0,1\,]$. Now, given $\varepsilon > 0$, suppose that we can find a polynomial p such that $\|g - p\| < \varepsilon$; in other words, suppose that

$$\max_{0\le x\le 1} \big| f\big(a + (b-a)x\big) - p(x)\big| < \varepsilon.$$

Then,

$$\max_{a\le t\le b} \left| f(t) - p\left(\frac{t-a}{b-a}\right)\right| < \varepsilon.$$

(Why?) But if $p(x)$ is a polynomial in x, then $q(t) = p\left(\frac{t-a}{b-a}\right)$ is a polynomial in t (again, why?) satisfying $\|f - q\| < \varepsilon$.

The proof of the converse is entirely similar: If $g(x)$ is an element of $C[\,0,1\,]$, then $f(t) = g\left(\frac{t-a}{b-a}\right)$, $a \le t \le b$, defines an element of $C[\,a,b\,]$. Moreover, if $q(t)$ is a polynomial in t approximating $f(t)$, then $p(x) = q(a + (b-a)x)$ is a polynomial in x approximating $g(x)$. The remaining details are left as an exercise. $\square$

The point to our first result is that it suffices to prove the Weierstrass theorem for any interval we like; $[\,0,1\,]$ and $[-1,1\,]$ are popular choices, but it hardly matters which interval we use.

Bernstein's Proof

The proof of the Weierstrass theorem we present here is due to the great Russian mathematician S. N. Bernstein in 1912. Bernstein's proof is of interest to us for a variety of reasons; perhaps most important is that Bernstein actually *displays* a sequence of polynomials that approximate a given $f \in C[\,0,1\,]$. Moreover, as we'll see later, Bernstein's proof generalizes to yield a powerful, unifying theorem, called the Bohman-Korovkin theorem.

If f is any *bounded* function on $[\,0,1\,]$, we define the sequence of *Bernstein polynomials* for f by

$$\big(B_n(f)\big)(x) = \sum_{k=0}^{n} f\left(\frac{k}{n}\right) \cdot \binom{n}{k} x^k (1-x)^{n-k}, \qquad 0 \le x \le 1.$$

Please note that $B_n(f)$ is a polynomial of degree at most n. Also, it's easy to see that $\big(B_n(f)\big)(0) = f(0)$, and $\big(B_n(f)\big)(1) = f(1)$. In general, $\big(B_n(f)\big)(x)$ is an *average* of the numbers $f(k/n)$, $k = 0, \ldots, n$. Bernstein's theorem states that $B_n(f) \rightrightarrows f$ for each $f \in C[\,0,1\,]$. Surprisingly, the proof actually only requires that we check three easy cases:

$$f_0(x) = 1, \quad f_1(x) = x, \quad \text{and} \quad f_2(x) = x^2.$$

This, and more, is the content of the following lemma.

Lemma. (i) $B_n(f_0) = f_0$ *and* $B_n(f_1) = f_1$.

(ii) $B_n(f_2) = \left(1 - \frac{1}{n}\right) f_2 + \frac{1}{n} f_1$, *and hence* $B_n(f_2) \rightrightarrows f_2$.

(iii) $\displaystyle\sum_{k=0}^{n} \left(\frac{k}{n} - x\right)^2 \binom{n}{k} x^k (1-x)^{n-k} = \frac{x(1-x)}{n} \le \frac{1}{4n}$, *if* $0 \le x \le 1$.

(iv) *Given* $\delta > 0$ *and* $0 \le x \le 1$, *let* F *denote the set of* k*'s in* $\{0, \ldots, n\}$ *for which* $\left|\frac{k}{n} - x\right| \ge \delta$. *Then* $\displaystyle\sum_{k \in F} \binom{n}{k} x^k (1-x)^{n-k} \le \frac{1}{4n\delta^2}$.

Proof. That $B_n(f_0) = f_0$ follows from the binomial formula:

$$\sum_{k=0}^{n} \binom{n}{k} x^k (1-x)^{n-k} = [x + (1-x)]^n = 1.$$

To see that $B_n(f_1) = f_1$, first notice that for $k \ge 1$ we have

$$\frac{k}{n}\binom{n}{k} = \frac{(n-1)!}{(k-1)!\,(n-k)!} = \binom{n-1}{k-1}.$$

Consequently,

$$\begin{aligned} \sum_{k=0}^{n} \frac{k}{n}\binom{n}{k} x^k (1-x)^{n-k} &= x \sum_{k=1}^{n} \binom{n-1}{k-1} x^{k-1} (1-x)^{n-k} \\ &= x \sum_{j=0}^{n-1} \binom{n-1}{j} x^j (1-x)^{(n-1)-j} = x. \end{aligned}$$

Next, to compute $B_n(f_2)$, we rewrite twice:

$$\begin{aligned} \left(\frac{k}{n}\right)^2 \binom{n}{k} = \frac{k}{n}\binom{n-1}{k-1} &= \frac{n-1}{n} \cdot \frac{k-1}{n-1}\binom{n-1}{k-1} + \frac{1}{n}\binom{n-1}{k-1}, \text{ if } k \ge 1 \\ &= \left(1 - \frac{1}{n}\right)\binom{n-2}{k-2} + \frac{1}{n}\binom{n-1}{k-1}, \text{ if } k \ge 2. \end{aligned}$$

Thus,

$$\begin{aligned}
\sum_{k=0}^{n} \left(\frac{k}{n}\right)^2 \binom{n}{k} x^k (1-x)^{n-k} \\
&= \left(1 - \frac{1}{n}\right) \sum_{k=2}^{n} \binom{n-2}{k-2} x^k (1-x)^{n-k} \;+\; \frac{1}{n} \sum_{k=1}^{n} \binom{n-1}{k-1} x^k (1-x)^{n-k} \\
&= \left(1 - \frac{1}{n}\right) x^2 + \frac{1}{n}\, x,
\end{aligned}$$

which establishes (ii) since $\|B_n(f_2) - f_2\| = \frac{1}{n}\|f_1 - f_2\| \to 0$ as $n \to \infty$.

To prove (iii) we combine the results in (i) and (ii) and simplify. Since $((k/n) - x)^2 = (k/n)^2 - 2x(k/n) + x^2$, we get

$$\begin{aligned}
\sum_{k=0}^{n} \left(\frac{k}{n} - x\right)^2 \binom{n}{k} x^k (1-x)^{n-k} \;&=\; \left(1 - \frac{1}{n}\right) x^2 + \frac{1}{n} x - 2x^2 + x^2 \\
&=\; \frac{1}{n}\, x(1-x) \;\le\; \frac{1}{4n},
\end{aligned}$$

for $0 \le x \le 1$.

Finally, to prove (iv), note that $1 \le ((k/n) - x)^2/\delta^2$ for $k \in F$, and hence

$$\begin{aligned}
\sum_{k \in F} \binom{n}{k} x^k (1-x)^{n-k} \;&\le\; \frac{1}{\delta^2} \sum_{k \in F} \left(\frac{k}{n} - x\right)^2 \binom{n}{k} x^k (1-x)^{n-k} \\
&\le\; \frac{1}{\delta^2} \sum_{k=0}^{n} \left(\frac{k}{n} - x\right)^2 \binom{n}{k} x^k (1-x)^{n-k} \\
&\le\; \frac{1}{4n\delta^2}, \quad \text{from (iii).} \quad \square
\end{aligned}$$

Now we're ready for *the proof of Bernstein's theorem*:

PROOF. Let $f \in C[0,1]$ and let $\varepsilon > 0$. Then, since f is uniformly continuous, there is a $\delta > 0$ such that $|f(x) - f(y)| < \varepsilon/2$ whenever $|x - y| < \delta$. Now we use the previous lemma to estimate $\|f - B_n(f)\|$. First notice that since the numbers $\binom{n}{k} x^k (1-x)^{n-k}$ are nonnegative and sum to 1, we have

$$\begin{aligned}
|f(x) - B_n(f)(x)| &= \left| f(x) - \sum_{k=0}^{n} \binom{n}{k} f\left(\frac{k}{n}\right) x^k(1-x)^{n-k} \right| \\
&= \left| \sum_{k=0}^{n} \left(f(x) - f\left(\frac{k}{n}\right)\right) \binom{n}{k} x^k(1-x)^{n-k} \right| \\
&\le \sum_{k=0}^{n} \left| f(x) - f\left(\frac{k}{n}\right)\right| \binom{n}{k} x^k(1-x)^{n-k},
\end{aligned}$$

Now fix n (to be specified in a moment) and let F denote the set of k's in $\{0, \ldots, n\}$ for which $|(k/n) - x| \ge \delta$. Then $|f(x) - f(k/n)| < \varepsilon/2$ for $k \notin F$, while $|f(x) - f(k/n)| \le 2\|f\|$ for $k \in F$. Thus,

$$\begin{aligned}
&\big| f(x) - \big(B_n(f)\big)(x) \big| \\
&\qquad \le \frac{\varepsilon}{2} \sum_{k \notin F} \binom{n}{k} x^k(1-x)^{n-k} + 2\|f\| \sum_{k \in F} \binom{n}{k} x^k(1-x)^{n-k} \\
&\qquad < \frac{\varepsilon}{2} \cdot 1 \;+\; 2\|f\| \cdot \frac{1}{4n\delta^2}, \quad \text{from (iv) of the Lemma,} \\
&\qquad < \varepsilon, \quad \text{provided that } n > \|f\|/\varepsilon\delta^2. \qquad \Box
\end{aligned}$$

Landau's Proof

Just because it's good for us, let's give *a second proof of Weierstrass's theorem*. This one is due to Landau in 1908. First, given $f \in C[\,0,1\,]$, notice that it suffices to approximate $f - p$, where p is any polynomial. (Why?) In particular, by subtracting the *linear* function $f(0) + x(f(1) - f(0))$, we may suppose that $f(0) = f(1) = 0$ and, hence, that $f \equiv 0$ outside $[\,0,1\,]$. That is, we may suppose that f is defined and uniformly continuous on all of $\mathbb{R}$.

Again we will display a sequence of polynomials that converge uniformly to f; this time we define

$$L_n(x) = c_n \int_{-1}^{1} f(x+t)\,(1-t^2)^n\,dt,$$

where c_n is chosen so that

$$c_n \int_{-1}^{1} (1-t^2)^n\,dt = 1.$$

Note that by our assumptions on f, we may rewrite this expression as

$$L_n(x) = c_n \int_{-x}^{1-x} f(x+t)\,(1-t^2)^n\,dt = c_n \int_{0}^{1} f(t)\,(1-(t-x)^2)^n\,dt.$$

Written this way, it's clear that L_n is a polynomial in x of degree at most n.

We first need to estimate c_n. An easy induction argument will convince you that $(1-t^2)^n \geq 1 - nt^2$, and so we get

$$\int_{-1}^{1} (1-t^2)^n \, dt \geq 2 \int_0^{1/\sqrt{n}} (1-nt^2)\, dt = \frac{4}{3\sqrt{n}} > \frac{1}{\sqrt{n}},$$

from which it follows that $c_n < \sqrt{n}$. In particular, for any $0 < \delta < 1$,

$$c_n \int_{\delta}^{1} (1-t^2)^n \, dt < \sqrt{n}\,(1-\delta^2)^n \to 0 \qquad (n \to \infty),$$

which is the inequality we'll need.

Next, let $\varepsilon > 0$ be given, and choose $0 < \delta < 1$ such that

$$|f(x) - f(y)| \leq \varepsilon/2 \text{ whenever } |x - y| \leq \delta.$$

Then, since $c_n(1-t^2)^n \geq 0$ and integrates to 1, we get

$$\begin{aligned} |L_n(x) - f(x)| &= \left| c_n \int_0^1 \big[f(x+t) - f(x) \big] (1-t^2)^n \, dt \right| \\ &\leq c_n \int_0^1 |f(x+t) - f(x)| (1-t^2)^n \, dt \\ &\leq \frac{\varepsilon}{2}\, c_n \int_0^{\delta} (1-t^2)^n \, dt + 2\|f\|\, c_n \int_{\delta}^{1} (1-t^2)^n \, dt \\ &\leq \frac{\varepsilon}{2} + 2\|f\|\,\sqrt{n}\,(1-\delta^2)^n < \varepsilon, \end{aligned}$$

provided that n is sufficiently large. $\square$

A third proof of the Weierstrass theorem, due to Lebesgue in 1898, is outlined in the exercises. Lebesgue's proof is of particular interest since it inspired Stone's version of the Weierstrass theorem; we'll discuss the Stone-Weierstrass theorem a bit later in the course.

Before we go on, let's stop and make an observation or two: While the Bernstein polynomials $B_n(f)$ offer a convenient and explicit polynomial approximation to f, they are by no means the best approximations. Indeed, recall that if $f_1(x) = x$ and $f_2(x) = x^2$, then $B_n(f_2) = (1 - \frac{1}{n})f_2 + \frac{1}{n} f_1 \neq f_2$. Clearly, the best approximation to f_2 out of $\mathcal{P}_n$ should be f_2 itself whenever $n \geq 2$. On the other hand, since we always have

$$E_n(f) \leq \|f - B_n(f)\| \qquad \text{(why?)},$$

a detailed understanding of Bernstein's proof will lend insight into the general problem of polynomial approximation. Our next project, then, is to improve upon our estimate of the error $\|f - B_n(f)\|$.

Improved Estimates

To begin, we will need a bit more notation. The *modulus of continuity* of a bounded function f on the interval $[\,a, b\,]$ is defined by

$$\omega_f(\delta) = \omega_f([\,a, b\,]; \delta) = \sup\big\{|f(x) - f(y)| : x, y \in [\,a, b\,],\ |x - y| \le \delta\big\}$$

for any $\delta > 0$. Note that $\omega_f(\delta)$ is a measure of the "ε" that goes along with δ (in the definition of uniform continuity); literally, we have written $\varepsilon = \omega_f(\delta)$ as a function of δ.

Here are a few easy facts about the modulus of continuity:

Exercises

1. We always have $|f(x) - f(y)| \le \omega_f(\,|x - y|\,)$ for any $x \ne y \in [\,a, b\,]$.
2. If $0 < \delta' \le \delta$, then $\omega_f(\delta') \le \omega_f(\delta)$.
3. f is *uniformly continuous* if and only if $\omega_f(\delta) \to 0$ as $\delta \to 0^+$.
4. If $f\,'$ exists and is bounded on $[\,a, b\,]$, then $\omega_f(\delta) \le K\delta$ for some constant K.
5. More generally, we say that f satisfies a *Lipschitz condition* of order α with constant K, where $0 < \alpha \le 1$ and $0 \le K < \infty$, if $|f(x) - f(y)| \le K|x - y|^\alpha$ for all x, y. We abbreviate this statement by the symbols: $f \in \text{lip}_K\alpha$. Check that if $f \in \text{lip}_K\alpha$, then $\omega_f(\delta) \le K\delta^\alpha$ for all $\delta > 0$.

For the time being, we actually only need one simple fact about $\omega_f(\delta)$:

Lemma. *Let f be a bounded function on $[\,a, b\,]$ and let $\delta > 0$. Then, $\omega_f(n\delta) \le n\,\omega_f(\delta)$ for $n = 1, 2, \ldots$. Consequently, $\omega_f(\lambda\delta) \le (1 + \lambda)\,\omega_f(\delta)$ for any $\lambda > 0$.*

Proof. Given $x < y$ with $|x - y| \le n\,\delta$, split the interval $[\,x, y\,]$ into n pieces, each of length at most δ. Specifically, if we set $z_k = x + k(y - x)/n$, for $k = 0, 1, \ldots, n$, then

$|z_k - z_{k-1}| \le \delta$ for any $k \ge 1$, and so

$$\begin{aligned} |f(x) - f(y)| &= \left| \sum_{k=1}^{n} f(z_k) - f(z_{k-1}) \right| \\ &\le \sum_{k=1}^{n} |f(z_k) - f(z_{k-1})| \\ &\le n\,\omega_f(\delta). \end{aligned}$$

Thus, $\omega_f(n\delta) \le n\,\omega_f(\delta)$.

The second assertion follows from the first (and one of our exercises). Given $\lambda > 0$, choose an integer n so that $n - 1 < \lambda \le n$. Then,

$$\omega_f(\lambda\delta) \le \omega_f(n\,\delta) \le n\,\omega_f(\delta) \le (1+\lambda)\,\omega_f(\delta). \qquad \square$$

We next repeat the proof of Bernstein's theorem, making a few minor adjustments here and there.

Theorem. *For any bounded function f on $[\,0,1\,]$ we have*

$$\|f - B_n(f)\| \le \frac{3}{2}\,\omega_f\left(\frac{1}{\sqrt{n}}\right).$$

In particular, if $f \in C[\,0,1\,]$, then $E_n(f) \le \frac{3}{2}\,\omega_f(\frac{1}{\sqrt{n}}) \to 0$ as $n \to \infty$.

Proof. We first do some term juggling:

$$\begin{aligned} |f(x) - B_n(f)(x)| &= \left| \sum_{k=0}^{n} \left(f(x) - f\left(\frac{k}{n}\right) \right) \binom{n}{k} x^k (1-x)^{n-k} \right| \\ &\le \sum_{k=0}^{n} \left| f(x) - f\left(\frac{k}{n}\right) \right| \binom{n}{k} x^k (1-x)^{n-k} \\ &\le \sum_{k=0}^{n} \omega_f\left(\left| x - \frac{k}{n} \right| \right) \binom{n}{k} x^k (1-x)^{n-k} \\ &\le \omega_f\left(\frac{1}{\sqrt{n}}\right) \sum_{k=0}^{n} \left[1 + \sqrt{n} \left| x - \frac{k}{n} \right| \right] \binom{n}{k} x^k (1-x)^{n-k} \\ &= \omega_f\left(\frac{1}{\sqrt{n}}\right) \left[1 \,+\, \sqrt{n} \sum_{k=0}^{n} \left| x - \frac{k}{n} \right| \binom{n}{k} x^k (1-x)^{n-k} \right], \end{aligned}$$

where the third inequality follows from our previous Lemma (by taking $\lambda = \sqrt{n}\,\left|x - \frac{k}{n}\right|$ and $\delta = \frac{1}{\sqrt{n}}$). All that remains is to estimate the sum, and for this we'll use Cauchy-Schwarz (and our earlier observations about Bernstein polynomials). Since each of the terms $\binom{n}{k} x^k (1-x)^{n-k}$ is nonnegative, we have

$$\begin{aligned}
&\sum_{k=0}^{n} \left| x - \frac{k}{n} \right| \binom{n}{k} x^k (1-x)^{n-k} \\
&\qquad \le \left[\sum_{k=0}^{n} \left| x - \frac{k}{n} \right|^2 \binom{n}{k} x^k (1-x)^{n-k} \right]^{1/2} \cdot \left[\sum_{k=0}^{n} \binom{n}{k} x^k (1-x)^{n-k} \right]^{1/2} \\
&\qquad \le \left[\frac{1}{4n} \right]^{1/2} = \frac{1}{2\sqrt{n}}.
\end{aligned}$$

Finally,

$$|f(x) - B_n(f)(x)| \le \omega_f\!\left(\frac{1}{\sqrt{n}}\right) \left[1 \;+\; \sqrt{n} \cdot \frac{1}{2\sqrt{n}} \right] = \frac{3}{2}\,\omega_f\!\left(\frac{1}{\sqrt{n}}\right). \qquad \square$$

Examples

1. If $f \in \mathrm{lip}_K\alpha$, it follows that $\|f - B_n(f)\| \le \frac{3}{2} K n^{-\alpha/2}$ and hence $E_n(f) \le \frac{3}{2} K n^{-\alpha/2}$.

2. As a particular case of the first example, consider $f(x) = \left|x - \frac{1}{2}\right|$ on $[0, 1]$. Then $f \in \mathrm{lip}_1 1$, and so $\|f - B_n(f)\| \le \frac{3}{2}\, n^{-1/2}$. But, as Rivlin points out (see Remark 3 on p. 16 of his book), $\|f - B_n(f)\| > \frac{1}{2}\, n^{-1/2}$. Thus, we can't hope to improve on the power of n in this estimate. Nevertheless, we will see an improvement in our estimate of $E_n(f)$.

The Bohman-Korovkin Theorem

The real value to us in Bernstein's approach is that the map $f \mapsto B_n(f)$, while providing a simple formula for an approximating polynomial, is also *linear* and *positive*. In other words,

$$B_n(f + g) = B_n(f) + B_n(g),$$

$$B_n(\alpha f) = \alpha B_n(f), \quad \alpha \in \mathbb{R},$$

and

$$B_n(f) \ge 0 \quad \text{whenever} \quad f \ge 0.$$

As it happens, any positive, linear map $T : C[\,0, 1\,] \to C[\,0, 1\,]$ is necessarily also continuous!

Lemma. *If $T : C[\,a,b\,] \to C[\,a,b\,]$ is both positive and linear, then T is continuous.*

PROOF. First note that a positive, linear map is also *monotone*. That is, T satisfies $T(f) \le T(g)$ whenever $f \le g$. (Why?) Thus, for any $f \in C[\,a,b\,]$, we have

$$-f,\ f \le |f| \implies -T(f),\ T(f) \le T(|f|);$$

that is, $|T(f)| \le T(|f|)$. But now $|f| \le \|f\| \cdot \mathbf{1}$, where $\mathbf{1}$ denotes the constant 1 function, and so we get

$$|T(f)| \le T(|f|) \le \|f\|\, T(\mathbf{1}).$$

Thus,

$$\|T(f)\| \le \|f\|\, \|T(\mathbf{1})\|$$

for any $f \in C[\,a,b\,]$. Finally, since T is linear, it follows that T is Lipschitz with constant $\|T(\mathbf{1})\|$:

$$\|T(f) - T(g)\| = \|T(f-g)\| \le \|T(\mathbf{1})\|\, \|f-g\|.$$

Consequently, T is continuous. $\square$

Now positive, linear maps abound in analysis, so this is a fortunate turn of events. What's more, Bernstein's theorem generalizes very nicely when placed in this new setting. The following elegant theorem was proved (independently) by Bohman and Korovkin in, roughly, 1952.

Theorem. *Let $T_n : C[\,0,1\,] \to C[\,0,1\,]$ be a sequence of positive, linear maps, and suppose that $T_n(f) \to f$ uniformly in each of the three cases*

$$f_0(x) = 1, \quad f_1(x) = x, \quad \text{and} \quad f_2(x) = x^2.$$

Then, $T_n(f) \to f$ uniformly for every $f \in C[\,0,1\,]$.

The proof of the Bohman-Korovkin theorem is essentially identical to the proof of Bernstein's theorem except, of course, we write $T_n(f)$ in place of $B_n(f)$. For full details, see Cheney's book *An Introduction to Approximation Theory*, Chelsea, 1982. Rather than proving the theorem, let's settle for a quick application.

Example

Let $f \in C[\,0,1\,]$ and, for each n, let $L_n(f)$ be the "polygonal" approximation to f with nodes at k/n, $k = 0, 1, \ldots, n$. That is, $L_n(f)$ is linear on each subinterval $[\,(k-1)/n,\, k/n\,]$ and agrees with f at each of the endpoints $L_n(f)(k/n) = f(k/n)$. Then, $L_n(f) \to f$ uniformly for each $f \in C[\,0,1\,]$. This is actually an easy calculation all by itself, but let's see why the Bohman-Korovkin theorem makes short work of it.

That $L_n(f)$ is positive and linear is (nearly) obvious; that $L_n(f_0) = f_0$ and $L_n(f_1) = f_1$ are really easy since, in fact, $L_n(f) = f$ for any linear function f. We just need to show that $L_n(f_2) \rightrightarrows f_2$. But a picture will convince you that the maximum distance between $L_n(f_2)$ and f_2 on the interval $[\,(k-1)/n,\, k/n\,]$ is at most

$$\left(\frac{k}{n}\right)^2 - \left(\frac{k-1}{n}\right)^2 = \frac{2k-1}{n^2} \le \frac{2}{n}.$$

That is, $\|f_2 - L_n(f_2)\| \le 2/n \to 0$ as $n \to \infty$.

[Note that L_n is a linear projection from $C[\,0,1\,]$ onto the subspace of polygonal functions based on the nodes k/n, $k = 0, \ldots, n$. An easy calculation, similar in spirit to the example above, will show that $\|f - L_n(f)\| \le 2\,\omega_f(1/n) \to 0$ as $n \to \infty$ for any $f \in C[\,0,1\,]$.]

One of our first tasks will be to give a constructive proof of *Weierstrass's Theorem*, stating that each $f \in C[\,a,b\,]$ is the uniform limit of a sequence of polynomials. As it happens, the choice of interval $[\,a,b\,]$ is inconsequential: If Weierstrass's theorem is true for one, then it's true for all.

$\triangleright$ $*$ **18.** Define $\sigma : [\,0,1\,] \to [\,a,b\,]$ by $\sigma(t) = a + t(b-a)$ for $0 \le t \le 1$, and define a transformation $T_\sigma : C[\,a,b\,] \to C[\,0,1\,]$ by $(T_\sigma(f))(t) = f(\sigma(t))$. Prove that T_σ satisfies:

(a) $T_\sigma(f+g) = T_\sigma(f) + T_\sigma(g)$ and $T_\sigma(cf) = c\,T_\sigma(f)$ for $c \in \mathbb{R}$.

(b) $T_\sigma(fg) = T_\sigma(f)\,T_\sigma(g)$. In particular, T_σ maps polynomials to polynomials.

(c) $T_\sigma(f) \le T_\sigma(g)$ if and only if $f \le g$.

(d) $\|T_\sigma(f)\| = \|f\|$.

(e) T_σ is both one-to-one and onto. Moreover, $(T_\sigma)^{-1} = T_{\sigma^{-1}}$.

The point to exercise **18** is that $C[\,a,b\,]$ and $C[\,0,1\,]$ are *identical* as vector spaces, metric spaces, algebras, and lattices. For all practical purposes, they are one and the same space. While Bernstein's proof of the Weierstrass theorem (below) will prove most useful for our purposes, there are *many* others; two of these (in the case of $C[\,0,1\,]$) are sketched below.

$*$ **19.** (Landau's proof): For each $n = 1, 2, \ldots$ and $0 \le \delta \le 1$, define $I_n(\delta) = \int_\delta^1 (1-x^2)^n\,dx$. Show that $I_n(\delta)/I_n(0) \to 0$ as $n \to \infty$ for any $\delta > 0$. Now, given $f \in C[\,0,1\,]$ with $f(0) = f(1) = 0$, show that the *polynomial* $L_n(x) = (2I_n(0))^{-1}\int_0^1 f(t)(1-(t-x)^2)^n\,dt$ converges uniformly to $f(x)$ on $[\,0,1\,]$ as $n \to \infty$. [Hint: You may assume that $f \equiv 0$ outside of $[\,0,1\,]$.] To get the result for general $f \in C[\,0,1\,]$, we simply need to subtract the *linear* function $f(0) + x(f(1) - f(0))$.

$*$ **20.** (Lebesgue's proof): Given $f \in C[\,0,1\,]$, first show that f can be uniformly approximated by a *polygonal* function. Specifically, given a positive integer N, define $L(x)$ by the conditions $L(k/N) = f(k/N)$ for $k = 0, 1, \ldots, N$, and $L(x)$ is linear for $k/N \le x \le (k+1)/N$; show that $\|f - L\|$ is small provided that N is sufficiently large. The function $L(x)$ can be written (uniquely) as a linear combination of the "angles" $\varphi_k(x) = |x - k/N| + x - k/N$ and $\varphi_N(x) = 1$; the equation $L(x) = \sum_{k=0}^N c_k\varphi_k(x)$ can be solved since the system of equations $L(k/N) = \sum_{k=0}^N c_k\varphi_k(k/N)$, $k = 0, \ldots, N$, can be solved (uniquely) for $c_0, \ldots, c_N$. (How?) To finish the proof, we need to show that $|x|$ can be approximated by polynomials on any interval $[\,a, b\,]$. (Why?)

$*$ **21.** Here's an elementary proof that there is a sequence of polynomials (P_n) converging uniformly to $|x|$ on $[-1, 1\,]$.

(a) Define (P_n) recursively by $P_{n+1}(x) = P_n(x) + [x - P_n(x)^2]/2$, where $P_0(x) = 0$. Clearly, each P_n is a polynomial.

(b) Check that $0 \le P_n(x) \le P_{n+1}(x) \le \sqrt{x}$ for $0 \le x \le 1$. Use Dini's theorem to conclude that $P_n(x) \rightrightarrows \sqrt{x}$ on $[\,0, 1\,]$.

(c) $P_n(x^2)$ is also a polynomial, and $P_n(x^2) \rightrightarrows |x|$ on $[-1, 1\,]$.

$\triangleright *$ **22.** The result in problem **19** (or **20**) shows that the polynomials are dense in $C[\,0, 1\,]$. Using the results in **18**, conclude that the polynomials are also dense in $C[\,a, b\,]$.

$\triangleright *$ **23.** How do we know that there are non-polynomial elements in $C[\,0, 1\,]$? In other words, is it possible that every element of $C[\,0, 1\,]$ agrees with some polynomial on $[\,0, 1\,]$?

24. Let (Q_n) be a sequence of polynomials of degree m_n, and suppose that (Q_n) converges uniformly to f on $[\,a, b\,]$, where f is *not* a polynomial. Show that $m_n \to \infty$.

25. If $f \in C[-1, 1\,]$ (or $C^{2\pi}$) is an *even* function, show that f may be uniformly approximated by even polynomials (or even trig polynomials).

26. If $f \in C[\,0, 1\,]$ and if $f(0) = f(1) = 0$, show that the sequence of polynomials $\sum_{k=0}^{n} \left[\binom{n}{k} f(k/n)\right] x^k (1-x)^{n-k}$ with *integer* coefficients converges uniformly to f (where $[x]$ denotes the greatest integer in x). The same trick works for any $f \in C[\,a, b\,]$ provided that $0 < a < b < 1$.

27. If p is a polynomial and $\varepsilon > 0$, prove that there is a polynomial q with *rational* coefficients such that $\|p - q\| < \varepsilon$ on $[\,0, 1\,]$. Conclude that $C[\,0, 1\,]$ is *separable*.

28. Let (x_i) be a sequence of numbers in $(0, 1)$ such that $\lim\limits_{n\to\infty} \frac{1}{n} \sum_{i=1}^{n} x_i^k$ exists for every $k = 0, 1, 2, \ldots$. Show that $\lim\limits_{n\to\infty} \frac{1}{n} \sum_{i=1}^{n} f(x_i)$ exists for every $f \in C[\,0, 1\,]$.

29. If $f \in C[\,0, 1\,]$ and if $\int_0^1 x^n f(x)\,dx = 0$ for each $n = 0, 1, 2, \ldots$, show that $f \equiv 0$. [Hint: Using the Weierstrass theorem, show that $\int_0^1 f^2 = 0$.]

The next proof of the Weierstrass theorem that we consider is quite explicit; we actually display a sequence of polynomials that converges uniformly to a given $f \in C[\,0,1\,]$. Given $f \in C[\,0,1\,]$, we define the sequence $\big(B_n(f)\big)_{n=1}^{\infty}$ of *Bernstein polynomials* for f by

$$\big(B_n(f)\big)(x) = \sum_{k=0}^{n} f\left(\frac{k}{n}\right) \cdot \binom{n}{k} x^k (1-x)^{n-k}.$$

Please note that $B_n(f)$ is a polynomial of degree at most n. Also, it's easy to see that $\big(B_n(f)\big)(0) = f(0)$ and $\big(B_n(f)\big)(1) = f(1)$. In general, $\big(B_n(f)\big)(x)$ is an *average* of the numbers $f(k/n)$, $k = 0, \ldots, n$. *Bernstein's theorem* states that the sequence $B_n(f)$ converges uniformly to f for each $f \in C[\,0,1\,]$; the proof is rather simple once we have a few facts about the Bernstein polynomials at our disposal. For later reference, let's write

$$f_0(x) = 1, \quad f_1(x) = x, \quad \text{and} \quad f_2(x) = x^2.$$

Among other things, the following exercise establishes Bernstein's theorem for these three polynomials. Curiously, these few special cases will imply the general result.

▷ **30.** (i) $B_n(f_0) = f_0$ and $B_n(f_1) = f_1$. [Hint: Use the binomial theorem.]

(ii) $B_n(f_2) = \left(1 - \frac{1}{n}\right) f_2 + \frac{1}{n} f_1$, and hence $(B_n(f_2))$ converges uniformly to f_2.

(iii) $\sum_{k=0}^{n} \left(\frac{k}{n} - x\right)^2 \binom{n}{k} x^k (1-x)^{n-k} = \frac{x(1-x)}{n} \le \frac{1}{4n}$, if $0 \le x \le 1$.

(iv) Given $\delta > 0$ and $0 \le x \le 1$, let F denote the set of k's in $\{0, \ldots, n\}$ for which $\left|\frac{k}{n} - x\right| \ge \delta$. Then $\sum_{k \in F} \binom{n}{k} x^k (1-x)^{n-k} \le \frac{1}{4n\delta^2}$.

▷ **31.** Show that $|B_n(f)| \le B_n(|f|)$, and that $B_n(f) \ge 0$ whenever $f \ge 0$. Conclude that $\|B_n(f)\| \le \|f\|$.

32. If f is a bounded function on $[\,0,1\,]$, show that $B_n(f)(x) \to f(x)$ at each point of continuity of f.

33. (Bohman, Korovkin) Let (T_n) be a sequence of monotone linear operators on $C[\,0,1\,]$; that is, each T_n is a linear map from $C[\,0,1\,]$ into itself satisfying $T_n(f) \le T_n(g)$ whenever $f \le g$. Suppose also that $T_n(f_0) \rightrightarrows f_0$, $T_n(f_1) \rightrightarrows f_1$, and $T_n(f_2) \rightrightarrows f_2$. Prove that $T_n(f) \rightrightarrows f$ for every $f \in C[\,0,1\,]$. [Hint: Mimic the proof of Bernstein's theorem.]

34. Find $B_n(f)$ for $f(x) = x^3$. [Hint: $k^2 = (k-1)(k-2) + 3(k-1) + 1$.] The same method of calculation can be used to show that $B_n(f) \in \mathcal{P}_m$ whenever $f \in \mathcal{P}_m$ and $n > m$.

∗ **35.** Let f be continuously differentiable on $[a,b]$, and let $\varepsilon > 0$. Show that there is a polynomial p such that $\|f - p\| < \varepsilon$ *and* $\|f' - p'\| < \varepsilon$.

36. Suppose that $f \in C[a,b]$ is twice continuously differentiable and has $f'' > 0$. Prove that the best linear approximation to f on $[a,b]$ is $a_0 + a_1 x$ where $a_0 = f'(c)$, $a_1 = [f(a) + f(c) + f'(c)(a+c)]/2$, and where c is the unique solution to $f'(c) = (f(b) - f(a))/(b-a)$.

The next several exercises concern the *modulus of continuity*. Given a bounded real-valued function f defined on some interval I, we define ω_f, the modulus of continuity of f, by

$$\omega_f(I;\delta) = \omega_f(\delta) = \sup\big\{|f(x) - f(y)| : x, y \in I,\ |x - y| \le \delta\big\}, \quad \delta \ge 0.$$

Note, for example, that if f is *uniformly continuous*, then $\omega_f(\delta) \to 0$ as $\delta \to 0$. Indeed, the statement that $|f(x) - f(y)| \le \varepsilon$ whenever $|x - y| \le \delta$ is equivalent to the statement that $\omega_f(\delta) \le \varepsilon$. On the other hand, if the graph of f has a jump of magnitude 1, say, then $\omega_f(\delta) \ge 1$ for all $\delta > 0$.

▷ **37.** If f satisfies the Lipschitz condition $|f(x) - f(y)| \le K|x - y|$, what can you say about ω_f? Calculate ω_g for $g(x) = \sqrt{x}$.

38. If $f \in C[a,b]$, show that $\omega_f(\delta_1 + \delta_2) \le \omega_f(\delta_1) + \omega_f(\delta_2)$ and that $\omega_f(\delta) \downarrow 0$ as $\delta \downarrow 0$. Use this to show that ω_f is continuous for $\delta \ge 0$. Finally, show that the modulus of continuity of ω_f is again ω_f.

▷ **39.** (a) If $x = \cos\theta$, where $-1 \le x \le 1$, and if $g(\theta) = f(\cos\theta)$, show that $\omega_g([-\pi,\pi],\delta) = \omega_g([0,\pi],\delta) \le \omega_f([-1,1];\delta)$.

(b) If $g(x) = f(ax+b)$ for $c \le x \le d$, show that $\omega_g([c,d];\delta) = \omega_f([ac+b, ad+b]; a\delta)$.

40. Let f be continuously differentiable on $[0,1]$. Show that $(B_n(f)')$ converges uniformly to f' by showing that $\|B_n(f') - (B_{n+1}(f))'\| \le \omega_{f'}(1/(n+1))$. [In order to see why this is of interest, find a uniformly convergent sequence of *polynomials* whose derivatives fail to converge uniformly. Compare this result with problem **35**.]

Introduction

A (real) *trigonometric polynomial*, or trig polynomial for short, is a function of the form

$$a_0 + \sum_{k=1}^{n} (\, a_k \cos kx + b_k \sin kx \,), \tag{$*$}$$

where $a_0, \ldots, a_n$ and $b_1, \ldots, b_n$ are real numbers. The *degree* of a trig polynomial is the highest frequency occurring in any representation of the form $(*)$; thus, $(*)$ has degree n provided that one of a_n or b_n is nonzero. We will use $\mathcal{T}_n$ to denote the collection of trig polynomials of degree at most n, and $\mathcal{T}$ to denote the collection of all trig polynomials (i.e., the union of the $\mathcal{T}_n$'s).

It is convenient to take the space of all continuous 2π-periodic functions on $\mathbb{R}$ as the containing space for $\mathcal{T}_n$; a space we denote by $C^{2\pi}$. The space $C^{2\pi}$ has several equivalent descriptions. For one, it's obvious that $C^{2\pi}$ is a subspace of $C(\mathbb{R})$, the space of all continuous functions on $\mathbb{R}$. But we might also consider $C^{2\pi}$ as a subspace of $C[\,0, 2\pi\,]$ in the following way: The 2π-periodic continuous functions on $\mathbb{R}$ may be identified with the set of functions $f \in C[\,0, 2\pi\,]$ satisfying $f(0) = f(2\pi)$. Each such f extends to a 2π-periodic element of $C(\mathbb{R})$ in an obvious way, and it's not hard to see that the condition $f(0) = f(2\pi)$ defines a subspace of $C[\,0, 2\pi\,]$. As a third description, it is often convenient to identify $C^{2\pi}$ with the collection $C(\mathbb{T})$, consisting of all the continuous real-valued functions on $\mathbb{T}$, where $\mathbb{T}$ is the unit circle in the complex plane $\mathbb{C}$. That is, we simply make the identifications

$$\theta \longleftrightarrow e^{i\theta} \qquad \text{and} \qquad f(\theta) \longleftrightarrow f(e^{i\theta}).$$

In any case, each $f \in C^{2\pi}$ is uniformly continuous and uniformly bounded on all of $\mathbb{R}$, and is completely determined by its values on any interval of length 2π. In particular, we may (and will) endow $C^{2\pi}$ with the sup norm:

$$\|f\| = \max_{0 \le x \le 2\pi} |f(x)| = \max_{x \in \mathbb{R}} |f(x)|.$$

Our goal in this chapter is to prove what is sometimes called *Weierstrass's second theorem* (also from 1885).

Theorem. (Weierstrass's Second Theorem, 1885) *Let $f \in C^{2\pi}$. Then, for every $\varepsilon > 0$, there exists a trig polynomial T such that $\|f - T\| < \varepsilon$.*

Ultimately, we will give several different proofs of this theorem. Weierstrass gave a separate proof of this result in the same paper containing his theorem on approximation by algebraic polynomials, but it was later pointed out by Lebesgue (1898) that the two theorems are, in fact, equivalent. Lebesgue's proof is based on several elementary observations. We will outline these elementary facts as "exercises with hints," supplying a few proofs here and there, but leaving full details to the reader.

We first justify the use of the word "polynomial" in describing $(*)$.

Lemma. *$\cos nx$ and $\sin(n+1)x/\sin x$ can be written as polynomials of degree exactly n in $\cos x$ for any integer $n \geq 0$.*

Proof. Using the recurrence formula $\cos kx + \cos(k-2)x = 2\cos(k-1)x\cos x$ it's not hard to see that $\cos 2x = 2\cos^2 x - 1$, $\cos 3x = 4\cos^3 x - 3\cos x$, and $\cos 4x = 8\cos^4 x - 8\cos^2 x + 1$. More generally, by induction, $\cos nx$ is a polynomial of degree n in $\cos x$ with leading coefficient 2^{n-1}. Using this fact and the identity $\sin(k+1)x - \sin(k-1)x = 2\cos kx \sin x$ (along with another easy induction argument), it follows that $\sin(n+1)x$ can be written as $\sin x$ times a polynomial of degree n in $\cos x$ with leading coefficient 2^n.

Alternatively, notice that by writing $(i\sin x)^{2k} = (\cos^2 x - 1)^k$ we have

$$\begin{aligned}\cos nx &= \operatorname{Re}\left[(\cos x + i\sin x)^n\right] = \operatorname{Re}\left[\sum_{k=0}^{n}\binom{n}{k}(i\sin x)^k\cos^{n-k} x\right]\\ &= \sum_{k=0}^{[n/2]}\binom{n}{2k}(\cos^2 x - 1)^k\cos^{n-2k} x.\end{aligned}$$

The coefficient of $\cos^n x$ in this expansion is then

$$\sum_{k=0}^{[n/2]}\binom{n}{2k} = \frac{1}{2}\sum_{k=0}^{n}\binom{n}{k} = 2^{n-1}.$$

(All the binomial coefficients together sum to $(1+1)^n = 2^n$, but the even or odd terms taken separately sum to exactly half this amount since $(1+(-1))^n = 0$.)

Similarly,

$$\sin(n+1)x = \operatorname{Im}\left[(\cos x + i\sin x)^{n+1}\right] = \operatorname{Im}\left[\sum_{k=0}^{n+1}\binom{n+1}{k}(i\sin x)^k\cos^{n+1-k}x\right]$$

$$= \sum_{k=0}^{[n/2]}\binom{n+1}{2k+1}(\cos^2 x - 1)^k \cos^{n-2k} x\sin x,$$

where we've written $(i\sin x)^{2k+1} = i(\cos^2 x - 1)^k \sin x$. The coefficient of $\cos^n x\sin x$ is

$$\sum_{k=0}^{[n/2]}\binom{n+1}{2k+1} = \frac{1}{2}\sum_{k=0}^{n+1}\binom{n+1}{k} = 2^n. \qquad \square$$

Corollary. *Any trig polynomial* $(*)$ *may be written as* $P(\cos x) + Q(\cos x)\sin x$*, where* P *and* Q *are algebraic polynomials of degree at most* n *and* $n-1$*, respectively. If* $(*)$ *represents an even function, then it can be written using only cosines.*

Corollary. *The collection* $\mathcal{T}$*, consisting of all trig polynomials, is both a subspace and a subring of* $C^{2\pi}$ *(that is,* $\mathcal{T}$ *is closed under both linear combinations and products). In other words,* $\mathcal{T}$ *is a subalgebra of* $C^{2\pi}$*.*

It's not hard to see that the procedure we've described above can be reversed; that is, each algebraic polynomial in $\cos x$ and $\sin x$ can be written in the form $(*)$. For example, $4\cos^3 x = 3\cos x + \cos 3x$. But, rather than duplicate our efforts, let's use a bit of linear algebra. First, the $2n+1$ functions

$$\mathcal{A} = \{\, 1,\ \cos x,\ \cos 2x,\ \ldots,\ \cos nx,\ \sin x,\ \sin 2x,\ \ldots,\ \sin nx \,\},$$

are *linearly independent*; the easiest way to see this is to notice that we may define an *inner product* on $C^{2\pi}$ under which these functions are *orthogonal*. Specifically,

$$\langle f, g\rangle = \int_0^{2\pi} f(x)\,g(x)\,dx = 0, \qquad \langle f, f\rangle = \int_0^{2\pi} f(x)^2\,dx \neq 0$$

for any pair of functions $f \neq g \in \mathcal{A}$. (We'll pursue this direction in greater detail later in the course.) Second, we've shown that each element of $\mathcal{A}$ lives in the space spanned by the $2n+1$ functions

$$\mathcal{B} = \{\, 1,\ \cos x,\ \cos^2 x,\ \ldots,\ \cos^n x,\ \sin x,\ \cos x\sin x,\ \ldots,\ \cos^{n-1} x\sin x \,\}.$$

That is,

$$\mathcal{T}_n \ \equiv \ \operatorname{span}\mathcal{A} \ \subset \ \operatorname{span}\mathcal{B}.$$

By comparing dimensions, we have

$$2n+1 = \dim \mathcal{T}_n = \dim(\operatorname{span}\mathcal{A}) \le \dim(\operatorname{span}\mathcal{B}) \le 2n+1,$$

and hence we must have $\operatorname{span}\mathcal{A} = \operatorname{span}\mathcal{B}$. The point here is that $\mathcal{T}_n$ is a finite-dimensional subspace of $C^{2\pi}$ of dimension $2n+1$, and we may use either one of these sets of functions as a basis for $\mathcal{T}_n$.

Before we leave these issues behind, let's summarize the situation for *complex* trig polynomials; i.e., the case where we allow complex coefficients in $(*)$. Now it's clear that every trig polynomial $(*)$, whether real or complex, can be written as

$$\sum_{k=-n}^{n} c_k e^{ikx}, \tag{$**$}$$

where the c_k's are *complex*; that is, a trig polynomial is actually a polynomial (over $\mathbb{C}$) in $z = e^{ix}$ and $\bar{z} = e^{-ix}$. Conversely, every polynomial $(**)$ can be written in the form $(*)$, using complex a_k's and b_k's. Thus, the complex trig polynomials of degree n form a vector space of dimension $2n+1$ over $\mathbb{C}$ (hence of dimension $2(2n+1)$ when considered as a vector space over $\mathbb{R}$). But, not every polynomial in z and $\bar{z}$ represents a *real* trig polynomial. Rather, the real trig polynomials are the *real parts* of the complex trig polynomials. To see this, notice that $(**)$ represents a real-valued function if and only if

$$\sum_{k=-n}^{n} c_k e^{ikx} = \overline{\sum_{k=-n}^{n} c_k e^{ikx}} = \sum_{k=-n}^{n} \bar{c}_{-k} e^{ikx};$$

that is, $c_k = \bar{c}_{-k}$ for each k. In particular, c_0 must be real, and hence

$$\begin{aligned}
\sum_{k=-n}^{n} c_k e^{ikx} &= c_0 + \sum_{k=1}^{n} (c_k e^{ikx} + c_{-k} e^{-ikx}) \\
&= c_0 + \sum_{k=1}^{n} (c_k e^{ikx} + \bar{c}_k e^{-ikx}) \\
&= c_0 + \sum_{k=1}^{n} \big[\, (c_k + \bar{c}_k) \cos kx + i(c_k - \bar{c}_k) \sin kx \,\big] \\
&= c_0 + \sum_{k=1}^{n} \big[\, 2\mathrm{Re}(c_k) \cos kx - 2\mathrm{Im}(c_k) \sin kx \,\big],
\end{aligned}$$

which is of the form $(*)$ with a_k and b_k real.

Conversely, given any real trig polynomial $(*)$, we have

$$a_0 + \sum_{k=1}^{n} \big(\, a_k \cos kx + b_k \sin kx \,\big) = a_0 + \sum_{k=1}^{n} \left[\left(\frac{a_k - ib_k}{2} \right) e^{ikx} + \left(\frac{a_k + ib_k}{2} \right) e^{-ikx} \right],$$

which of of the form $(**)$ with $c_k = \bar{c}_{-k}$ for each k.

It's time we returned to approximation theory! Since we've been able to identify $C^{2\pi}$ with a subspace of $C[\,0, 2\pi\,]$, and since $\mathcal{T}_n$ is a finite-dimensional subspace of $C^{2\pi}$, we have

Corollary. *Each $f \in C^{2\pi}$ has a best approximation (on all of $\mathbb{R}$) out of $\mathcal{T}_n$. If f is an even function, then it has a best approximation which is also even.*

Proof. We only need to prove the second claim, so suppose that $f \in C^{2\pi}$ is even and that $T^* \in \mathcal{T}_n$ satisfies

$$\|f - T^*\| = \min_{T \in \mathcal{T}_n} \|f - T\|.$$

Then, since f is even, $\widetilde{T}(x) = T^*(-x)$ is also a best approximation to f out of $\mathcal{T}_n$; indeed,

$$\begin{aligned} \|f - \widetilde{T}\,\| &= \max_{x \in \mathbb{R}} |f(x) - T^*(-x)| \\ &= \max_{x \in \mathbb{R}} |f(-x) - T^*(x)| \\ &= \max_{x \in \mathbb{R}} |f(x) - T^*(x)| = \|f - T^*\|. \end{aligned}$$

But now, the *even* trig polynomial

$$\widehat{T}(x) = \frac{\widetilde{T}(x) + T^*(x)}{2} = \frac{T^*(-x) + T^*(x)}{2}$$

is also a best approximation out of $\mathcal{T}_n$ since

$$\|f - \widehat{T}\,\| = \left\| \frac{(f - \widetilde{T}\,) + (f - T^*)}{2} \right\| \le \frac{\|f - \widetilde{T}\,\| + \|f - T^*\|}{2} = \min_{T \in \mathcal{T}_n} \|f - T\|. \qquad \square$$

We next give (de la Vallée Poussin's version of) Lebesgue's proof of *Weierstrass's second theorem*; that is, we will deduce the second theorem from the first.

Theorem. *Let $f \in C^{2\pi}$ and let $\varepsilon > 0$. Then, there is a trig polynomial T such that $\|f - T\| = \max_{x\in\mathbb{R}} |f(x) - T(x)| < \varepsilon$.*

PROOF. We will prove that Weierstrass's first theorem for $C[-1,1]$ implies his second theorem for $C^{2\pi}$.

Step 1. If f is even, then f may be uniformly approximated by even trig polynomials.

If f is even, then it's enough to approximate f on the interval $[0,\pi]$. In this case, we may consider the function $g(y) = f(\arccos y)$, $-1 \le y \le 1$, in $C[-1,1]$. By Weierstrass's first theorem, there is an algebraic polynomial $p(y)$ such that

$$\max_{-1\le y\le 1} |f(\arccos y) - p(y)| = \max_{0\le x\le \pi} |f(x) - p(\cos x)| < \varepsilon.$$

But $T(x) = p(\cos x)$ is an even trig polynomial! Hence,

$$\|f - T\| = \max_{x\in\mathbb{R}} |f(x) - T(x)| < \varepsilon.$$

Let's agree to abbreviate $\|f - T\| < \varepsilon$ as $f \approx T$.

Step 2. Given $f \in C^{2\pi}$, there is a trig polynomial T such that $2f(x)\sin^2 x \approx T(x)$.

Each of the functions $f(x) + f(-x)$ and $[f(x) - f(-x)]\sin x$ is *even*. Thus, we may choose even trig polynomials T_1 and T_2 such that

$$f(x) + f(-x) \approx T_1(x) \qquad \text{and} \qquad [f(x) - f(-x)]\sin x \approx T_2(x).$$

Multiplying the first expression by $\sin^2 x$, the second by $\sin x$, and adding, we get

$$2f(x)\sin^2 x \approx T_1(x)\sin^2 x + T_2(x)\sin x \equiv T_3(x),$$

where $T_3(x)$ is still a trig polynomial, and where "$\approx$" now means "within 2ε" (since $|\sin x| \le 1$).

Step 3. Given $f \in C^{2\pi}$, there is a trig polynomial T such that $2f(x)\cos^2 x \approx T(x)$, where "$\approx$" means "within 2ε."

Repeat Step 2 for $f(x - \pi/2)$ and translate: We first choose a trig polynomial $T_4(x)$ such that

$$2f\left(x - \frac{\pi}{2}\right)\sin^2 x \approx T_4(x).$$

That is,

$$2f(x)\cos^2 x \approx T_5(x),$$

where $T_5(x)$ is a trig polynomial.

Finally, by combining the conclusions of Steps 2 and 3, we find that there is a trig polynomial $T_6(x)$ such that $f \approx T_6(x)$, where, again, "$\approx$" means "within 2ε." $\square$

Just for fun, let's complete the circle and show that Weierstrass's second theorem for $C^{2\pi}$ implies his first theorem for $C[-1,1]$. Since, as we'll see, it's possible to give an independent proof of the second theorem, this is a meaningful exercise.

Theorem. *Given $f \in C[-1,1]$ and $\varepsilon > 0$, there exists an algebraic polynomial p such that $\|f - p\| < \varepsilon$.*

Proof. Given $f \in C[-1,1]$, the function $f(\cos x)$ is an even function in $C^{2\pi}$. By our Corollary to Weierstrass's second theorem, we may approximate $f(\cos x)$ by an even trig polynomial:

$$f(\cos x) \approx a_0 + a_1 \cos x + a_2 \cos 2x + \cdots + a_n \cos nx.$$

But, as we've seen, $\cos kx$ can be written as an algebraic polynomial in $\cos x$. Hence, there is some algebraic polynomial p such that $f(\cos x) \approx p(\cos x)$. That is,

$$\max_{0 \le x \le \pi} |f(\cos x) - p(\cos x)| = \max_{-1 \le t \le 1} |f(t) - p(t)| < \varepsilon. \qquad \square$$

The algebraic polynomials $T_n(x)$ satisfying

$$T_n(\cos x) = \cos nx, \text{ for } n = 0, 1, 2, \ldots,$$

are called the *Chebyshev polynomials of the first kind.* Please note that this formula *uniquely* defines T_n as a polynomial of degree exactly n, and hence uniquely determines the values of $T_n(x)$ for $|x| > 1$, too. The algebraic polynomials $U_n(x)$ satisfying

$$U_n(\cos x) = \frac{\sin(n+1)x}{\sin x}, \text{ for } n = 0, 1, 2, \ldots,$$

are called the *Chebyshev polynomials of the second kind.* Likewise, note that this formula uniquely defines U_n as a polynomial of degree exactly n.

We will discover many intriguing properties of the Chebyshev polynomials in the next chapter. For now, let's settle for just one: The recurrence formula we gave earlier

$$\cos nx = 2\cos x\,\cos(n-1)x - \cos(n-2)x$$

now becomes

$$T_n(x) = 2x\,T_{n-1}(x) - T_{n-2}(x), \quad n \ge 2,$$

where $T_0(x) = 1$ and $T_1(x) = x$. This recurrence relation (along with the initial cases T_0 and T_1) may be taken as a definition for the Chebyshev polynomials of the first kind. At any rate, it's now easy to list any number of the Chebyshev polynomials T_n; for example, the next few are $T_2(x) = 2x^2 - 1$, $T_3(x) = 4x^3 - 3x$, $T_4(x) = 8x^4 - 8x^2 + 1$, and $T_5(x) = 16x^5 - 20x^3 + 5x$.

A (real) trigonometric polynomial, or trig polynomial for short, is a function of the form

$$a_0 + \sum_{k=1}^{n} (\, a_k \cos kx + b_k \sin kx \,), \tag{$*$}$$

where $a_0, \ldots, a_n$ and $b_1, \ldots, b_n$ are real numbers. We will use $\mathcal{T}_n$ to denote the collection of trig polynomials of degree at most n, considered as a subspace $C^{2\pi}$, the space of all continuous 2π-periodic functions on $\mathbb{R}$. The space $C^{2\pi}$ may, in turn, be considered as a subspace of $C[\,0, 2\pi\,]$. Indeed, the 2π-periodic continuous functions on $\mathbb{R}$ may be identified with the subspace of $C[\,0, 2\pi\,]$ consisting of those f's which satisfy $f(0) = f(2\pi)$. As an alternate description, it is often convenient to instead identify $C^{2\pi}$ with the collection $C(\mathbb{T})$, consisting of all continuous real-valued functions on $\mathbb{T}$, where $\mathbb{T}$ is the unit circle in the complex plane $\mathbb{C}$. In this case, we simply make the identifications

$$\theta \longleftrightarrow e^{i\theta} \qquad \text{and} \qquad f(\theta) \longleftrightarrow f(e^{i\theta}).$$

▷ **41.** (a) By using the recurrence formulas $\cos kx + \cos(k-2)x = 2\cos(k-1)x\,\cos x$ and $\sin(k+1)x - \sin(k-1)x = 2\cos kx\,\sin x$, show that each of the functions $\cos kx$ and $\sin(k+1)x/\sin x$ may be written as algebraic polynomials of degree exactly k in $\cos x$. In each case, what is the coefficient of $\cos^k x$?

(b) Equivalently, use the binomial formula to write the real and imaginary parts of $(\cos x + i \sin x)^n = \cos nx + i \sin nx$ as algebraic polynomials in $\cos x$ and $\sin x$. Again, what are the leading coefficients of these polynomials?

(c) If $P(x, y)$ is an algebraic polynomial (in two variables) of degree at most n, show that $P(\cos x, \sin x)$ may be written as $Q(\cos x) + R(\cos x)\sin x$, where Q and R are algebraic polynomials (in one variable) of degrees at most n and $n-1$, respectively.

(d) Show that $\cos^n x$ can be written as a linear combination of the functions $\cos kx$, $k = 1, \ldots, n$, and that $\cos^{n-1} x \sin x$ can be written as a linear combinations of the functions $\sin kx$, $k = 1, \ldots, n$. Thus, each polynomial $P(\cos x, \sin x)$ in $\cos x$ and $\sin x$ can be written in the form $(*)$.

(e) If $(*)$ represents an even function, show that it can be written using only cosines. Conversely, if $P(x, y)$ is an even polynomial, show that $P(\cos x, \sin x)$ can be written using only cosines.

42. Show that $\mathcal{T}_n$ has dimension exactly $2n + 1$ (as a vector space over $\mathbb{R}$).

43. We might also consider complex trig polynomials; that is, functions of the form $(*)$ in which we now allow the a_k's and b_k's to be complex numbers.

(a) Show that every trig polynomial, whether real or complex, may be written as

$$\sum_{k=-n}^{n} c_k e^{ikx}, \tag{$**$}$$

where the c_k's are complex. Thus, complex trig polynomials are just algebraic polynomials in z and $\bar{z}$, where $z = e^{ix} \in \mathbb{T}$.

(b) Show that $(**)$ is real-valued if and only if $\bar{c}_k = c_{-k}$ for any k.

(c) If $(**)$ is a real-valued function, show that it may be written as a real trig polynomial; that is, it may be written in the form $(*)$ using only real coefficients.

We next discuss Chebyshev's solution to the problem of best polynomial approximation from 1854. Given that there was no reason to believe that the problem even had a solution, let alone a unique solution, Chebyshev's accomplishment should not be underestimated. Chebyshev might very well have been able to prove Weierstrass's result—30 years early—had the thought simply occurred to him! Chebyshev's original papers are apparently rather sketchy. It wasn't until 1903 that full details were given by Kirchberger. Curiously, Kirchberger's proofs foreshadow very modern techniques such as convexity and separation arguments. The presentation we'll give owes much to Haar and to de la Vallée Poussin (both from around 1918).

We begin with an easy observation:

Lemma. *Let $f \in C[\,a,b\,]$ and let $p = p_n^*$ be a best approximation to f out of $\mathcal{P}_n$. Then, there are at least two distinct points $x_1,\, x_2 \in [\,a,b\,]$ such that*

$$f(x_1) - p(x_1) = -(f(x_2) - p(x_2)) = \|f - p\|.$$

That is, $f - p$ attains both of the values $\pm\|f - p\|$.

PROOF. Let's write $E = E_n(f) = \|f - p\| = \max\limits_{a \le x \le b} |f(x) - p(x)|$. If the conclusion of the Lemma is false, then we might as well suppose that $f(x_1) - p(x_1) = E$, for some x_1, but that

$$e = \min_{a \le x \le b} \,(f(x) - p(x)) > -E.$$

In particular, $E + e \ne 0$ and so $q = p + (E + e)/2$ is an element of $\mathcal{P}_n$ with $q \ne p$. We claim that q is a better approximation to f than p. Here's why:

$$E - \left(\frac{E+e}{2}\right) \;\ge\; f(x) - p(x) - \left(\frac{E+e}{2}\right) \;\ge\; e - \left(\frac{E+e}{2}\right),$$

or

$$\left(\frac{E-e}{2}\right) \;\ge\; f(x) - q(x) \;\ge\; -\left(\frac{E-e}{2}\right).$$

That is,

$$\|f - q\| \;\le\; \left(\frac{E-e}{2}\right) \;<\; E \;=\; \|f - p\|,$$

a contradiction. $\square$

Corollary. *The best approximating constant to $f \in C[\,a,b\,]$ is*

$$p_0^* = \frac{1}{2}\left[\max_{a\le x\le b} f(x) + \min_{a\le x\le b} f(x)\right],$$

and

$$E_0(f) = \frac{1}{2}\left[\max_{a\le x\le b} f(x) - \min_{a\le x\le b} f(x)\right].$$

PROOF. **Exercise.**

Now all of this is meant as motivation for the general case, which essentially repeats the observation of our first Lemma inductively. A little experimentation will convince you that a best linear approximation, for example, would imply the existence of *three* points (at least) at which $f - p_1^*$ alternates between $\pm\|f - p_1^*\|$.

A bit of notation will help us set up the argument for the general case: Given g in $C[\,a,b\,]$, we'll say that $x \in [\,a,b\,]$ is a $(+)$ *point* for g (respectively, a $(-)$ *point* for g) if $g(x) = \|g\|$ (respectively, $g(x) = -\|g\|$). A set of distinct point $a \le x_0 < x_1 < \cdots < x_n \le b$ will be called an *alternating set* for g if the x_i's are alternately $(+)$ points and $(-)$ points; that is, if

$$|g(x_i)| = \|g\|, \qquad i = 0, 1, \ldots, n,$$

and

$$g(x_i) = -g(x_{i-1}), \qquad i = 1, 2, \ldots, n.$$

Using this notation, we will be able to characterize the polynomial of best approximation. Since the following three theorems are particularly important, we will number them for future reference. Our first result is where all the fighting takes place:

Theorem 1. *Let $f \in C[\,a,b\,]$, and suppose that $p = p_n^*$ is a best approximation to f out of $\mathcal{P}_n$. Then, there is an alternating set for $f - p$ consisting of at least $n + 2$ points.*

PROOF. If $f \in \mathcal{P}_n$, there's nothing to show. (Why?) Thus, we may suppose that $f \notin \mathcal{P}_n$ and, hence, that $E = E_n(f) = \|f - p\| > 0$.

Now consider the (uniformly) continuous function $\varphi = f - p$. We may partition $[\,a, b\,]$ by way of $a = t_0 < t_1 < \cdots < t_n = b$ into sufficiently small intervals so that

$$|\varphi(x) - \varphi(y)| < E/2 \qquad \text{whenever} \qquad x, y \in [\,t_i, t_{i+1}\,].$$

Here's why we'd want to do such a thing: If $[\,t_i, t_{i+1}\,]$ contains a $(+)$ point for $\varphi = f - p$, then φ is positive on all of $[\,t_i, t_{i+1}\,]$. Indeed,

$$x, y \in [\,t_i, t_{i+1}\,] \text{ and } \varphi(x) = E \quad \Longrightarrow \quad \varphi(y) > E/2 > 0.$$

Similarly, if $[\,t_i, t_{i+1}\,]$ contains a $(-)$ point for φ, then φ is negative on all of $[\,t_i, t_{i+1}\,]$. Consequently, **no interval** $[\,t_i, t_{i+1}\,]$ **can contain both** $(+)$ **points and** $(-)$ **points**.

Call $[\,t_i, t_{i+1}\,]$ a $(+)$ *interval* (respectively, a $(-)$ *interval*) if it contains a $(+)$ point (respectively, a $(-)$ point) for $\varphi = f - p$. Notice that **no** $(+)$ **interval can even touch a** $(-)$ **interval**. In other words, a $(+)$ interval and a $(-)$ interval must be strictly separated (by some interval containing a zero for φ).

We now relabel the $(+)$ and $(-)$ intervals from left to right, ignoring the "neither" intervals. There's no harm in supposing that the first "signed" interval is a $(+)$ interval. Thus, we suppose that our relabeled intervals are written

$$\begin{array}{ll} I_1, I_2, \ldots, I_{k_1} & (+) \text{ intervals,} \\ I_{k_1+1}, I_{k_1+2}, \ldots, I_{k_2} & (-) \text{ intervals,} \\ \quad \cdots\cdots\cdots\cdots\cdots\cdots & \\ I_{k_{m-1}+1}, I_{k_1+2}, \ldots, I_{k_m} & (-1)^{m-1} \text{ intervals,} \end{array}$$

where I_{k_1} is the last $(+)$ interval before we reach the first $(-)$ interval, I_{k_1+1}. And so on.

For later reference, we let S denote the union of all the "signed" intervals $[\,t_i, t_{i+1}\,]$; that is, $S = \bigcup_{j=1}^m I_{k_j}$, and we let N denote the union of all the "neither" intervals $[\,t_i, t_{i+1}\,]$. Thus, S and N are compact sets with $S \cup N = [\,a, b\,]$ (note that while S and N aren't quite disjoint, they are at least "non-overlapping"—their interiors are disjoint).

Our goal here is to show that $m \geq n + 2$. (So far we only know that $m \geq 2$!) Let's suppose that $m < n + 2$ and see what goes wrong.

Since any (+) interval is strictly separated from any (−) interval, we can find points $z_1, \ldots, z_{m-1} \in N$ such that

$$\max I_{k_1} < z_1 < \min I_{k_1+1}$$
$$\max I_{k_2} < z_2 < \min I_{k_2+1}$$
$$\cdots\cdots\cdots\cdots\cdots\cdots\cdots$$
$$\max I_{k_{m-1}} < z_{m-1} < \min I_{k_{m-1}+1}$$

And now we construct the offending polynomial:

$$q(x) = (z_1 - x)(z_2 - x)\cdots(z_{m-1} - x).$$

Notice that $q \in \mathcal{P}_n$ since $m - 1 \le n$. (Here is the *only* use we'll make of the assumption $m < n + 2$!) We're going to show that $p + \lambda q \in \mathcal{P}_n$ is a better approximation to f than p, for some suitable scalar λ.

We first claim that q and $f - p$ have the same sign. Indeed, q has no zeros in any of the (±) intervals, hence is of constant sign on any such interval. Thus, $q > 0$ on $I_1, \ldots, I_{k_1}$ because each $(z_j - x) > 0$ on these intervals; $q < 0$ on $I_{k_1+1}, \ldots, I_{k_2}$ because here $(z_1 - x) < 0$, while $(z_j - x) > 0$ for $j > 1$; and so on.

We next find λ. Let $e = \max_{x \in N} |f(x) - p(x)|$, where N is the union of all the subintervals $[\,t_i, t_{i+1}\,]$ which are neither (+) intervals nor (−) intervals. Then, $e < E$. (Why?) Now choose $\lambda > 0$ so that $\lambda\|q\| < \min\{E - e, E/2\}$. We claim that $p + \lambda q$ is a better approximation to f than p. One case is easy: If $x \in N$, then

$$|f(x) - (p(x) + \lambda q(x))| \;\le\; |f(x) - p(x)| + \lambda |q(x)| \;\le\; e + \lambda\|q\| \;<\; E.$$

On the other hand, if $x \notin N$, then x is in either a (+) interval or a (−) interval. In particular, we know that $|f(x) - p(x)| > E/2 > \lambda\|q\|$ and that $f(x) - p(x)$ and $\lambda q(x)$ have the same sign. Thus,

$$\begin{aligned} |f(x) - (p(x) + \lambda q(x))| \;&=\; |f(x) - p(x)| - \lambda |q(x)| \\ &\le\; E - \lambda \min_{x \in S} |q(x)| \;<\; E, \end{aligned}$$

since q is nonzero on S. This contradiction finishes the proof. (Phew!) $\square$

Remarks

1. It should be pointed out that the number $n+2$ here is actually $1 + \dim \mathcal{P}_n$.
2. Notice, too, that if $f - p_n^*$ alternates in sign $n+2$ times, then $f - p_n^*$ must have at least $n+1$ zeros. Thus, p_n^* actually agrees with f (or "interpolates" f) at $n+1$ points.

We're now ready to establish the uniqueness of the polynomial of best approximation.

Theorem 2. *Let $f \in C[\,a,b\,]$. Then, the polynomial of best approximation to f out of $\mathcal{P}_n$ is unique.*

Proof. Suppose that p, $q \in \mathcal{P}_n$ both satisfy $\|f - p\| = \|f - q\| = E_n(f) = E$. Then, as we've seen, their average $r = (p+q)/2 \in \mathcal{P}_n$ is also best: $\|f - r\| = E$ since $f - r = (f-p)/2 + (f-q)/2$.

By Theorem 1, $f - r$ has an alternating set $x_0, x_1, \ldots, x_{n+1}$, containing $n+2$ points. Thus, for each i,

$$(f-p)(x_i) + (f-q)(x_i) = \pm 2E \qquad \text{(alternating)},$$

while

$$-E \le (f-p)(x_i),\ (f-q)(x_i) \le E.$$

But this means that

$$(f-p)(x_i) = (f-q)(x_i) = \pm E \qquad \text{(alternating)}$$

for each i. (Why?) That is, $x_0, x_1, \ldots, x_{n+1}$ is an alternating set for both $f-p$ and $f-q$. In particular, the polynomial $q - p = (f-p) - (f-q)$ has $n+2$ zeros! Since $q - p \in \mathcal{P}_n$, we must have $p = q$. $\square$

Finally, we come full circle:

Theorem 3. *Let $f \in C[\,a,b\,]$, and let $p \in \mathcal{P}_n$. If $f - p$ has an alternating set containing $n+2$ (or more) points, then p is the best approximation to f out of $\mathcal{P}_n$.*

PROOF. Let $x_0, x_1, \ldots, x_{n+1}$ be an alternating set for $f-p$, and suppose that some $q \in \mathcal{P}_n$ is a better approximation to f than p; that is, $\|f-q\| < \|f-p\|$. In particular, then, we must have

$$|f(x_i) - p(x_i)| = \|f-p\| > \|f-q\| \geq |f(x_i) - q(x_i)|$$

for each $i = 0, 1, \ldots, n+1$. Now the inequality $|a| > |b|$ implies that a and $a-b$ have the same sign (why?), hence $q - p = (f-p) - (f-q)$ alternates in sign $n+2$ times (because $f-p$ does). But then, $q-p$ would have at least $n+1$ zeros. Since $q - p \in \mathcal{P}_n$, we must have $q = p$, which is a contradiction. Thus, p is the best approximation to f out of $\mathcal{P}_n$. □

Example (taken from Rivlin)

While an alternating set for $f - p_n^*$ is supposed to have at least $n+2$ points, it may well have more than $n+2$ points; thus, alternating sets need not be unique. For example, consider the function $f(x) = \sin 4x$ on $[-\pi, \pi\,]$. Since there are 8 points where f alternates between ± 1, it follows that $p_0^* = 0$ and that there are $4 \times 4 = 16$ different alternating sets consisting of exactly 2 points (not to mention all those with more than 2 points). In addition, notice that we actually have $p_1^* = \cdots = p_6^* = 0$, but that $p_7^* \neq 0$. (Why?)

Exercise

Show that $y = x - 1/8$ is the best linear approximation to $y = x^2$ on $[\,0, 1\,]$.

Essentially repeating the proof given for Theorem 3 yields a lower bound for $E_n(f)$.

Theorem. *Let $f \in C[\,a, b\,]$, and suppose that $q \in \mathcal{P}_n$ is such that $f(x_i) - q(x_i)$ alternates in sign at $n+2$ points $a \leq x_0 \leq x_1 \leq \ldots \leq x_{n+1} \leq b$. Then,*

$$E_n(f) \;\geq\; \min_{i=0,\ldots,n+1} |f(x_i) - q(x_i)|.$$

PROOF. If the inequality fails, then the best approximation $p = p_n^*$ would satisfy

$$\max_{1 \leq i \leq n+1} |f(x_i) - p(x_i)| \;\leq\; E_n(f) \;<\; \min_{1 \leq i \leq n+1} |f(x_i) - q(x_i)|.$$

Now we could repeat (essentially) the same argument used in the proof of Theorem 3 to arrive at a contradiction. The details are left as an **exercise**. □

Even for relatively simple functions, the problem of actually finding the polynomial of best approximation is genuinely difficult (even computationally). We end this section by stating two important problems that Chebyshev was able to solve.

Problem

Find the polynomial $p_{n-1}^* \in \mathcal{P}_{n-1}^*$, of degree at most $n-1$, that best approximates $f(x) = x^n$ on the interval $[-1, 1]$. (This particular choice of interval makes for a tidy solution; we'll discuss the general situation later.)

Since p_{n-1}^* is to minimize $\max\limits_{|x|\le 1} |x^n - p_{n-1}^*(x)|$, our first problem is equivalent to:

Problem

Find the monic polynomial of degree n which deviates least from 0 on $[-1, 1]$. In other words, find the monic polynomial of degree n which has smallest norm in $C[-1, 1]$.

We'll give two solutions to this problem (which we know has a unique solution, of course). First, let's simplify our notation. We write

$$p(x) = x^n - p_{n-1}^*(x) \quad \text{(the solution)},$$

and

$$M = \|p\| = E_{n-1}(x^n; [-1, 1]).$$

All we know about p is that it has an alternating set $-1 \le x_0 < x_1 < \cdots < x_n \le 1$ containing $(n-1) + 2 = n + 1$ points; that is, $|p(x_i)| = M$ and $p(x_{i+1}) = -p(x_i)$ for all i. Using this tiny bit of information, Chebyshev was led to compare the polynomials p^2 and p'. Watch closely!

Step 1. At any x_i in $(-1, 1)$, we must have $p'(x_i) = 0$ (because $p(x_i)$ is a relative extreme value for p). But, p' is a polynomial of degree $n-1$ and so can have at most $n-1$ zeros. Thus, we must have

$$x_i \in (-1, 1) \quad \text{and} \quad p'(x_i) = 0, \quad \text{for} \quad i = 1, \ldots, n-1,$$

(in fact, $x_1, \ldots, x_{n-1}$ are *all* the zeros of p') and

$$x_0 = -1, \quad p'(x_0) \ne 0, \quad x_{n-1} = 1, \quad p'(x_{n-1}) \ne 0.$$

Step 2. Now consider the polynomial $M^2 - p^2 \in \mathcal{P}_{2n}$. We know that $M^2 - (p(x_i))^2 = 0$ for $i = 0, 1, \ldots, n$, and that $M^2 - p^2 \geq 0$ on $[-1, 1]$. Thus, $x_1, \ldots, x_{n-1}$ must be double roots (at least) of $M^2 - p^2$. But this makes for $2(n-1) + 2 = 2n$ roots already, so we must have them all. Hence, $x_1, \ldots, x_{n-1}$ are double roots, x_0 and x_n are simple roots, and these are all the roots of $M^2 - p^2$.

Step 3. Next consider $(p')^2 \in \mathcal{P}_{2(n-1)}$. We know that $(p')^2$ has a double root at each of $x_1, \ldots, x_{n-1}$ (and no other roots), hence $(1 - x^2)(p'(x))^2$ has a double root at each $x_1, \ldots, x_{n-1}$, and simple roots at x_0 and x_n. Since $(1 - x^2)(p'(x))^2 \in \mathcal{P}_{2n}$, we've found all of its roots.

Here's the point to all this "rooting":

Step 4. Since $M^2 - (p(x))^2$ and $(1 - x^2)(p'(x))^2$ are polynomials of the same degree with the same roots, they are, up to a constant multiple, the same polynomial! It's easy to see what constant, too: The leading coefficient of p is 1 while the leading coefficient of p' is n; thus,

$$M^2 - (p(x))^2 \;=\; \frac{(1 - x^2)(p'(x))^2}{n^2}.$$

After tidying up,

$$\frac{p'(x)}{\sqrt{M^2 - (p(x))^2}} \;=\; \frac{n}{\sqrt{1 - x^2}}.$$

We really should have an extra $\pm$ here, but we know that p' is positive on *some* interval; we'll simply assume that it's positive on $[-1, x_1]$. Now, upon integrating,

$$\arccos\left(\frac{p(x)}{M}\right) \;=\; n \arccos x + C$$

or

$$p(x) = M \cos(n \arccos x + C).$$

But $p(-1) = -M$ (because $p'(-1) \geq 0$), so

$$\begin{aligned} \cos(n\pi + C) = -1 &\implies C = m\pi \quad (\text{with } n + m \text{ odd}) \\ &\implies p(x) = \pm M \cos(n \arccos x) \\ &\implies p(\cos x) = \pm M \cos nx. \end{aligned}$$

Look familiar? Since we know that $\cos nx$ is a polynomial of degree n with leading coefficient 2^{n-1} (the n-th Chebyshev polynomial T_n), the solution to our problem must be

$$p(x) \;=\; 2^{-n+1}\, T_n(x).$$

Since $|T_n(x)| \le 1$ for $|x| \le 1$ (why?), the minimum norm is $M = 2^{-n+1}$. $\square$

Next we give a "fancy" solution, based on our characterization of best approximations (Theorem 3) and a few simple properties of the Chebyshev polynomials.

Theorem. *For any $n \ge 1$, the formula $p(x) = x^n - 2^{-n+1}\,T_n(x)$ defines a polynomial $p \in \mathcal{P}_{n-1}$ satisfying*

$$2^{-n+1} \;=\; \max_{|x|\le 1} |x^n - p(x)| \;<\; \max_{|x|\le 1} |x^n - q(x)|$$

for any other $q \in \mathcal{P}_{n-1}$.

PROOF. We know that $2^{-n+1}\,T_n(x)$ has leading coefficient 1, and so $p \in \mathcal{P}_{n-1}$. Now set $x_k = \cos((n-k)\pi/n)$ for $k = 0, 1, \ldots, n$. Then, $-1 = x_0 < x_1 < \cdots < x_n = 1$ and

$$T_n(x_k) = T_n(\cos((n-k)\pi/n)) = \cos((n-k)\pi) = (-1)^{n-k}.$$

Since $|T_n(x)| = |T_n(\cos\theta)| = |\cos n\theta| \le 1$, for $-1 \le x \le 1$, we've found an alternating set for T_n containing $n+1$ points.

In other words, $x^n - p(x) = 2^{-n+1}\,T_n(x)$ satisfies $|x^n - p(x)| \le 2^{-n+1}$ and, for each $k = 0, 1, \ldots, n$, has $x_k^n - p(x_k) = 2^{-n+1}\,T_n(x_k) = (-1)^{n-k}2^{-n+1}$. By our characterization of best approximations (Theorem 3), p must be the best approximation to x^n out of $\mathcal{P}_{n-1}$. $\square$

Corollary. *The monic polynomial of degree exactly n having smallest norm in $C[\,a, b\,]$ is*

$$\frac{(b-a)^n}{2^n 2^{n-1}} \cdot T_n\left(\frac{2x-b-a}{b-a}\right).$$

PROOF. **Exercise**. [Hint: If $p(x)$ is a polynomial of degree n with leading coefficient 1, then $\tilde{p}(x) = p((2x-b-a)/(b-a))$ is a polynomial of degree n with leading coefficient $2^n/(b-a)^n$. Moreover, $\max\limits_{a\le x\le b} |p(x)| = \max\limits_{-1\le x\le 1} |\tilde{p}(x)|$.]

Properties of the Chebyshev Polynomials

As we've seen, the Chebyshev polynomial $T_n(x)$ is the (unique, real) polynomial of degree n (having leading coefficient 1 if $n = 0$, and 2^{n-1} if $n \ge 1$) such that $T_n(\cos\theta) = \cos n\theta$

for all θ. The Chebyshev polynomials have dozens of interesting properties and satisfy all sorts of curious equations. We'll catalogue just a few.

C1. $T_n(x) = 2x\,T_{n-1}(x) - T_{n-2}(x)$ for $n \ge 2$.

PROOF. It follows from the trig identity $\cos n\theta = 2\cos\theta\,\cos(n-1)\theta - \cos(n-2)\theta$ that $T_n(\cos\theta) = 2\cos\theta\,T_{n-1}(\cos\theta) - T_{n-2}(\cos\theta)$ for all θ; that is, the equation $T_n(x) = 2x\,T_{n-1}(x) - T_{n-2}(x)$ holds for all $-1 \le x \le 1$. But since both sides are polynomials, equality must hold for all x. $\square$

The next two properties are proved in essentially the same way:

C2. $T_m(x) + T_n(x) = \frac{1}{2}\big[T_{m+n}(x) + T_{m-n}(x)\big]$ for $m > n$.

C3. $T_m(T_n(x)) = T_{mn}(x)$.

C4. $T_n(x) = \frac{1}{2}\big[(x + \sqrt{x^2-1}\,)^n + (x - \sqrt{x^2-1}\,)^n\big]$.

PROOF. First notice that the expression on the right-hand side is actually a polynomial since, on combining the binomial expansions of $(x + \sqrt{x^2-1}\,)^n$ and $(x - \sqrt{x^2-1}\,)^n$, the odd powers of $\sqrt{x^2-1}$ cancel. Next, for $x = \cos\theta$,

$$\begin{aligned} T_n(x) = T_n(\cos\theta) = \cos n\theta &= \frac{1}{2}(e^{in\theta} + e^{-in\theta}) \\ &= \frac{1}{2}\big[(\cos\theta + i\sin\theta)^n + (\cos\theta - i\sin\theta)^n\big] \\ &= \frac{1}{2}\Big[\big(x + i\sqrt{1-x^2}\,\big)^n + \big(x - i\sqrt{1-x^2}\,\big)^n\Big] \\ &= \frac{1}{2}\Big[\big(x + \sqrt{x^2-1}\,\big)^n + \big(x - \sqrt{x^2-1}\,\big)^n\Big]. \end{aligned}$$

We've shown that these two polynomials agree for $|x| \le 1$, hence they must agree for all x (real or complex, for that matter). $\square$

For real x with $|x| \ge 1$, the expression $\frac{1}{2}\big[(x + \sqrt{x^2-1}\,)^n + (x - \sqrt{x^2-1}\,)^n\big]$ equals $\cosh(n\cosh^{-1}x)$. In other words,

C5. $T_n(\cosh x) = \cosh nx$ for all real x.

The next property also follows from property **C4**.

C6. $T_n(x) \le (|x| + \sqrt{x^2 - 1}\,)^n$ for $|x| \ge 1$.

An approach similar to the proof of property **C4** allows us to write x^n in terms of the Chebyshev polynomials $T_0, T_1, \ldots, T_n$.

C7. For n odd, $2^n x^n = \sum_{k=0}^{[n/2]} \binom{n}{k} 2\,T_{n-2k}(x)$; for n even, $2\,T_0$ should be replaced by T_0.

PROOF. For $-1 \le x \le 1$,

$$
\begin{aligned}
2^n x^n &= 2^n(\cos\theta)^n = (e^{i\theta} + e^{-i\theta})^n \\
&= e^{in\theta} + \binom{n}{1} e^{i(n-2)\theta} + \binom{n}{2} e^{i(n-4)\theta} + \cdots \\
&\qquad \cdots + \binom{n}{n-2} e^{-i(n-4)\theta} + \binom{n}{n-1} e^{-i(n-2)\theta} + e^{-in\theta} \\
&= 2\cos n\theta + \binom{n}{1} 2\cos(n-2)\theta + \binom{n}{2} 2\cos(n-4)\theta + \cdots \\
&= 2\,T_n(x) + \binom{n}{1} 2\,T_{n-2}(x) + \binom{n}{2} 2\,T_{n-4}(x) + \cdots,
\end{aligned}
$$

where, if n is even, the last term in this last sum is $\binom{n}{[n/2]} T_0$ (since the central term in the binomial expansion, namely $\binom{n}{[n/2]} = \binom{n}{[n/2]} T_0$, isn't doubled in this case). $\square$

C8. The zeros of T_n are $x_k^{(n)} = \cos((2k-1)\pi/2n)$, $k = 1, \ldots, n$. They're real, simple, and lie in the open interval $(-1, 1)$.

PROOF. Just check! But notice, please, that the zeros are listed here in *decreasing* order (because cosine decreases). $\square$

C9. Between two consecutive zeros of T_n, there is precisely one root of T_{n-1}.

PROOF. It's not hard to check that

$$
\frac{2k-1}{2n} < \frac{2k-1}{2\,(n-1)} < \frac{2k+1}{2n},
$$

for $k = 1, \ldots, n-1$, which means that $x_k^{(n)} > x_k^{(n-1)} > x_{k+1}^{(n)}$. $\square$

C10. T_n and T_{n-1} have no common zeros.

PROOF. Although this is immediate from property **C9**, there's another way to see it: $T_n(x_0) = 0 = T_{n-1}(x_0)$ implies that $T_{n-2}(x_0) = 0$ by property **C1**. Repeating this observation, we would have $T_k(x_0) = 0$ for every $k < n$, including $k = 0$. No good! $T_0(x) = 1$ has no zeros. $\square$

C11. The set $\{x_k^{(n)} : 1 \le k \le n,\ n = 1, 2, \ldots\}$ is dense in $[-1, 1]$.

PROOF. Since $\cos x$ is (strictly) monotone on $[0, \pi]$, it's enough to know that the set $\{(2k-1)\pi/2n\}_{k,n}$ is dense in $[0, \pi]$, and for this it's enough to know that $\{(2k-1)/2n\}_{k,n}$ is dense in $[0, 1]$. (Why?) But

$$\frac{2k-1}{2n} = \frac{k}{n} - \frac{1}{2n} \approx \frac{k}{n}$$

for n large; that is, the set $\{(2k-1)/2n\}_{k,n}$ is dense among the rationals in $[0, 1]$. $\square$

It's interesting to note here that the *distribution* of the roots $\{x_k^{(n)}\}_{k,n}$ can be estimated (see Natanson, *Constructive Function Theory*, Vol. I, pp. 48–51). For large n, the number of roots of T_n that lie in an interval $[x, x + \Delta x] \subset [-1, 1]$ is approximately

$$\frac{n\Delta x}{\pi\sqrt{1 - x^2}}.$$

In particular, for n large, the roots of T_n are "thickest" near the endpoints ± 1.

In probabilistic terms, this means that if we assign equal probability to each of the roots $x_0^{(n)}, \ldots, x_n^{(n)}$ (that is, if we think of each root as the position of a point with mass $1/n$), then the *density* of this probability distribution (or the density of the system of point masses) at a point x is approximately $1/\pi\sqrt{1 - x^2}$ for large n. In still other words, this tells us that the probability that a root of T_n lies in the interval $[a, b]$ is approximately

$$\frac{1}{\pi} \int_a^b \frac{1}{\sqrt{1 - x^2}}\, dx\,.$$

C12. The Chebyshev polynomials are mutually *orthogonal* relative to the weight $w(x) = (1 - x^2)^{-1/2}$ on $[-1, 1]$.

PROOF. For $m \ne n$ the substitution $x = \cos\theta$ yields

$$\int_{-1}^{1} T_n(x)\, T_m(x)\, \frac{dx}{\sqrt{1 - x^2}} = \int_0^{\pi} \cos m\theta\, \cos n\theta\, d\theta = 0,$$

while for $m = n$ we get

$$\int_{-1}^{1} T_n^2(x)\,\frac{dx}{\sqrt{1-x^2}} \;=\; \int_0^{\pi} \cos^2 n\theta\,d\theta = \begin{cases} \pi & \text{if } n = 0 \\ \pi/2 & \text{if } n > 0. \end{cases} \qquad \square$$

C13. $|T_n'(x)| \le n^2$ for $-1 \le x \le 1$, and $|T_n'(\pm 1)| = n^2$.

PROOF. For $-1 < x < 1$ we have

$$\frac{d}{dx} T_n(x) \;=\; \frac{\frac{d}{d\theta}\, T_n(\cos\theta)}{\frac{d}{d\theta}\,\cos\theta} \;=\; \frac{n \sin n\theta}{\sin\theta}.$$

Thus, $|T_n'(x)| \le n^2$ because $|\sin n\theta| \le n|\sin\theta|$ (which can be easily checked by induction, for example). At $x = \pm 1$, we interpret this derivative formula as a limit (as $\theta \to 0$ and $\theta \to \pi$) and find that $|T_n'(\pm 1)| = n^2$. $\square$

As we'll see later, each $p \in \mathcal{P}_n$ satisfies $|p'(x)| \le \|p\| n^2 = \|p\|\, T_n'(1)$ for $-1 \le x \le 1$, and this is, of course, best possible. As it happens, $T_n(x)$ has the largest possible rate of growth outside of $[-1, 1]$ among all polynomials of degree n. Specifically:

Theorem. *Let $p \in \mathcal{P}_n$ and let $\|p\| = \max\limits_{-1 \le x \le 1} |p(x)|$. Then, for any x_0 with $|x_0| \ge 1$ and any $k = 0, 1, \ldots, n$ we have*

$$|p^{(k)}(x_0)| \le \|p\|\, |T_n^{(k)}(x_0)|,$$

where $p^{(k)}$ is the k-th derivative of p.

We'll prove only the case $k = 0$. In other words, we'll check that $|p(x_0)| \le \|p\|\, |T_n(x_0)|$. The more general case is in Rivlin, Theorem 1.10, p. 31.

PROOF. Since all the zeros of T_n lie in $(-1, 1)$, we know that $T_n(x_0) \ne 0$. Thus, we may consider the polynomial

$$q(x) \;=\; \frac{p(x_0)}{T_n(x_0)}\, T_n(x) \;-\; p(x) \in \mathcal{P}_n.$$

If the claim is *false*, then

$$\|p\| < \left| \frac{p(x_0)}{T_n(x_0)} \right|.$$

Now at each of the points $y_k = \cos(k\pi/n)$, $k = 0, 1, \ldots, n$, we have $T_n(y_k) = (-1)^k$ and, hence,

$$q(y_k) \;=\; (-1)^k \frac{p(x_0)}{T_n(x_0)} \;-\; p(y_k).$$

Since $|p(y_k)| \le \|p\|$, it follows that q alternates in sign at these $n+1$ points. In particular, q must have at least n zeros in $(-1, 1)$. But $q(x_0) = 0$, by design, and $|x_0| \ge 1$. That is, we've found $n+1$ zeros for a polynomial of degree n. So, $q \equiv 0$; that is,

$$p(x) \;=\; \frac{p(x_0)}{T_n(x_0)}\, T_n(x).$$

But then,

$$|p(1)| = \left| \frac{p(x_0)}{T_n(x_0)} \right| > \|p\|,$$

since $T_n(1) = T_n(\cos 0) = 1$, which is a contradiction. $\square$

Corollary. *Let $p \in \mathcal{P}_n$ and let $\|p\| = \max\limits_{-1 \le x \le 1} |p(x)|$. Then, for any x_0 with $|x_0| \ge 1$, we have*

$$|p(x_0)| \le \|p\| \left(|x_0| + \sqrt{x_0^2 - 1} \right)^n .$$

Rivlin's proof of our last Theorem in the general case uses the following observation:

C14. For $x \ge 1$ and $k = 0, 1, \ldots, n$, we have $T_n^{(k)}(x) > 0$.

PROOF. **Exercise**. [Hint: It follows from Rolle's theorem that $T_n^{(k)}$ is never zero for $x \ge 1$. (Why?) Now just compute $T_n^{(k)}(1)$.]

Uniform Approximation by Trig Polynomials

We end this section by summarizing (without proofs) the analogues of Theorems 1–3 for uniform approximation by trig polynomials. Throughout, $f \in C^{2\pi}$ and $\mathcal{T}_n$ denotes the collection of trig polynomials of degree at most n.

1. f has a best approximation $T^* \in \mathcal{T}_n$.

2. $f - T^*$ has an alternating set containing $2n+2$ (or more) points in $[0, 2\pi)$. (Note here that $2n + 2 = 1 + \dim \mathcal{T}_n$.)

3. T^* is unique.

4. If $T \in \mathcal{T}_n$ is such that $f - T$ has an alternating set containing $2n+2$ or more points in $[\,0, 2\pi)$, then $T = T^*$.

The proofs of **1**–**4** are very similar to the corresponding results for algebraic polynomials. As you might imagine, **2** is where all the fighting takes place, and there are a few technical difficulties to cope with. Nevertheless, we'll swallow these facts whole and apply them with a clear conscience to a few examples.

Example

For $m > n$, the best approximation to $f(x) = A\cos mx + B\sin mx$ out of $\mathcal{T}_n$ is 0!

PROOF. We may write $f(x) = R\cos m(x - x_0)$ for some R and x_0. (How?) Now we need only display a sufficiently large alternating set for f (in some interval of length 2π).

Setting $x_k = x_0 + k\pi/m$, $k = 1, 2, \ldots, 2m$, we get $f(x_k) = R\cos k\pi = R(-1)^k$ and $x_k \in (x_0, x_0 + 2\pi]$. Since $m > n$, it follows that $2m \geq 2n + 2$. $\square$

Example

The best approximation to

$$f(x) = a_0 + \sum_{k=1}^{n+1} \big(a_k \cos kx + b_k \sin kx\big)$$

out of $\mathcal{T}_n$ is

$$T(x) = a_0 + \sum_{k=1}^{n} \big(a_k \cos kx + b_k \sin kx\big),$$

and $\|f - T\| = \sqrt{a_{n+1}^2 + b_{n+1}^2}$ in $C^{2\pi}$.

PROOF. By our last example, the best approximation to $f - T$ out of $\mathcal{T}_n$ is 0, hence T must be the best approximation to f. (Why?) The last assertion is easy to check: Since we can always write $A\cos mx + B\sin mx = \sqrt{A^2 + B^2} \cdot \cos m(x - x_0)$, for some x_0, it follows that $\|f - T\| = \sqrt{a_{n+1}^2 + b_{n+1}^2}$. $\square$

Finally, let's make a simple connection between the two types of polynomial approximation:

Theorem. *Let $f \in C[-1,1]$ and define $\varphi \in C^{2\pi}$ by $\varphi(\theta) = f(\cos\theta)$. Then,*

$$E_n(f) = \min_{p \in \mathcal{P}_n} \|f - p\| = \min_{T \in \mathcal{T}_n} \|\varphi - T\| \equiv E_n^T(\varphi).$$

PROOF. Suppose that $p^*(x) = \sum_{k=0}^n a_k x^k$ is the best approximation to f out of $\mathcal{P}_n$. Then, $\widehat{T}(\theta) = p^*(\cos\theta)$ is in $\mathcal{T}_n$ and, clearly,

$$\max_{-1 \le x \le 1} |f(x) - p^*(x)| = \max_{0 \le \theta \le 2\pi} |f(\cos\theta) - p^*(\cos\theta)|.$$

Thus, $\|f - p^*\| = \|\varphi - \widehat{T}\| \ge \min_{T \in \mathcal{T}_n} \|\varphi - T\|$.

On the other hand, since φ is *even*, we know that T^*, its best approximation out of $\mathcal{T}_n$, is also even. Thus, $T^*(\theta) = q(\cos\theta)$ for some algebraic polynomial $q \in \mathcal{P}_n$. Consequently, $\|\varphi - T^*\| = \|f - q\| \ge \min_{p \in \mathcal{P}_n} \|f - p\|$. $\square$

Remarks

1. Once we know that $\min_{p \in \mathcal{P}_n} \|f - p\| = \min_{T \in \mathcal{T}_n} \|\varphi - T\|$, it follows that we must also have $T^*(\theta) = p^*(\cos\theta)$.

2. Each *even* $\varphi \in C^{2\pi}$ corresponds to an $f \in C[-1,1]$ by $f(x) = \varphi(\arccos x)$ and, of course, the conclusions of the Theorem and of Remark 1 hold in this case, too.

3. Whenever we speak of even trig polynomials, the Chebyshev polynomials are lurking somewhere in the background. Indeed, let $T(\theta)$ be an even trig polynomial, write $x = \cos\theta$, as usual, and consider the following cryptic equation:

$$T(\theta) = \sum_{k=0}^n a_k \cos k\theta = \sum_{k=0}^n a_k T_k(\cos\theta) = p(\cos\theta),$$

where $p(x) = \sum_{k=0}^n a_k T_k(x) \in \mathcal{P}_n$.

We've shown that $\cos n\theta$ and $\sin(n+1)\theta/\sin\theta$ can be written as algebraic polynomials of degree n in $\cos\theta$; we use this observation to define the *Chebyshev polynomials*. The Chebyshev polynomials of the *first kind* $(T_n(x))$ are defined by $T_n(\cos\theta) = \cos n\theta$, for $n = 0, 1, 2, \ldots$, while the Chebyshev polynomials of the *second kind* $(U_n(x))$ are defined by $U_n(\cos\theta) = \sin(n+1)\theta/\sin\theta$ for $n = 0, 1, 2, \ldots$.

▷ **44.** Establish the following properties of $T_n(x)$.

(i) $T_0(x) = 1$, $T_1(x) = x$, and $T_n(x) = 2xT_{n-1}(x) - T_{n-2}(x)$ for $n \geq 2$.

(ii) $T_n(x)$ is a polynomial of degree n having leading coefficient 2^{n-1} for $n \geq 1$, and containing only even (resp., odd) powers of x if n is even (resp., odd).

(iii) $|T_n(x)| \leq 1$ for $-1 \leq x \leq 1$; when does equality occur? Where are the zeros of $T_n(x)$? Show that between two consecutive zeros of $T_n(x)$ there is exactly one zero of $T_{n-1}(x)$. Can $T_n(x)$ and $T_{n-1}(x)$ have a common zero?

(iv) $|T_n'(x)| \leq n^2$ for $-1 \leq x \leq 1$, and $|T_n'(\pm 1)| = n^2$.

(v) $T_m(x) + T_n(x) = \frac{1}{2}\big[T_{m+n}(x) + T_{m-n}(x)\big]$ for $m > n$.

(vi) $T_m(T_n(x)) = T_{mn}(x)$.

(vii) Evaluate $\displaystyle\int_{-1}^{1} T_n(x)\,T_m(x)\,\frac{dx}{\sqrt{1-x^2}}$.

(viii) Show that T_n is a solution to $(1-x^2)y'' - xy' + n^2y = 0$.

(ix) $T_n(x) = \frac{1}{2}\big[(x + \sqrt{x^2-1}\,)^n + (x - \sqrt{x^2-1}\,)^n\big]$ for any x, real or complex.

(x) $\mathrm{Re}\left(\sum_{n=0}^\infty t^n e^{in\theta}\right) = \sum_{n=0}^\infty t^n \cos n\theta = \dfrac{1 - t\cos\theta}{1 - 2t\cos\theta + t^2}$ for $-1 < t < 1$; that is, $\sum_{n=0}^\infty t^n T_n(x) = \dfrac{1 - tx}{1 - 2tx + t^2}$ (this is a *generating function* for T_n; it's closely related to the Poisson kernel).

(xi) Find analogues of (i)–(x) (if possible) for $U_n(x)$.

▷ **45.** Show that every $p \in \mathcal{P}_n$ has a unique representation as $p = a_0 + a_1T_1 + \cdots + a_nT_n$. Find this representation in the case $p(x) = x^n$.

▷ **46.** The polynomial of degree n having leading coefficient 1 and deviating least from 0 on $[-1, 1]$ is given by $T_n(x)/2^{n-1}$. On an arbitrary interval $[a, b]$ we would instead take

$$\frac{(b-a)^n}{2^{2n-1}}\, T_n\left(\frac{2x-b-a}{b-a}\right).$$

Is this solution unique? Explain.

47. If p is a polynomial on $[a, b]$ of degree n having leading coefficient $a_n > 0$, then $\|p\| \geq a_n(b-a)^n/2^{2n-1}$. If $b - a \geq 4$, then no polynomial of degree exactly n with *integer* coefficients can satisfy $\|p\| < 2$ (compare this with problem 26 on the "Uniform Approximation by Polynomials" problem set).

48. Given $p \in \mathcal{P}_n$, show that $|p(x)| \leq \|p\|\, |T_n(x)|$ for $|x| > 1$.

49. If $p \in \mathcal{P}_n$ with $\|p\| = 1$ on $[-1, 1]$, and if $|p(x_i)| = 1$ at $n+1$ distinct point $x_0, \ldots, x_n$ in $[-1, 1]$, show that either $p = \pm 1$, or else $p = \pm T_n$. [Hint: One approach is to compare the polynomials $1 - p^2$ and $(1 - x^2)(p\,')^2$.]

50. Compute $T_n^{(k)}(1)$ for $k = 0, 1, \ldots, n$, where $T_n^{(k)}$ is the k-th derivative of T_n. For $x \geq 1$ and $k = 0, 1, \ldots, n$, show that $T_n^{(k)}(x) > 0$.

The following examples are cribbed from the book *Chebyshev Polynomials*, by L. Fox and I. B. Parker (Oxford University Press, 1968).

As we've seen, the Chebyshev polynomals can be generated by a recurrence relation. By reversing the procedure, we could solve for x^n in terms of $T_0, T_1, \ldots, T_n$ (we'll do this calculation in class). Here are the first few terms in each of these relations:

$$\begin{aligned}
T_0(x) &= 1 & 1 &= T_0(x) \\
T_1(x) &= x & x &= T_1(x) \\
T_2(x) &= 2x^2 - 1 & x^2 &= (T_0(x) + T_2(x))/2 \\
T_3(x) &= 4x^3 - 3x & x^3 &= (3\,T_1(x) + T_3(x))/4 \\
T_4(x) &= 8x^4 - 8x^2 + 1 & x^4 &= (3\,T_0(x) + 4\,T_2(x) + T_4(x))/8 \\
T_5(x) &= 16x^5 - 20x^3 + 5x & x^5 &= (10\,T_1(x) + 5\,T_3(x) + T_5(x))/16
\end{aligned}$$

Note the separation of even and odd terms in each case. Writing ordinary, garden variety polynomials in their equivalent Chebyshev form has some distinct advantages for numerical computations. Here's why:

$$1 - x + x^2 - x^3 + x^4 = \frac{15}{6}\,T_0(x) - \frac{7}{4}\,T_1(x) + T_2(x) - \frac{1}{4}\,T_3(x) + \frac{1}{8}\,T_4(x)$$

(after some simplification). Now we see at once that we can get a cubic approximation to $1 - x + x^2 - x^3 + x^4$ on $[-1, 1]$ with error at most $1/8$ by simply dropping the T_4 term on the right-hand side (since $|T_4(x)| \le 1$), whereas simply using $1 - x + x^2 - x^3$ as our cubic approximation could cause an error as big as 1. Pretty slick! This gimmick of truncating the equivalent Chebyshev form is called *economization*.

As a second example we note that a polynomial with small norm on $[-1, 1]$ may have annoyingly large coefficients:

$$\begin{aligned}
(1 - x^2)^{10} = 1 - 10x^2 + 45x^4 - 120x^6 + 210x^8 - 252x^{10} \\
+ 210x^{12} - 120x^{14} + 45x^{16} - 10x^{18} + x^{20}
\end{aligned}$$

but in Chebyshev form (look out!):

$$(1-x^2)^{10} = \frac{1}{524{,}288}\{92{,}378\, T_0(x) - 167{,}960\, T_2(x) + 125{,}970\, T_4(x) - 77{,}520\, T_6(x)$$
$$+ 38{,}760\, T_8(x) - 15{,}504\, T_{10}(x) + 4{,}845\, T_{12}(x) - 1{,}140\, T_{14}(x)$$
$$+ 190\, T_{16}(x) - 20\, T_{18(x)} + T_{20}(x)\}$$

The largest coefficient is now only about 0.3, and the omission of the last three terms produces a maximum error of about 0.0004. Not bad.

As a last example, consider the Taylor polynomial $e^x = \sum_{k=0}^{n} x^k/k! + x^{n+1}e^{\xi}/(n+1)!$ (with remainder), where $-1 \le x,\ \xi \le 1$. Taking $n = 6$, the truncated series has error no greater than $e/7! \approx 0.0005$. But if we "economize" the first six terms, then:

$$\sum_{k=0}^{6} x^k/k = 1.26606\, T_0(x) + 1.13021\, T_1(x) + 0.27148\, T_2(x) + 0.04427\, T_3(x)$$
$$+ 0.00547\, T_4(x) + 0.00052\, T_5(x) + 0.00004\, T_6(x).$$

The initial approximation already has an error of about 0.0005, so we can certainly drop the T_6 term without any additional error. Even dropping the T_5 term causes an error of no more than 0.001 (or thereabouts). The resulting approximation has a far smaller error than the corresponding truncated Taylor series: $e/5! \approx 0.023$.

The approach used in our last example has the decided disadvantage that we must first decide where to truncate the Taylor series—which might converge very slowly. A better approach would be to write e^x as a series involving Chebyshev polynomials directly. That is, if possible, we want to write $e^x = \sum_{k=0}^{\infty} a_k T_k(x)$. If the a_k's are absolutely summable, it will be very easy to estimate any truncation error. We'll get some idea on how to go about this when we talk about "least-squares" approximation. As it happens, such a series is easy to find (it's rather like a Fourier series), and its partial sums are remarkably good uniform approximations.

Our goal in this section is to prove the following result (as well as discuss its ramifications). In fact, this result is so fundamental that we will present three proofs!

Theorem. *Let $x_0, x_1, \ldots, x_n$ be distinct points, and let $y_0, y_1, \ldots, y_n$ be arbitrary points in $\mathbb{R}$. Then, there exists a unique polynomial $p \in \mathcal{P}_n$ satisfying $p(x_i) = y_i$, $i = 0, 1, \ldots, n$.*

First notice that uniqueness is obvious. Indeed, if two polynomials p, $q \in \mathcal{P}_n$ agree at $n+1$ points, then $p \equiv q$. (Why?) The real work comes in proving existence.

FIRST PROOF. (Vandermonde's determinant.) We seek $c_0, c_1, \ldots, c_n$ so that $p(x) = \sum_{k=0}^n c_k x^k$ satisfies

$$p(x_i) = \sum_{k=0}^{n} c_k x_i^k = y_i, \quad i = 0, 1, \ldots, n.$$

That is, we need to solve a system of $n+1$ *linear* equations for the c_i's. In matrix form:

$$\begin{bmatrix} 1 & x_0 & x_0^2 & \cdots & x_0^n \\ 1 & x_1 & x_1^2 & \cdots & x_1^n \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 1 & x_n & x_n^2 & \cdots & x_n^n \end{bmatrix} \begin{bmatrix} c_0 \\ c_1 \\ \vdots \\ c_n \end{bmatrix} = \begin{bmatrix} y_0 \\ y_1 \\ \vdots \\ y_n \end{bmatrix}.$$

This equation always has a unique solution because the coefficient matrix has determinant

$$D = \prod_{0 \le j < i \le n} (x_i - x_j) \neq 0.$$

(D is called *Vandermonde's determinant*; note that $D > 0$ if $x_0 < x_1 < \cdots < x_n$.) Since this fact is of independent interest, we'll sketch a short proof below. $\square$

Lemma. $D = \prod_{0 \le j < i \le n} (x_i - x_j)$.

PROOF. Consider

$$V(x_0, x_1, \ldots, x_{n-1}, x) = \begin{vmatrix} 1 & x_0 & x_0^2 & \cdots & x_0^n \\ 1 & x_1 & x_1^2 & \cdots & x_1^n \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 1 & x_{n-1} & x_{n-1}^2 & \cdots & x_{n-1}^n \\ 1 & x & x^2 & \cdots & x^n \end{vmatrix}.$$

$V(x_0, x_1, \ldots, x_{n-1}, x)$ is a *polynomial* of degree n in x, and it's 0 whenever $x = x_i$, $i = 0, 1, \ldots, n-1$. Thus, $V(x_0, \ldots, x) = c \prod_{i=0}^{n-1}(x - x_i)$, by comparing roots and degree. However, it's easy to see that the coefficient of x^n in $V(x_0, \ldots, x)$ is $V(x_0, \ldots, x_{n-1})$. Thus, $V(x_0, \ldots, x) = V(x_0, \ldots, x_{n-1}) \prod_{i=0}^{n-1}(x - x_i)$. The result now follows by induction and the obvious case $\begin{vmatrix} 1 & x_0 \\ 1 & x_1 \end{vmatrix} = x_1 - x_0$. $\square$

Second Proof. (Lagrange interpolation.) We could define p immediately if we had polynomials $\ell_i(x) \in \mathcal{P}_n$, $i = 0, \ldots, n$, such that $\ell_i(x_j) = \delta_{i,j}$ (where $\delta_{i,j}$ is *Kronecker's delta*; that is, $\delta_{i,j} = 0$ for $i \neq j$, and $\delta_{i,j} = 1$ for $i = j$). Indeed, $p(x) = \sum_{i=0}^{n} y_i\, \ell_i(x)$ would then work as our interpolating polynomial. In short, notice that the polynomials $\{\ell_0, \ell_1, \ldots, \ell_n\}$ would form a (particularly convenient) *basis* for $\mathcal{P}_n$.

We'll give two formulas for $\ell_i(x)$:

(a). Clearly, $\ell_i(x) = \displaystyle\prod_{j \neq i} \frac{x - x_j}{x_i - x_j}$ works.

(b). Start with $W(x) = (x - x_0)(x - x_1)\cdots(x - x_n)$, and notice that the polynomial we need satisfies

$$\ell_i(x) = a_i \cdot \frac{W(x)}{x - x_i}$$

for some $a_i \in \mathbb{R}$. (Why?) But then, $1 = \ell_i(x_i) = a_i W'(x_i)$ (again, why?); that is,

$$\ell_i(x) = \frac{W(x)}{(x - x_i)\, W'(x_i)}.$$

Please note that $\ell_i(x)$ is a multiple of the polynomial $\prod_{j \neq i}(x - x_j)$, for $i = 0, \ldots, n$, and that $p(x)$ is then a suitable linear combination of such polynomials. $\square$

Third Proof. (Newton's formula.) We seek $p(x)$ of the form

$$p(x) = a_0 + a_1(x - x_0) + a_2(x - x_0)(x - x_1) + \cdots + a_n(x - x_0)\cdots(x - x_{n-1}).$$

(Please note that x_n does not appear on the right-hand side.) This form makes it almost effortless to solve for the a_i's by plugging-in the x_i's, $i = 0, \ldots, n-1$.

$$y_0 = p(x_0) = a_0$$
$$y_1 = p(x_1) = a_0 + a_1(x_1 - x_0) \implies a_1 = \frac{y_1 - a_0}{x_1 - x_0}.$$

Continuing, we find

$$a_2 = \frac{y_2 - a_0 - a_1(x_2 - x_0)}{(x_2 - x_0)(x_2 - x_1)}$$

$$a_3 = \frac{y_3 - a_0 - a_1(x_3 - x_0) - a_2(x_3 - x_0)(x_3 - x_1)}{(x_3 - x_0)(x_3 - x_1)(x_3 - x_2)}$$

and so on. [Natanson, Vol. III, gives another formula for the a_i's.] $\square$

Example

As a quick means of comparing these three solutions, let's find the interpolating polynomial (quadratic) passing through $(1, 2)$, $(2, -1)$, and $(3, 1)$. You're invited to check the following:

Vandermonde: $p(x) = 10 - \frac{21}{2}\,x + \frac{5}{2}\,x^2$.

Lagrange: $p(x) = (x-2)(x-3) + (x-1)(x-3) + \frac{1}{2}(x-1)(x-2)$.

Newton: $p(x) = 2 - 3(x-1) + \frac{5}{2}(x-1)(x-2)$.

As you might have already surmised, Lagrange's method is the easiest to apply by hand, although Newton's formula has much to recommend it too (it's especially well-suited to situations where we introduce additional nodes). We next set up the necessary notation to discuss the finer points of Lagrange's method.

Given $n+1$ distinct points $a \le x_0 < x_1 < \cdots < x_n \le b$ (sometimes called *nodes*), we first form the polynomials

$$W(x) \;=\; \prod_{i=0}^{n} (x - x_i)$$

and

$$\ell_i(x) \;=\; \prod_{j \ne i} \frac{x - x_j}{x_i - x_j} \;=\; \frac{W(x)}{(x - x_i)\, W'(x_i)}.$$

The *Lagrange interpolation formula* is

$$L_n(f)(x) \;=\; \sum_{i=0}^{n} f(x_i)\, \ell_i(x).$$

That is, $L_n(f)$ is the unique polynomial in $\mathcal{P}_n$ that agrees with f at the x_i's. In particular, notice that we must have $L_n(p) = p$ whenever $p \in \mathcal{P}_n$. In fact, L_n is a *linear projection* from $C[\,a, b\,]$ onto $\mathcal{P}_n$. [Why is $L_n(f)$ linear in f?]

Typically we're given (or construct) *an array of nodes*:

$$X \quad \begin{cases} x_0^{(0)} & & & \\ x_0^{(1)} & x_1^{(1)} & & \\ x_0^{(2)} & x_1^{(2)} & x_2^{(2)} & \\ \vdots & & & \ddots \end{cases}$$

and form the corresponding sequence of projections

$$L_n(f)(x) \;=\; \sum_{i=0}^{n} f(x_i^{(n)})\,\ell_i^{(n)}(x).$$

An easy (but admittedly pointless) observation is that for a *given* $f \in C[\,a,b\,]$ we can always find an array X so that $L_n(f) = p_n^*$, the polynomial of best approximation to f out of $\mathcal{P}_n$ (since $f - p_n^*$ has $n+1$ zeros, we may use these for the x_i's). Thus, $\|L_n(f) - f\| = E_n(f) \to 0$ in this case. However, the problem of convergence changes character dramatically if we *first choose* X and then consider $L_n(f)$. In general, there's no reason to believe that $L_n(f)$ converges to f. In fact, quite the opposite is true:

Theorem. (Faber, 1914) *Given any array of nodes X in $[\,a,b\,]$, there is some $f \in C[\,a,b\,]$ for which $\|L_n(f) - f\|$ is unbounded.*

The problem here has little to do with interpolation and everything to do with projections:

Theorem. (Kharshiladze, Lozinski, 1941) *For each n, let L_n be a continuous, linear projection from $C[\,a,b\,]$ onto $\mathcal{P}_n$. Then, there is some $f \in C[\,a,b\,]$ for which $\|L_n(f) - f\|$ is unbounded.*

Evidently, the operators L_n aren't positive (monotone), for otherwise the Bohman-Korovkin theorem (and the fact that L_n is a projection onto $\mathcal{P}_n$) would imply that $L_n(f)$ converges uniformly to f for every $f \in C[\,a,b\,]$.

The proofs of these theorems are long and difficult—we'll save them for another day. (Some of you may recognize the Principle of Uniform Boundedness at work here.) The real point here is that we can't have everything: A positive result about convergence of interpolation will require that we impose some extra conditions on the functions f we want to approximate. As a first step in this direction, we prove that if f has sufficiently many derivatives, then the error $\|L_n(f) - f\|$ can at least be measured.

Theorem. *Suppose that f has $n+1$ continuous derivatives on $[a, b]$. Let $a \le x_0 < x_1 < \cdots < x_n \le b$, let $p \in \mathcal{P}_n$ be the polynomial that interpolates f at the x_i's, and let $W(x) = \prod_{i=0}^{n}(x - x_i)$. Then,*

$$\|f - p\| \le \frac{1}{(n+1)!}\, \|f^{(n+1)}\|\, \|W\|.$$

Proof. We'll prove the Theorem by showing that, given x in $[a, b]$, there is a ξ in (a, b) with

$$f(x) - p(x) = \frac{1}{(n+1)!}\, f^{(n+1)}(\xi)\, W(x). \qquad (*)$$

If x is one of the x_i's, then both sides of this formula are 0 and we're done. Otherwise, $W(x) \ne 0$ and we may set $\lambda = [f(x) - p(x)]/W(x)$. Now consider

$$\varphi(t) = f(t) - p(t) - \lambda W(t).$$

Clearly, $\varphi(x_i) = 0$ for each $i = 0, 1, \ldots, n$ and, by our choice of λ, we also have $\varphi(x) = 0$. Here comes Rolle's theorem! Since φ has $n+2$ distinct zeros in $[a, b]$, we must have $\varphi^{(n+1)}(\xi) = 0$ for some ξ in (a, b). (Why?) Hence,

$$\begin{aligned} 0 = \varphi^{(n+1)}(\xi) &= f^{(n+1)}(\xi) - p^{(n+1)}(\xi) - \lambda W^{(n+1)}(\xi) \\ &= f^{(n+1)}(\xi) - \left(\frac{f(x) - p(x)}{W(x)}\right) \cdot (n+1)! \end{aligned}$$

because p has degree at most n and W is monic and degree $n+1$. $\square$

Observations

1. Equation $(*)$ is called the *Lagrange formula with remainder*. [Compare this result to Taylor's formula with remainder.]

2. The term $f^{(n+1)}(\xi)$ is actually a continuous function of x. That is, $[f(x)-p(x)]/W(x)$ is continuous; its value at an x_i is $[f'(x_i) - p'(x_i)]/W'(x_i)$ (why?) and $W'(x_i) = \prod_{j \ne i}(x_i - x_j) \ne 0$.

3. On *any* interval $[a, b]$, using *any* nodes, the sequence of Lagrange interpolating polynomials for e^x converge uniformly to e^x. In this case,

$$\|e^x - L_n(e^x)\| \le \frac{c}{(n+1)!}\,(b-a)^n$$

where $c = \|e^x\|$ in $C[\,a,b\,]$. The same would hold true for any infinitely differentiable function satisfying, say, $\|f^{(n)}\| \le M^n$ (any entire function, for example).

4. On $[-1,1]$, the norm of $\prod_{i=1}^{n}(x-x_i)$ is minimized by taking $x_i = \cos((2i-1)\pi/2n)$, the zeros of the n-th Chebyshev polynomial T_n. (Why?) As Rivlin points out, the zeros of the Chebyshev polynomials are a *nearly optimal* choice for the nodes if good uniform approximation is desired.

The question of convergence of interpolation is actually very closely related to the analogous question for the convergence of Fourier series—and the answer here is nearly the same. We'll have more to say about this analogy later. First, let's note that L_n is continuous (bounded); this will give us our first bit of insight into Faber's negative result.

Lemma. $\|L_n(f)\| \le \|f\|\,\left\|\sum_{i=0}^{n}|\ell_i(x)|\,\right\|$ *for any* $f \in C[\,a,b\,]$.

PROOF. **Exercise.**

The numbers $\Lambda_n = \left\|\sum_{i=0}^{n}|\ell_i(x)|\,\right\|$ are called the *Lebesgue numbers* associated to this process. It's not hard to see that Λ_n is the smallest possible constant that will work in this inequality (in other words, $\|L_n\| = \Lambda_n$). Indeed, if

$$\left\|\sum_{i=0}^{n}|\ell_i(x)|\right\| = \sum_{i=0}^{n}|\ell_i(x_0)|,$$

then we can find an $f \in C[\,a,b\,]$ with $\|f\| = 1$ and $f(x_i) = \operatorname{sgn}(\ell_i(x_0))$ for all i. (How?) Then,

$$\|L_n(f)\| \ge |L_n(f)(x_0)| = \left|\sum_{i=0}^{n}\operatorname{sgn}(\ell_i(x_0))\,\ell_i(x_0)\right| = \sum_{i=0}^{n}|\ell_i(x_0)| = \Lambda_n\|f\|.$$

As it happens, $\Lambda_n \ge c\,\log n$ (this is where the hard work comes in; see Rivlin or Natanson for further details), and, in particular, $\Lambda_n \to \infty$ as $n \to \infty$.

A simple application of the triangle inequality will allow us to bring $E_n(f)$ back into the picture:

Lemma. (Lebesgue's theorem) $\|f - L_n(f)\| \le (1 + \Lambda_n)\, E_n(f)$, *for any* $f \in C[\,a, b\,]$.

PROOF. Let p^* be the best approximation to f out of $\mathcal{P}_n$. Then, since $L_n(p^*) = p^*$, we have

$$\begin{aligned} \|f - L_n(f)\| &\le \|f - p^*\| + \|L_n(f - p^*)\| \\ &\le (1 + \Lambda_n)\,\|f - p^*\| \;=\; (1 + \Lambda_n)\, E_n(f). \end{aligned} \qquad \square$$

Appendix

Although we won't need anything quite so fancy, it is of some interest to discuss more general problems of interpolation. We again suppose that we are given distinct points $x_0 < \cdots < x_n$ in $[\,a, b\,]$, but now we suppose that we are given an array of information

$$\begin{array}{ccccc} y_0 & y_0' & y_0'' & \ldots & y_0^{(m_0)} \\ y_1 & y_1' & y_1'' & \ldots & y_1^{(m_1)} \\ \vdots & \vdots & \vdots & & \vdots \\ y_n & y_n' & y_n'' & \ldots & y_n^{(m_n)}, \end{array}$$

where each m_i is a nonnegative integer. Our problem is to find the polynomial p of *least degree* that incorporates all of this data by satisfying

$$\begin{array}{llll} p(x_0) = y_0 & p\,'(x_0) = y_0' & \ldots & p^{(m_0)}(x_0) = y_0^{(m_0)} \\ p(x_1) = y_1 & p\,'(x_1) = y_1' & \ldots & p^{(m_1)}(x_1) = y_1^{(m_1)} \\ \quad\vdots & \quad\vdots & & \quad\vdots \\ p(x_n) = y_n & p\,'(x_n) = y_n' & \ldots & p^{(m_n)}(x_n) = y_n^{(m_n)}. \end{array}$$

In other words, we specify not only the value of p at each x_i, but also the first m_i derivatives of p at x_i. This is often referred to as the problem of *Hermite interpolation.*

Since the problem has a total of $m_0 + m_1 + \cdots + m_n + n + 1$ "degrees of freedom," it won't come as any surprise that is has a (unique) solution p of degree (at most) $N = m_0 + m_1 + \cdots + m_n + n$. Rather than discuss this particular problem any further, let's instead discuss the general problem of *linear interpolation.*

The notational framework for our problem is an n-dimensional vector space X on which m linear, real-valued functions (or *linear functionals*) $L_1, \ldots, L_m$ are defined. The general problem of linear interpolation asks whether the system of equations

$$L_i(f) = y_i, \quad i = 1, \ldots, m \qquad (*)$$

has a (unique) solution $f \in X$ for any given set of scalars $y_1, \ldots, y_m \in \mathbb{R}$. Since a linear functional is completely determined by its values on any basis, we would next be led to consider a basis $f_1, \ldots, f_n$ for X, and from here it is a small step to rewrite $(*)$ as a *matrix equation*. That is, we seek a solution $f = a_1 f_1 + \cdots + a_n f_n$ satisfying

$$\begin{aligned} a_1 L_1(f_1) + \cdots + a_n L_1(f_n) &= y_1 \\ a_1 L_2(f_1) + \cdots + a_n L_2(f_n) &= y_2 \\ &\vdots \\ a_1 L_m(f_1) + \cdots + a_n L_m(f_n) &= y_m. \end{aligned}$$

If we are to guarantee a solution $a_1, \ldots, a_n$ for each choice of $y_1, \ldots, y_m$, then we'll need to have $m = n$ and, moreover, the matrix $[L_i(f_j)]$ will have to be nonsingular.

Lemma. *Let X be an n-dimensional vector space with basis $f_1, \ldots, f_n$, and let $L_1, \ldots, L_n$ be linear functionals on X. Then, $L_1, \ldots, L_n$ are linearly independent if and only if the matrix $[L_i(f_j)]$ is nonsingular; that is, if and only if $\det\big(L_i(f_j)\big) \neq 0$.*

Proof. If $[L_i(f_j)]$ is singular, then the matrix equation

$$\begin{aligned} c_1 L_1(f_1) + \cdots + c_n L_n(f_1) &= 0 \\ c_1 L_1(f_2) + \cdots + c_n L_n(f_2) &= 0 \\ &\vdots \\ c_1 L_1(f_n) + \cdots + c_n L_n(f_n) &= 0 \end{aligned}$$

has a nontrivial solution $c_1, \ldots, c_n$. Thus, the functional $c_1 L_1 + \cdots + c_n L_n$ satisfies

$$(c_1 L_1 + \cdots + c_n L_n)(f_i) = 0, \quad i = 1, \ldots, n.$$

Since $f_1, \ldots, f_n$ form a basis for X, this means that

$$(c_1 L_1 + \cdots + c_n L_n)(f) = 0$$

for all $f \in X$. That is, $c_1 L_1 + \cdots + c_n L_n = 0$ (the zero functional), and so $L_1, \ldots, L_n$ are linearly dependent.

Conversely, if $L_1, \ldots, L_n$ are linearly dependent, just reverse the steps in the first part of the proof to see that $[L_i(f_j)]$ is singular. $\square$

Theorem. *Let X be an n-dimensional vector space and let $L_1, \ldots, L_n$ be linear functionals on X. Then, the interpolation problem*

$$L_i(f) = y_i, \quad i = 1, \ldots, n \tag{$*$}$$

always has a (unique) solution $f \in X$ for any choice of scalars $y_1, \ldots, y_n$ if and only if $L_1, \ldots, L_n$ are linearly independent.

Proof. Let $f_1, \ldots, f_n$ be a basis for X. Then, $(*)$ is equivalent to the system of equations

$$\begin{aligned} a_1 L_1(f_1) + \cdots + a_n L_1(f_n) &= y_1 \\ a_1 L_2(f_1) + \cdots + a_n L_2(f_n) &= y_2 \\ &\vdots \\ a_1 L_n(f_1) + \cdots + a_n L_n(f_n) &= y_n \end{aligned} \tag{$**$}$$

by taking $f = a_1 f_1 + \cdots + a_n f_n$. Thus, $(*)$ always has a solution if and only if $(**)$ always has a solution if and only if $[L_i(f_j)]$ is nonsingular if and only if $L_1, \ldots, L_n$ are linearly independent. In any of these cases, note that the solution must be unique. $\square$

In the case of Lagrange interpolation, $X = \mathcal{P}_n$ and L_i is evaluation at x_i; i.e., $L_i(f) = f(x_i)$, which is easily seen to be linear in f. Moreover, $L_0, \ldots, L_n$ are linearly independent provided that $x_0, \ldots, x_n$ are distinct. (Why?)

In the case of Hermite interpolation, the linear functionals are of the form $L_{x,k}(f) = f^{(k)}(x)$, differentiation composed with a point evaluation. If $x \neq y$, then $L_{x,k}$ and $L_{y,m}$ are independent for any k and m; if $k \neq m$, then $L_{x,k}$ and $L_{x,m}$ are independent. (How would you check this?)

Throughout, $x_0, x_1, \ldots, x_n$ are distinct points in some interval $[\,a,b\,]$, and $V(x_0, x_1, \ldots, x_n)$ denotes the *Vandermonde determinant*:

$$V(x_0, x_1, \ldots, x_n) \;=\; \begin{vmatrix} 1 & x_0 & x_0^2 & \cdots & x_0^n \\ 1 & x_1 & x_1^2 & \cdots & x_1^n \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 1 & x_n & x_n^2 & \cdots & x_n^n \end{vmatrix}.$$

▷ ∗ **51.** Show by induction that $V(x_0, x_1, \ldots, x_n) = \displaystyle\prod_{0 \le j < i \le n} (x_i - x_j)$.

[Hint: In order to reduce to the $n \times n$ case, replace $\mathbf{c}_j$, the j-th column, by $\mathbf{c}_j - x_0\mathbf{c}_{j-1}$, starting on the right with $j = n$. Factor and use the induction hypothesis.]

52. Let $y_0, y_1, \ldots, y_n \in \mathbb{R}$ be given. Show that the polynomial $p \in \mathcal{P}_n$ satisfying $p(x_i) = y_i$, $i = 0, 1, \ldots, n$, may be written as

$$p(x) \;=\; c\begin{vmatrix} 0 & 1 & x & x^2 & \cdots & x^n \\ y_0 & 1 & x_0 & x_0^2 & \cdots & x_0^n \\ \vdots & \vdots & \vdots & \ddots & \vdots & \\ y_n & 1 & x_n & x_n^2 & \cdots & x_n^n \end{vmatrix},$$

where c is a certain constant. Find c and prove the formula.

53. Given $f \in C[\,a,b\,]$, let $L_n(f)$ denote the polynomial of degree at most n that agrees with f at the x_i's. Prove that L_n is a linear projection onto $\mathcal{P}_n$. That is, show that $L_n(\alpha f + \beta g) = \alpha L_n(f) + \beta L_n(g)$, and that $L_n(f) = f$ if and only if $f \in \mathcal{P}_n$.

54. Let $\ell_i(x)$, $i = 0, 1, \ldots, n$, denote the Lagrange interpolating polynomials of degree at most n associated with the nodes $x_0, x_1, \ldots, x_n$; that is, $\ell_i(x_j) = \delta_{i,j}$. Show that $\sum_{i=0}^n \ell_i(x) \equiv 1$ and, more generally, that $\sum_{i=0}^n x_i^k \ell_i(x) = x^k$, for $k = 0, 1, \ldots, n$.

55. If ℓ_i and L_n are as above, show that the error in the Lagrange interpolation formula is $(L_n(f) - f)(x) = \sum_{i=0}^n [\, f(x_i) - f(x)\,]\,\ell_i(x)$.

56. With ℓ_i and L_n as above, show that $\|L_n(f)\| \le \Lambda_n \|f\|$, where $\Lambda_n = \big\| \sum_{i=0}^n |\ell_i(x)| \,\big\|$. Show that no smaller number Λ has this property for all $f \in C[\,a,b\,]$.

In this section we consider a question of computational interest: Since best approximations are often very hard to find, how might we *approximate* the best approximation? The answer to this question lies in approximations over *finite* sets. Here's the plan:

(1) Fix a finite subset X_m of $[\,a,b\,]$ consisting of m distinct points $a \le x_1 < \cdots < x_m \le b$, and find the best approximation to f out of $\mathcal{P}_n$ *considered as a subspace of* $C(X_m)$. In other words, if we call the best approximation $p_n^*(X_m)$, then

$$\max_{1\le i\le m} |f(x_i) - p_n^*(X_m)(x_i)| = \min_{p\in\mathcal{P}_n} \max_{1\le i\le m} |f(x_i) - p(x_i)| \equiv E_n(f; X_m).$$

(2) Argue that this process converges (in some sense) to the best approximation on all of $[\,a,b\,]$ provided that X_m "gets big" as $m \to \infty$. In actual practice, there's no need to worry about $p_n^*(X_m)$ converging to p_n^* (the best approximation on all of $[\,a,b\,]$); rather, we will argue that $E_n(f; X_m) \to E_n(f)$ and appeal to "abstract nonsense."

(3) Find an efficient *strategy* for carrying out items (1) and (2).

Observations

1. If $m \le n+1$, then $E_n(f; X_m) = 0$. That is, we can always find a polynomial $p \in \mathcal{P}_n$ that agrees with f at $n+1$ (or fewer) points. (How?) Of course, p won't be unique if $m < n+1$. (Why?) In any case, we might as well assume that $m \ge n+2$. In fact, as we'll see, the case $m = n+2$ is all that we really need to worry about.

2. If $X \subset Y \subset [\,a,b\,]$, then $E_n(f; X) \le E_n(f; Y) \le E_n(f)$. Indeed, if $p \in \mathcal{P}_n$ is the best approximation on Y, then

$$E_n(f; X) \le \max_{x\in X} |f(x) - p(x)| \le \max_{x\in Y} |f(x) - p(x)| = E_n(f; Y).$$

Consequently, we expect $E_n(f; X_m)$ to *increase* to $E_n(f)$ as X_m "gets big."

Now if we were to repeat our earlier work on characterizing best approximations, restricting ourselves to X_m everywhere, here's what we'd get:

Theorem. *Let $m \geq n+2$. Then,*

(i) *$p \in \mathcal{P}_n$ is a best approximation to f on X_m if and only if $f-p$ has an alternating set containing $n+2$ points* **out of** *X_m; that is, $f - p = \pm E_n(f; X_m)$, alternately, on X_m.*

(ii) *$p_n^*(X_m)$ is unique.*

Next let's see how this reduces our study to the case $m = n+2$.

Theorem. *Fix n, $m \geq n+2$, and $f \in C[\,a,b\,]$.*

(i) *If $p_n^* \in \mathcal{P}_n$ is best on all of $[\,a,b\,]$, then there is a subset X_{n+2}^* of $[\,a,b\,]$, containing $n+2$ points, such that $p_n^* = p_n^*(X_{n+2}^*)$. Moreover, $E_n(f; X_{n+2}) \leq E_n(f) = E_n(f; X_{n+2}^*)$ for any other subset X_{n+2} of $[\,a,b\,]$, with equality if and only if $p_n^*(X_{n+2}) = p_n^*$.*

(ii) *If $p_n^*(X_m) \in \mathcal{P}_n$ is best on X_m, then there is a subset X_{n+2}^* of X_m such that $p_n^*(X_m) = p_n^*(X_{n+2}^*)$ and $E_n(f; X_m) = E_n(f; X_{n+2}^*)$. For any other $X_{n+2} \subset X_m$ we have $E_n(f; X_{n+2}) \leq E_n(f; X_{n+2}^*) = E_n(f; X_m)$, with equality if and only if $p_n^*(X_{n+2}) = p_n^*(X_m)$.*

Proof. (i): Let X_{n+2}^* be an alternating set for $f - p_n^*$ over $[\,a,b\,]$ containing exactly $n+2$ points. Then, X_{n+2}^* is also an alternating set for $f - p_n^*$ over X_{n+2}^*. That is, for $x \in X_{n+2}^*$,

$$\pm(f(x) - p_n^*(x)) = E_n(f) = \max_{y \in X_{n+2}^*} |f(y) - p_n^*(y)|.$$

So, by uniqueness of best approximations on X_{n+2}^*, we must have $p_n^* = p_n^*(X_{n+2}^*)$ and $E_n(f) = E_n(f; X_{n+2}^*)$. The second assertion follows from a similar argument using the uniqueness of p_n^* on $[\,a,b\,]$.

(ii): This is just (i) with $[\,a,b\,]$ replaced everywhere by X_m. $\square$

Here's the point: Through some as yet undisclosed method, we choose X_m with $m \geq n+2$ (in fact, $m >> n+2$) such that $E_n(f; X_m) \leq E_n(f) \leq E_n(f; X_m) + \varepsilon$, and then we search for the "best" $X_{n+2} \subset X_m$, meaning the *largest* value of $E_n(f; X_{n+2})$. We then take $p_n^*(X_{n+2})$ as an approximation for p_n^*. As we'll see momentarily, $p_n^*(X_{n+2})$ can be computed directly and explicitly.

Now suppose that the elements of X_{n+2} are $a \le x_0 < x_1 < \cdots < x_{n+1} \le b$, let $p = p_n^*(X_{n+2})$ be $p(x) = a_0 + a_1 x + \cdots + a_n x^n$, and let

$$E = E_n(f; X_{n+2}) = \max_{0 \le i \le n+1} |f(x_i) - p(x_i)|.$$

In order to compute p and E, we use the fact that $f(x_i) - p(x_i) = \pm E$, alternately, and write (for instance)

$$\begin{aligned} f(x_0) &= E + p(x_0) \\ f(x_1) &= -E + p(x_1) \\ &\vdots \\ f(x_{n+1}) &= (-1)^{n+1} E + p(x_{n+1}) \end{aligned}$$

(where the "E column" might, instead, read $-E$, E, ..., $(-1)^n E$). That is, in order to find p and E, we need to solve a system of $n+2$ *linear* equations in the $n+2$ unknowns E, $a_0, \ldots, a_n$. The determinant of this system is (up to sign)

$$\begin{vmatrix} 1 & 1 & x_0 & \cdots & x_0^n \\ -1 & 1 & x_1 & \cdots & x_1^n \\ \vdots & \vdots & & \ddots & \vdots \\ (-1)^{n+1} & 1 & x_n & \cdots & x_n^n \end{vmatrix} = A_0 + A_1 + \cdots + A_{n+1} > 0,$$

where we have expanded by cofactors along the first column and have used the fact that each minor A_k is a Vandermonde determinant (and hence each $A_k > 0$). If we apply Cramer's rule to find E we get

$$\begin{aligned} E &= \frac{f(x_0)A_0 - f(x_1)A_1 + \cdots + (-1)^{n+1} f(x_{n+1})A_{n+1}}{A_0 + A_1 + \cdots + A_{n+1}} \\ &= \lambda_0 f(x_0) - \lambda_1 f(x_1) + \cdots + (-1)^{n+1} \lambda_{n+1} f(x_{n+1}), \end{aligned}$$

where $\lambda_i > 0$ and $\sum_{i=0}^{n+1} \lambda_1 = 1$. Moreover, these same λ_i's satisfy $\sum_{i=0}^{n+1} (-1)^i \lambda_i q(x_i) = 0$ for every polynomial $q \in \mathcal{P}_n$ since $E = E_n(q; X_{n+2}) = 0$ for polynomials of degree at most n (and since Cramer's rule supplies the same coefficients for all f's).

It may be instructive to see a more explicit solution to this problem. For this, recall that since we have $n+2$ points we may interpolate exactly out of $\mathcal{P}_{n+1}$. Given this, our original problem can be rephrased quite succinctly.

Let p be the (unique) polynomial in $\mathcal{P}_{n+1}$ satisfying $p(x_i) = f(x_i)$, $i = 0, 1, \ldots, n+1$, and let e be the (unique) polynomial in $\mathcal{P}_{n+1}$ satisfying $e(x_i) = (-1)^i$, $i = 0, 1, \ldots, n+1$. If it is possible to find a scalar λ so that $p - \lambda e \in \mathcal{P}_n$, then $p - \lambda e = p_n^*(X_{n+2})$ and $|\lambda| = E_n(f; X_{n+2})$. Why? Because $f - (p - \lambda e) = \lambda e = \pm\lambda$, alternately, on X_{n+2} and so $|\lambda| = \max\limits_{x \in X_{n+2}} |f(x) - (p(x) - \lambda e(x))|$. Thus, we need to compare leading coefficients of p and e.

Now if p has degree less than $n+1$, then $p = p_n^*(X_{n+2})$ and $E_n(f; X_{n+2}) = 0$. Thus, $\lambda = 0$ would do nicely in this case. Otherwise, p has degree exactly $n+1$ and the question is whether e does too. Now,

$$e(x) = \sum_{i=0}^{n+1} \frac{(-1)^i}{W'(x_i)} \cdot \frac{W(x)}{(x - x_i)},$$

where $W(x) = \prod_{i=0}^{n+1}(x - x_i)$, and so the leading coefficient of e is $\sum_{i=0}^{n+1}(-1)^i/W'(x_i)$. We'll be done if we can convince ourselves that this is nonzero. But

$$W'(x_i) \;=\; \prod_{j \neq i}(x_i - x_j) \;=\; (-1)^{n-i+1} \prod_{j=0}^{i-1}(x_i - x_j) \prod_{j=i+1}^{n+1}(x_j - x_i),$$

hence $(-1)^i/W'(x_i)$ is of constant sign $(-1)^{n+1}$. Finally, since

$$p(x) = \sum_{i=0}^{n+1} \frac{f(x_i)}{W'(x_i)} \cdot \frac{W(x)}{(x - x_i)},$$

p has leading coefficient $\sum_{i=0}^{n+1} f(x_i)/W'(x_i)$ and it's easy to find the value of λ.

Conclusion. *$p_n^*(X_{n+2}) = p - \lambda e$, where*

$$\lambda \;=\; \frac{\sum_{i=0}^{n+1} f(x_i)/W'(x_i)}{\sum_{i=0}^{n+1}(-1)^i/W'(x_i)} \;=\; \sum_{i=0}^{n+1}(-1)^i \lambda_i f(x_i)$$

and

$$\lambda_i = \frac{1/|W'(x_i)|}{\sum_{j=0}^{n+1} 1/|W'(x_j)|},$$

and $|\lambda| = E_n(f; X_{n+2})$. Moreover, $\sum_{i=0}^{n+1}(-1)^i \lambda_i q(x_i) = 0$ for every $q \in \mathcal{P}_n$.

Example

Find the best linear approximation to $f(x) = x^2$ on $X_4 = \{0, 1/3, 2/3, 1\} \subset [\,0, 1\,]$.

We seek $p(x) = a_0 + a_1 x$ and we need only consider subsets of X_4 of size $1 + 2 = 3$. There are four:

$$X_{4,1} = \{0, 1/3, 2/3\},\ X_{4,2} = \{0, 1/3, 1\},\ X_{4,3} = \{0, 2/3, 1\},\ X_{4,4} = \{1/3, 2/3, 1\}.$$

In each case we find a p and a λ ($= E$ in our earlier setup). For instance, in the case of $X_{4,2}$ we would solve the system of equations $f(x) = \pm\lambda + p(x)$ for $x = 0, 1/3, 1$.

$$\left.\begin{aligned} 0 &= \lambda^{(2)} + a_0 \\ \frac{1}{9} &= -\lambda^{(2)} + a_0 + \frac{1}{3}a_1 \\ 1 &= \lambda^{(2)} + a_0 + a_1 \end{aligned}\right\} \Longrightarrow \qquad \begin{aligned} \lambda^{(2)} &= \frac{1}{9} \\ a_0 &= -\frac{1}{9} \\ a_1 &= 1 \end{aligned}$$

In the other three cases you would find that $\lambda^{(1)} = 1/18$, $\lambda^{(3)} = 1/9$, and $\lambda^{(4)} = 1/18$. Since we need the largest λ, we're done: $X_{4,2}$ (or $X_{4,3}$) works, and $p_1^*(X_4)(x) = x - 1/9$. (Recall that the best approximation on all of $[\,0, 1\,]$ is $p_1^*(x) = x - 1/8$.)

Where does this leave us? We still need to know that there is some hope of finding an initial set X_m with $E_n(f) - \varepsilon \le E_n(f; X_m) \le E_n(f)$, and we need a more efficient means of searching through the $\binom{m}{n+2}$ subsets $X_{n+2} \subset X_m$. In order to attack the problem of finding an initial X_m, we'll need a few classical inequalities. We won't directly attack the second problem; instead, we'll outline an algorithm that begins with an initial set X_{n+2}^0, containing exactly $n+2$ points, which is then "improved" to some X_{n+2}^1 by changing only a single point.

The Inequalities of Markov and Bernstein

In order to discuss the convergence of approximations over finite sets, we will need to know that differentiation is bounded on $\mathcal{P}_n$ (a fact that is nearly obvious by itself).

The inequality we'll use is due to A. A. Markov from 1889:

Theorem. *If $p \in \mathcal{P}_n$, and if $|p(x)| \le 1$ for $|x| \le 1$, then $|p'(x)| \le n^2$ for $|x| \le 1$. Moreover, $|p'(x)| = n^2$ can only occur at $x = \pm 1$, and only when $p = \pm T_n$, the Chebyshev polynomial of degree n.*

Markov's brother, V. A. Markov, later improved on this, in 1916, by showing that $|p^{(k)}(x)| \le T_n^{(k)}(1)$. We've alluded to this fact already (see Rivlin, p. 31), and even more is true. For our purposes, it's enough to have *some* bound on differentiation; in particular, we'll only use

$$\|p'\| \le n^2 \|p\| \qquad \text{and} \qquad \|p''\| \le n^4 \|p\|,$$

where $\|\cdot\|$ is the norm in $C[-1,1]$.

About 20 years after Markov, in 1912, Bernstein asked for a similar bound for the derivative of a complex polynomial over the unit disk $|z| \le 1$. Now the maximum modulus theorem tells us that we may reduce to the case $|z| = 1$, that is, $z = e^{i\theta}$, and so Bernstein was able to restate the problem in terms of trig polynomials.

Theorem. *If $S \in \mathcal{T}_n$, and if $|S(\theta)| \le 1$, then $|S'(\theta)| \le n$. Equality is only possible for $S(\theta) = \sin n(\theta - \theta_0)$.*

Our plan is to deduce Markov's inequality from Bernstein's inequality by a method of proof due to Pólya and Szegö in 1928. To begin, let's consider the Lagrange interpolation formula in the case where $x_i = \cos((2i-1)\pi/2n)$, $i = 1, \ldots, n$, are the zeros of the Chebyshev polynomial T_n. Recall that we have $-1 < x_n < x_{n-1} < \cdots < x_1 < 1$.

Lemma 1. *Each polynomial $p \in \mathcal{P}_{n-1}$ may be written*

$$p(x) \;=\; \frac{1}{n} \sum_{i=1}^{n} p(x_i) \cdot (-1)^{i-1} \sqrt{1 - x_i^2} \cdot \frac{T_n(x)}{x - x_i}.$$

Proof. We know that the Lagrange interpolation formula is exact for polynomials of degree $< n$, and we know that, up to a constant multiple, $T_n(x)$ is the product $W(x) = (x - x_1) \cdots (x - x_n)$. All that remains is to compute $T_n'(x_i)$. But recall that for $x = \cos\theta$ we have

$$T_n'(x) = \frac{n \sin n\theta}{\sin\theta} = \frac{n \sin n\theta}{\sqrt{1 - \cos^2\theta}} = \frac{n \sin n\theta}{\sqrt{1 - x^2}}.$$

But for $x_i = \cos((2i-1)\pi/2n)$, i.e., for $\theta_i = (2i-1)\pi/2n$, it follows that $\sin n\theta_i = \sin((2i-1)\pi/2) = (-1)^{i-1}$; that is,

$$\frac{1}{T_n'(x_i)} = \frac{(-1)^{i-1}\sqrt{1-x_i^2}}{n}. \qquad \square$$

Lemma 2. *For any polynomial $p \in \mathcal{P}_{n-1}$, we have*

$$\max_{-1\le x\le 1} |p(x)| \le \max_{-1\le x\le 1} \left| n\sqrt{1-x^2}\, p(x)\right|.$$

Proof. To save wear and tear, let's write $M = \max_{-1\le x\le 1} \left| n\sqrt{1-x^2}\, p(x)\right|$.

First consider an x in the interval $[\,x_n, x_1\,]$; that is, $|x| \le \cos(\pi/2n) = x_1$. In this case we can estimate $\sqrt{1-x^2}$ from below:

$$\sqrt{1-x^2} \ge \sqrt{1-x_1^2} = \sqrt{1-\cos^2\left(\frac{\pi}{2n}\right)} = \sin\left(\frac{\pi}{2n}\right) \ge \frac{1}{n},$$

because $\sin\theta \ge 2\theta/\pi$ for $0 \le \theta \le \pi/2$ (from the mean value theorem). Hence, for $|x| \le \cos(\pi/2n)$, we get $|p(x)| \le n\sqrt{1-x^2}\,|p(x)| \le M$.

Now, for x's outside the interval $[\,x_n, x_1\,]$, we apply our interpolation formula. In this case, each of the factors $x - x_i$ is of the *same sign*. Thus,

$$\begin{aligned} |p(x)| &= \frac{1}{n}\left|\sum_{i=1}^{n} p(x_i)\,\frac{(-1)^{i-1}\sqrt{1-x_i^2}\;T_n(x)}{x-x_i}\right| \\ &\le \frac{M}{n^2}\sum_{i=1}^{n}\left|\frac{T_n(x)}{x-x_i}\right| = \frac{M}{n^2}\left|\sum_{i=1}^{n}\frac{T_n(x)}{x-x_i}\right|. \end{aligned}$$

But,

$$\sum_{i=1}^{n}\frac{T_n(x)}{x-x_i} = T_n'(x) \qquad \text{(why?)}$$

and we know that $|T_n'(x)| \le n^2$. Thus, $|p(x)| \le M$. $\square$

We next turn our attention to trig polynomials. As usual, given an algebraic polynomial $p \in \mathcal{P}_n$, we will sooner or later consider $S(\theta) = p(\cos\theta)$. In this case, $S'(\theta) = p'(\cos\theta)\sin\theta$ is an odd trig polynomial of degree at most n and $|S'(\theta)| = |p'(\cos\theta)\sin\theta| =$

$|p'(x)\sqrt{1-x^2}\,|$. Conversely, if $S \in \mathcal{T}_n$ is an odd trig polynomial, then $S(\theta)/\sin\theta$ is even, and so may be written $S(\theta)/\sin\theta = p(\cos\theta)$ for some algebraic polynomial p of degree at most $n-1$. From Lemma 2,

$$\max_{0\le\theta\le 2\pi}\left|\frac{S(\theta)}{\sin\theta}\right| \;=\; \max_{0\le\theta\le 2\pi}|p(\cos\theta)| \;\le\; n\max_{0\le\theta\le 2\pi}|p(\cos\theta)\sin\theta| \;=\; n\max_{0\le\theta\le 2\pi}|S(\theta)|.$$

This proves

Corollary. *If $S \in \mathcal{T}_n$ is an odd trig polynomial, then*

$$\max_{0\le\theta\le 2\pi}\left|\frac{S(\theta)}{\sin\theta}\right| \;\le\; n\max_{0\le\theta\le 2\pi}|S(\theta)|.$$

Now we're ready for Bernstein's inequality.

Bernstein's Inequality. *If $S \in \mathcal{T}_n$, then*

$$\max_{0\le\theta\le 2\pi}|S'(\theta)| \;\le\; n\max_{0\le\theta\le 2\pi}|S(\theta)|.$$

Proof. We first define an auxiliary function $f(\alpha,\theta) = \big[S(\alpha+\theta) - S(\alpha-\theta)\big]/2$. For α fixed, $f(\alpha,\theta)$ is an odd trig polynomial in θ of degree at most n. Consequently,

$$\left|\frac{f(\alpha,\theta)}{\sin\theta}\right| \;\le\; n\max_{0\le\theta\le 2\pi}|f(\alpha,\theta)| \;\le\; n\max_{0\le\theta\le 2\pi}|S(\theta)|.$$

But

$$S'(\alpha) \;=\; \lim_{\theta\to 0}\frac{S(\alpha+\theta)-S(\alpha-\theta)}{2\theta} \;=\; \lim_{\theta\to 0}\frac{f(\alpha,\theta)}{\sin\theta},$$

and hence $|S'(\alpha)| \le n\max\limits_{0\le\theta\le 2\pi}|S(\theta)|$. $\square$

Finally, we prove Markov's inequality.

Markov's Inequality. *If $p \in \mathcal{P}_n$, then*

$$\max_{-1\le x\le 1}|p'(x)| \;\le\; n^2\max_{-1\le x\le 1}|p(x)|.$$

PROOF. We know that $S(\theta) = p(\cos\theta)$ is a trig polynomial of degree at most n satisfying

$$\max_{-1\le x\le 1} |p(x)| \;=\; \max_{0\le\theta\le 2\pi} |p(\cos\theta)|.$$

Since $S'(\theta) = p'(\cos\theta)\sin\theta$ is also trig polynomial of degree at most n, Bernstein's inequality yields

$$\max_{0\le\theta\le 2\pi} |p'(\cos\theta)\sin\theta| \;\le\; n \max_{0\le\theta\le 2\pi} |p(\cos\theta)|.$$

In other words,

$$\max_{-1\le x\le 1} \left|p'(x)\sqrt{1-x^2}\right| \;\le\; n \max_{-1\le x\le 1} |p(x)|.$$

Since $p' \in \mathcal{P}_{n-1}$, the desired inequality now follows easily from Lemma 2:

$$\max_{-1\le x\le 1} |p'(x)| \;\le\; n \max_{-1\le x\le 1} \left|p'(x)\sqrt{1-x^2}\right| \;\le\; n^2 \max_{-1\le x\le 1} |p(x)|. \qquad \square$$

Convergence of Approximations over Finite Sets

In order to simplify things here, we will make several assumptions: For one, we will consider only approximation over the interval $I = [-1, 1]$. As before, we consider a fixed $f \in C[-1, 1]$ and a fixed integer $n = 0, 1, 2, \ldots$. For each integer $m \ge 1$ we choose a finite subset $X_m \subset I$, consisting of m points $-1 \le x_1 < \cdots < x_m \le 1$; in addition, we will assume that $x_1 = -1$ and $x_m = 1$. If we put

$$\delta_m \;=\; \max_{x\in I}\, \min_{1\le i\le m} |x - x_i| > 0,$$

then each $x \in I$ is within δ_m of some x_i. If X_m consists of equally spaced points, for example, it's easy to see that $\delta_m = 1/(m-1)$.

Our goal is to prove

Theorem. *If $\delta_m \to 0$, then $E_n(f; X_m) \to E_n(f)$.*

And we would hope to accomplish this in such a way that δ_m is a measurable quantity, depending on f, m, and a prescribed tolerance $\varepsilon = E_n(f; X_m) - E_n(f)$.

As a first step in this direction, let's bring Markov's inequality into the picture.

Lemma. *Suppose that* $\tau_m \equiv \delta_m^2 n^4/2 < 1$. *Then, for any* $p \in \mathcal{P}_n$, *we have*

$$(1) \qquad \max_{-1\le x\le 1} |p(x)| \;\le\; (1-\tau_m)^{-1} \max_{1\le i\le m} |p(x_i)|$$

and

$$(2) \qquad \omega_p([-1,1];\delta_m) \;\le\; \delta_m n^2 (1-\tau_m)^{-1} \max_{1\le i\le m} |p(x_i)|.$$

Proof. (1): Take a in $[-1,1]$ with $|p(a)| = \|p\|$. If $a = \pm 1 \in X_m$, we're done (since $(1-\tau_m)^{-1} > 1$). Otherwise, we'll have $-1 < a < 1$ and $p'(a) = 0$. Next, choose $x_i \in X_m$ with $|a - x_i| \le \delta_m$ and apply Taylor's theorem:

$$p(x_i) = p(a) + (x_i - a)\, p'(a) + \frac{(x_i - a)^2}{2}\, p''(c),$$

for some c in $(-1,1)$. Re-writing, we have

$$|p(a)| \le |p(x_i)| + \frac{\delta_m^2}{2}\, |p''(c)|.$$

And now we bring in Markov:

$$\|p\| \le \max_{1\le i\le m} |p(x_i)| \;+\; \frac{\delta_m^2 n^4}{2}\, \|p\|,$$

which is what we need.

(2): The real point here is that each $p \in \mathcal{P}_n$ is *Lipschitz* with constant $n^2\|p\|$. Indeed,

$$|p(s) - p(t)| = |(s-t)\, p'(c)| \le |s-t|\, \|p'\| \le n^2 \|p\|\, |s-t|$$

(from the mean value theorem and Markov's inequality). Thus, $\omega_p(\delta) \le \delta n^2 \|p\|$ and, combining this with (1), we get

$$\omega_p(\delta_m) \;\le\; \delta_m n^2 \|p\| \;\le\; \delta_m n^2 (1-\tau_m)^{-1} \max_{1\le i\le m} |p(x_i)|. \qquad \square$$

Now we're ready to compare $E_n(f; X_m)$ to $E_n(f)$. Our result won't be as good as Rivlin's (he uses a fancier version of Markov's inequality), but it will be a bit easier to prove. As in the Lemma, we'll suppose that

$$\tau_m = \frac{\delta_m^2 n^4}{2} < 1,$$

and we'll set

$$\Delta_m = \frac{\delta_m n^2}{1 - \tau_m}.$$

[Note that as $\delta_m \to 0$ we also have $\tau_m \to 0$ and $\Delta_m \to 0$.]

Theorem. *For $f \in C[-1,1]$,*

$$E_n(f;X_m) \ \le\ E_n(f) \ \le\ (1+\Delta_m)\,E_n(f;X_m) \ +\ \omega_f([-1,1];\delta_m) \ +\ \Delta_m\,\|f\|.$$

Consequently, if $\delta_m \to 0$, then $E_n(f;X_m) \to E_n(f)$ (as $m \to \infty$).

Proof. Let $p = p_n^*(X_m) \in \mathcal{P}_n$ be the best approximation to f on X_m. Recall that

$$\max_{1\le i\le m} |f(x_i) - p(x_i)| \ =\ E_n(f;X_m) \ \le\ E_n(f) \ \le\ \|f - p\|.$$

Our plan is to estimate $\|f - p\|$.

Let $x \in [-1,1]$ and choose $x_i \in X_m$ with $|x - x_i| \le \delta_m$. Then,

$$\begin{aligned} |f(x) - p(x)| &\le |f(x) - f(x_i)| + |f(x_i) - p(x_i)| + |p(x_i) - p(x)| \\ &\le \omega_f(\delta_m) \ +\ E_n(f;X_m) \ +\ \omega_p(\delta_m) \\ &\le \omega_f(\delta_m) \ +\ E_n(f;X_m) \ +\ \Delta_m \max_{1\le i\le m} |p(x_i)|, \end{aligned}$$

where we've used (2) from the previous Lemma to estimate $\omega_p(\delta_m)$. All that remains is to revise this last estimate, eliminating reference to p. For this we use the triangle inequality again:

$$\begin{aligned} \max_{1\le i\le m} |p(x_i)| &\le \max_{1\le i\le m} |f(x_i) - p(x_i)| \ +\ \max_{1\le i\le m} |f(x_i)| \\ &\le E_n(f;X_m) \ +\ \|f\|. \end{aligned}$$

Putting all the pieces together gives us our result:

$$E_n(f) \ \le\ \omega_f(\delta_m) \ +\ E_n(f;X_m) \ +\ \Delta_m\big[\,E_n(f;X_m) \ +\ \|f\|\,\big]. \qquad \square$$

As Rivlin points out, it is quite possible to give a lower bound on m in the case of, say, equally spaced points, which will give $E_n(f;X_m) \le E_n(f) \le E_n(f;X_m) + \varepsilon$, but this is surely an inefficient approach to the problem. Instead, we'll discuss the *one point exchange algorithm*.

The One Point Exchange Algorithm

We're given $f \in C[-1,1]$, n, and $\varepsilon > 0$.

1. Pick a starting "reference" X_{n+2}. A convenient choice is the set $x_i = \cos\left(\frac{n+1-i}{n+1}\,\pi\right)$, $i = 0, 1, \ldots, n+1$. These are the "peak points" of T_{n+1}; that is, $T_{n+1}(x_i) = (-1)^{n+1-i}$ (and so T_{n+1} is the polynomial e from our "Conclusion").

2. Find $p = p_n^*(X_{n+2})$ and λ (by solving a system of linear equations). Recall that
$$|\lambda| \;=\; |f(x_i) - p(x_i)| \;\le\; \|f - p^*\| \;\le\; \|f - p\|,$$
where p^* is the best approximation to f on all of $[-1, 1]$.

3. Find (approximately, if necessary) the "error function" $e(x) = f(x) - p(x)$ and any point η where $|f(\eta) - p(\eta)| = \|f - p\|$. (According to Powell, this can be accomplished using "local quadratic fits.")

4. Replace an appropriate x_i by η so that the new reference $X'_{n+2} = \{x'_1, x'_2, \ldots\}$ has the properties that $f(x'_i) - p(x'_i)$ alternates in sign, and that $|f(x'_i) - p(x'_i)| \ge |\lambda|$ for all i. The new polynomial $p' = p_n^*(X'_{n+2})$ and new λ' must then satisfy
$$|\lambda| \;=\; \min_{0\le i\le n+1} |f(x'_i) - p(x'_i)| \;\le\; \max_{0\le i\le n+1} |f(x'_i) - p'(x'_i)| \;=\; |\lambda'|.$$
This is an observation due to de la Vallée Poussin: Since $f - p$ alternates in sign on an alternating set for $f - p'$, it follows that $f - p'$ increases the minimum error over this set. (See the Theorem on page 53 of "Characterization of Best Approximation" for a precise statement.) Again according to Powell, the new p' and λ' can be found quickly through matrix "updating" techniques. (Since we've only changed one of the x_i's, only one row of the matrix on page 82 needs to be changed.)

5. The new λ' satisfies $|\lambda'| \le \|f - p^*\| \le \|f - p'\|$, and the calculation stops when
$$\|f - p'\| - |\lambda'| \;=\; |f(\eta') - p'(\eta')| - |\lambda'| \;<\; \varepsilon.$$

The *Fourier series* of a 2π-periodic (bounded, integrable) function f is

$$\frac{a_0}{2} + \sum_{k=1}^{\infty} \big(a_k \cos kx + b_k \sin kx \big),$$

where the coefficients are defined by

$$a_k = \frac{1}{\pi} \int_{-\pi}^{\pi} f(t) \, \cos kt \, dt \qquad \text{and} \qquad b_k = \frac{1}{\pi} \int_{-\pi}^{\pi} f(t) \, \sin kt \, dt.$$

Please note that if f is Riemann integrable on $[-\pi, \pi]$, then each of these integrals is well-defined and finite; indeed,

$$|a_k| \le \frac{1}{\pi} \int_{-\pi}^{\pi} |f(t)| \, dt$$

and so, for example, we would have $|a_k| \le 2\|f\|$ for $f \in C^{2\pi}$.

We write the partial sums of the series as

$$s_n(f)(x) = \frac{a_0}{2} + \sum_{k=1}^{n} \big(a_k \cos kx + b_k \sin kx \big).$$

Now while $s_n(f)$ need not converge pointwise to f (in fact, it may even diverge at a given point), and while $s_n(f)$ is not typically a good uniform approximation to f, it is still a very natural choice for an approximation to f in the "least-squares" sense (which we'll make precise shortly). Said in other words, the Fourier series for f provides a useful *representation* for f even if it fails to converge pointwise to f.

Observations

1. The functions $1, \cos x, \cos 2x, \ldots, \sin x, \sin 2x, \ldots$, are *orthogonal* on $[-\pi, \pi]$. That is,

$$\int_{-\pi}^{\pi} \cos mx \, \cos nx \, dx = \int_{-\pi}^{\pi} \sin mx \, \sin nx \, dx = \int_{-\pi}^{\pi} \cos mx \, \sin nx \, dx = 0$$

for any $m \ne n$ (and the last equation even holds for $m = n$),

$$\int_{-\pi}^{\pi} \cos^2 mx \, dx = \int_{-\pi}^{\pi} \sin^2 mx \, dx = \pi$$

for any $m \neq 0$, and, of course, $\int_{-\pi}^{\pi} 1\,dx = 2\pi$.

2. What this means is that if $T(x) = \frac{\alpha_0}{2} + \sum_{k=1}^{n}\big(\alpha_k \cos kx + \beta_k \sin kx\big)$, then

$$\frac{1}{\pi}\int_{-\pi}^{\pi} T(x)\,\cos mx\,dx = \frac{\alpha_m}{\pi}\int_{-\pi}^{\pi}\cos^2 mx\,dx = \alpha_m$$

for $m \neq 0$, while

$$\frac{1}{\pi}\int_{-\pi}^{\pi} T(x)\,dx = \frac{\alpha_0}{2\pi}\int_{-\pi}^{\pi} dx = \alpha_0.$$

That is, if $T \in \mathcal{T}_n$, then T is actually equal to its own Fourier series.

3. The partial sum operator $s_n(f)$ is a *linear projection* from $C^{2\pi}$ onto $\mathcal{T}_n$.

4. If $T(x) = \frac{\alpha_0}{2} + \sum_{k=1}^{n}\big(\alpha_k \cos kx + \beta_k \sin kx\big)$ is a trig polynomial, then

$$\begin{aligned}\frac{1}{\pi}\int_{-\pi}^{\pi} f(x)\,T(x)\,dx &= \frac{\alpha_0}{2\pi}\int_{-\pi}^{\pi} f(x)\,dx + \sum_{k=1}^{n}\frac{\alpha_k}{\pi}\int_{-\pi}^{\pi} f(x)\,\cos kx\,dx \\ &\qquad + \sum_{k=1}^{n}\frac{\beta_k}{\pi}\int_{-\pi}^{\pi} f(x)\,\sin kx\,dx \\ &= \frac{\alpha_0 a_0}{2} + \sum_{k=1}^{n}\big(\alpha_k a_k + \beta_k b_k\big),\end{aligned}$$

where (a_k) and (b_k) are the Fourier coefficients for f. [This should remind you of the *dot product* of the coefficients.]

5. Motivated by **1**, **2**, and **4**, we define the *inner product* of two elements $f,\ g \in C^{2\pi}$ by

$$\langle f, g\rangle = \frac{1}{\pi}\int_{-\pi}^{\pi} f(x)\,g(x)\,dx.$$

Note that from **4** we have $\langle f, s_n(f)\rangle = \langle s_n(f), s_n(f)\rangle$ for any n. (Why?)

6. If some $f \in C^{2\pi}$ has $a_k = b_k = 0$ for all k, then $f \equiv 0$.

Indeed, by **4** (or linearity of the integral), this means that

$$\int_{-\pi}^{\pi} f(x)\,T(x)\,dx = 0$$

for any trig polynomial T. But from Weierstrass's second theorem we know that f is the uniform limit of some sequence of trig polynomials (T_n). Thus,

$$\int_{-\pi}^{\pi} f(x)^2\,dx = \lim_{n\to\infty}\int_{-\pi}^{\pi} f(x)\,T_n(x)\,dx = 0.$$

Since f is continuous, this easily implies that $f \equiv 0$.

7. If $f,\ g \in C^{2\pi}$ have the same Fourier series, then $f \equiv g$. Hence, the Fourier series for an $f \in C^{2\pi}$ provides a *representation* for f (even if the series fails to converge to f).

8. The coefficients $a_0, a_1, \ldots, a_n$ and $b_1, b_2, \ldots, b_n$ minimize the expression

$$\varphi(a_0, a_1, \ldots, b_n) = \int_{-\pi}^{\pi} \big[f(x) - s_n(f)(x)\big]^2\,dx.$$

It's not hard to see, for example, that

$$\frac{\partial\varphi}{\partial a_k} = \int_{-\pi}^{\pi} 2\big[f(x) - s_n(f)(x)\big]\,\cos kx\,dx = 0$$

precisely when a_k satisfies

$$\int_{-\pi}^{\pi} f(x)\,\cos kx\,dx = a_k \int_{-\pi}^{\pi} \cos^2 kx\,dx.$$

9. The partial sum $s_n(f)$ is the best approximation to f out of $\mathcal{T}_n$ *relative to the L_2 norm*

$$\|f\|_2 \;=\; \sqrt{\langle f, f\rangle} \;=\; \left(\frac{1}{\pi}\int_{-\pi}^{\pi} f(x)^2\,dx\right)^{1/2}.$$

(Be forewarned: Some authors prefer $1/2\pi$ in place of $1/\pi$.) That is,

$$\|f - s_n(f)\|_2 = \min_{T\in\mathcal{T}_n} \|f - T\|_2.$$

Moreover, using **4** and **5**, we have

$$\begin{aligned}
\|f - s_n(f)\|_2^2 &= \langle f - s_n(f), f - s_n(f)\rangle \\
&= \langle f, f\rangle - 2\,\langle f, s_n(f)\rangle + \langle s_n(f), s_n(f)\rangle \\
&= \|f\|_2^2 - \|s_n(f)\|_2^2 \\
&= \frac{1}{\pi}\int_{-\pi}^{\pi} f(x)^2\,dx - \frac{a_0^2}{2} - \sum_{k=1}^{n}(a_k^2 + b_k^2).
\end{aligned}$$

[This should remind you of the *Pythagorean theorem*.]

10. It follows from **9** that

$$\frac{1}{\pi}\int_{-\pi}^{\pi} s_n(f)(x)^2\,dx \;=\; \frac{a_0^2}{2} + \sum_{k=1}^{n}\big(a_k^2 + b_k^2\big) \;\le\; \frac{1}{\pi}\int_{-\pi}^{\pi} f(x)^2\,dx.$$

In other symbols, $\|s_n(f)\|_2 \le \|f\|_2$. In particular, the Fourier coefficients of any $f \in C^{2\pi}$ are square summable. (Why?)

11. If $f \in C^{2\pi}$, then its Fourier coefficients (a_n) and (b_n) tend to zero as $n \to \infty$.

12. It follows from **10** and Weierstrass's second theorem that $s_n(f) \to f$ *in the* L_2 *norm* whenever $f \in C^{2\pi}$. Indeed, given $\varepsilon > 0$, choose a trig polynomial T such that $\|f - T\| < \varepsilon$. Then, since $s_n(T) = T$ for large enough n, we have

$$\begin{aligned} \|f - s_n(f)\|_2 &\le \|f - T\|_2 + \|s_n(T - f)\|_2 \\ &\le 2\|f - T\|_2 \le 2\sqrt{2}\,\|f - T\| < 2\sqrt{2}\,\varepsilon. \end{aligned}$$

(Compare this calculation with Lebesgue's Theorem, page 74.)

By way of comparison, let's give a simple class of functions whose Fourier partial sums provide good *uniform* approximations.

Theorem. *If $f'' \in C^{2\pi}$, then the Fourier series for f converges absolutely and uniformly to f.*

PROOF. First notice that integration by-parts leads to an estimate on the order of growth of the Fourier coefficients of f:

$$\pi a_k = \int_{-\pi}^{\pi} f(x)\, \cos kx\, dx = \int_{-\pi}^{\pi} f(x)\, d\left(\frac{\sin kx}{k}\right) = -\frac{1}{k}\int_{-\pi}^{\pi} f'(x)\, \sin kx\, dx$$

(because f is 2π-periodic). Thus, $|a_k| \le 2\|f'\|/k \to 0$ as $k \to \infty$. Now we integrate by-parts again:

$$-\pi k a_k = \int_{-\pi}^{\pi} f'(x)\, \sin kx\, dx = \int_{-\pi}^{\pi} f'(x)\, d\left(\frac{\cos kx}{k}\right) = \frac{1}{k}\int_{-\pi}^{\pi} f''(x)\, \cos kx\, dx$$

(because f' is 2π-periodic). Thus, $|a_k| \le 2\|f''\|/k^2 \to 0$ as $k \to \infty$. More importantly, this inequality (along with the Weierstrass M-test) implies that the Fourier series for f is both uniformly and absolutely convergent:

$$\left|\frac{a_0}{2} + \sum_{k=1}^{\infty}\big(a_k \cos kx + b_k \sin kx\big)\right| \le \left|\frac{a_0}{2}\right| + \sum_{k=1}^{\infty}\big(|a_k| + |b_k|\big) \le C\sum_{k=1}^{\infty}\frac{1}{k^2}.$$

But why should the series actually converge to f? Well, if we call the sum

$$g(x) = \frac{a_0}{2} + \sum_{k=1}^{\infty} \big(a_k \cos kx + b_k \sin kx \big),$$

then $g \in C^{2\pi}$ (why?) and g has the same Fourier coefficients as f (why?). Hence (by **7**), $g = f$. $\square$

Our next chore is to find a closed expression for $s_n(f)$. For this we'll need a couple of trig identities; the first two need no explanation.

$$\cos kt \cos kx + \sin kt \sin kx = \cos k(t - x)$$

$$2 \cos \alpha \sin \beta = \sin(\alpha + \beta) - \sin(\alpha - \beta)$$

$$\tfrac{1}{2} + \cos \theta + \cos 2\theta + \cdots + \cos n\theta = \frac{\sin\,(n + \frac{1}{2})\,\theta}{2 \sin \frac{1}{2}\theta}$$

Here's a short proof for the third:

$$\sin \tfrac{1}{2}\theta + \sum_{k=1}^{n} 2 \cos k\theta \, \sin \tfrac{1}{2}\theta = \sin \tfrac{1}{2}\theta + \sum_{k=1}^{n} \big[\sin\,(k + \tfrac{1}{2})\,\theta - \sin\,(k - \tfrac{1}{2})\,\theta \big] = \sin\,(n + \tfrac{1}{2})\,\theta.$$

The function

$$D_n(t) = \frac{\sin\,(n + \frac{1}{2})\,t}{2 \sin \frac{1}{2}\,t}$$

is called *Dirichlet's kernel*. It plays an important role in our next calculation.

Now we're ready to re-write our formula for $s_n(f)$.

$$\begin{aligned}
s_n(f)(x) &= \tfrac{1}{2} a_0 + \sum_{k=1}^{n} \big(a_k \cos kx + b_k \sin kx \big) \\
&= \frac{1}{\pi} \int_{-\pi}^{\pi} f(t) \left[\tfrac{1}{2} + \sum_{k=1}^{n} \cos kt \cos kx + \sin kt \sin kx \right] dt \\
&= \frac{1}{\pi} \int_{-\pi}^{\pi} f(t) \left[\tfrac{1}{2} + \sum_{k=1}^{n} \cos k(t - x) \right] dt \\
&= \frac{1}{\pi} \int_{-\pi}^{\pi} f(t) \cdot \frac{\sin\,(n + \frac{1}{2})\,(t - x)}{2 \sin \frac{1}{2}\,(t - x)}\, dt \\
&= \frac{1}{\pi} \int_{-\pi}^{\pi} f(t)\, D_n(t - x)\, dt = \frac{1}{\pi} \int_{-\pi}^{\pi} f(x + t)\, D_n(t)\, dt.
\end{aligned}$$

It now follows easily that $s_n(f)$ is linear in f (because integration against D_n is linear), that $s_n(f) \in \mathcal{T}_n$ (because $D_n \in \mathcal{T}_n$), and, in fact, that $s_n(\mathcal{T}_m) = \mathcal{T}_{\min(m,n)}$. In other words, s_n is indeed a linear projection onto $\mathcal{T}_n$.

While we know that $s_n(f)$ is a good approximation to f in the L_2 norm, a better understanding of its effectiveness as a uniform approximation will require a better understanding of the Dirichlet kernel D_n. Here are a few pertinent facts:

Lemma. (a) *D_n is even,*

(b) $\displaystyle \frac{1}{\pi}\int_{-\pi}^{\pi} D_n(t)\,dt = \frac{2}{\pi}\int_0^{\pi} D_n(t)\,dt = 1,$

(c) $|D_n(t)| \le n + \frac{1}{2}$ *and* $D_n(0) = n + \frac{1}{2}$,

(d) $\displaystyle \frac{|\sin(n+\frac{1}{2})t|}{t} \le |D_n(t)| \le \frac{\pi}{2t}$ *for* $0 < t < \pi$,

(e) *If* $\displaystyle \lambda_n = \frac{1}{\pi}\int_{-\pi}^{\pi} |D_n(t)|\,dt,$ *then* $\displaystyle \frac{4}{\pi^2}\log n \le \lambda_n \le 3 + \log n.$

Proof. (a), (b), and (c) are relatively clear from the fact that

$$D_n(t) = \tfrac{1}{2} + \cos t + \cos 2t + \cdots + \cos nt.$$

(Notice, too, that (b) follows from the fact that $s_n(1) = 1$.) For (d) we use a more delicate estimate: Since $2\theta/\pi \le \sin\theta \le \theta$ for $0 < \theta < \pi/2$, it follows that $2t/\pi \le 2\sin(t/2) \le t$ for $0 < t < \pi$. Hence,

$$\frac{\pi}{2t} \ge \frac{|\sin(n+\frac{1}{2})t|}{2\sin\frac{1}{2}t} \ge \frac{|\sin(n+\frac{1}{2})t|}{t}$$

for $0 < t < \pi$. Next, the upper estimate in (e) is easy:

$$\begin{aligned}
\frac{2}{\pi}\int_0^{\pi} |D_n(t)|\,dt &= \frac{2}{\pi}\int_0^{\pi} \frac{|\sin(n+\frac{1}{2})t|}{2\sin\frac{1}{2}t}\,dt \\
&\le \frac{2}{\pi}\int_0^{1/n} (n+\tfrac{1}{2})\,dt + \frac{2}{\pi}\int_{1/n}^{\pi} \frac{\pi}{2t}\,dt \\
&= \frac{2n+1}{\pi n} + \log\pi + \log n < 3 + \log n.
\end{aligned}$$

The lower estimate takes some work:

$$\frac{2}{\pi}\int_0^{\pi} |D_n(t)|\,dt = \frac{2}{\pi}\int_0^{\pi} \frac{|\sin(n+\frac{1}{2})t|}{2\sin\frac{1}{2}t}\,dt \ge \frac{2}{\pi}\int_0^{\pi} \frac{|\sin(n+\frac{1}{2})t|}{t}\,dt$$

$$\begin{aligned}
&= \frac{2}{\pi}\int_0^{(n+\frac{1}{2})\pi} \frac{|\sin x\,|}{x}\,dx \;\ge\; \frac{2}{\pi}\int_0^{n\pi} \frac{|\sin x\,|}{x}\,dx \\
&= \frac{2}{\pi}\sum_{k=1}^{n}\int_{(k-1)\pi}^{k\pi} \frac{|\sin x\,|}{x}\,dx \\
&\ge \frac{2}{\pi}\sum_{k=1}^{n}\frac{1}{k\pi}\int_{(k-1)\pi}^{k\pi} |\sin x\,|\,dx \;=\; \frac{4}{\pi^2}\sum_{k=1}^{n}\frac{1}{k} \\
&\ge \frac{4}{\pi^2}\log n,
\end{aligned}$$

because $\sum_{k=1}^{n}\frac{1}{k} \ge \log n$. □

The numbers $\lambda_n = \|D_n\|_1 = \frac{1}{\pi}\int_{-\pi}^{\pi}|D_n(t)|\,dt$ are called the *Lebesgue numbers* associated to this process (compare this to the terminology we used for interpolation). The point here is that λ_n gives the norm of the partial sum operator (projection) on $C^{2\pi}$ and (just as with interpolation) $\lambda_n \to \infty$ as $n \to \infty$. As a matter of no small curiosity, notice that, from Observation **10**, the norm of s_n as an operator on L_2 is 1.

Corollary. *If $f \in C^{2\pi}$, then*

$$|s_n(f)(x)| \le \frac{1}{\pi}\int_{-\pi}^{\pi}|f(x+t)|\,|D_n(t)|\,dt \le \lambda_n\|f\|. \tag{$*$}$$

In particular, $\|s_n(f)\| \le \lambda_n\|f\| \le (3+\log n)\|f\|$.

If we approximate the function $\operatorname{sgn} D_n$ by a continuous function f of norm one, then

$$s_n(f)(0) \approx \frac{1}{\pi}\int_{-\pi}^{\pi}|D_n(t)|\,dt = \lambda_n.$$

Thus, λ_n is the smallest constant that works in $(*)$. The fact that the partial sum operators are not uniformly bounded on $C^{2\pi}$, along with the Baire category theorem, tells us that there must be some $f \in C^{2\pi}$ for which $\|s_n(f)\|$ is unbounded. But, as we've seen, this has more to do with projections than it does with Fourier series:

Theorem. (Kharshiladze, Lozinski) *For each n, let L_n be a continuous, linear projection from $C^{2\pi}$ onto $\mathcal{T}_n$. Then, there is some $f \in C^{2\pi}$ for which $\|L_n(f) - f\|$ is unbounded.*

Although our last Corollary may not look very useful, it does give us some information about the effectiveness of $s_n(f)$ as a uniform approximation to f. Specifically, we have *Lebesgue's theorem*:

Theorem. *If* $f \in C^{2\pi}$*, and if we set* $E_n^T(f) = \min_{T \in \mathcal{T}_n} \|f - T\|$*, then*

$$E_n^T(f) \ \le\ \|f - s_n(f)\| \ \le\ (4 + \log n)\, E_n^T(f).$$

Proof. Let T^* be the best uniform approximation to f out of $\mathcal{T}_n$. Then, since $s_n(T^*) = T^*$, we get

$$\|f - s_n(f)\| \le \|f - T^*\| + \|s_n(T^* - f)\| \le (4 + \log n)\, \|f - T^*\|. \qquad \square$$

As an application of Lebesgue's theorem, let's speak briefly about "Chebyshev series," a notion that fits neatly in between our discussions of approximation by algebraic polynomials and by trig polynomials.

Theorem. *Suppose that* $f \in C[-1,1]$ *is twice continously differentiable. Then,* f *may be written as a uniformly and absolutely convergent Chebyshev series; that is,* $f(x) = \sum_{k=0}^\infty a_k T_k(x)$*, where* $\sum_{k=0}^\infty |a_k| < \infty$.

Proof. As usual, consider $\varphi(\theta) = f(\cos\theta) \in C^{2\pi}$. Since φ is *even* and twice differentiable, its Fourier series is an absolutely and uniformly convergent cosine series:

$$f(\cos\theta) = \varphi(\theta) = \sum_{k=0}^\infty a_k \cos k\theta = \sum_{k=0}^\infty a_k T_k(\cos\theta),$$

where $|a_k| \le 2\|\varphi''\|/k^2$. Thus, $f(x) = \sum_{k=0}^\infty a_k T_k(x)$. $\square$

If we write $S_n(f)(x) = \sum_{k=0}^n a_k T_k(x)$, we get an interesting consequence of this Theorem. First, notice that

$$S_n(f)(\cos\theta) = s_n(\varphi)(\theta).$$

Thus, from Lebesgue's theorem,

$$\begin{aligned} E_n(f) \ \le\ \|f - S_n(f)\|_{C[-1,1]} \ &=\ \|\varphi - s_n(\varphi)\|_{C^{2\pi}} \\ &\le\ (4 + \log n)\, E_n^T(\varphi) \ =\ (4 + \log n)\, E_n(f). \end{aligned}$$

For $n < 400$, this reads

$$E_n(f) \ \le\ \|f - S_n(f)\| \ \le\ 10\, E_n(f).$$

That is, for numerical purposes, the error incurred by using $\sum_{k=0}^{n} a_k T_k(x)$ to approximate f is *within one decimal place accuracy* of the best approximation! Notice, too, that $E_n(f)$ would be very easy to estimate in this case since

$$E_n(f) \;\le\; \|f - S_n(f)\| \;=\; \left\| \sum_{k>n} a_k T_k \right\| \;\le\; \sum_{k>n} |a_k| \;\le\; 2\,\|\varphi''\| \sum_{k>n} \frac{1}{k^2}.$$

Lebesgue's theorem should remind you of our "fancy" version of Bernstein's theorem; if we knew that $E_n^T(f)\log n \to 0$ as $n \to \infty$, then we'd know that $s_n(f)$ converged uniformly to f. Our goal, then, is to improve our estimates on $E_n^T(f)$, and the idea behind these improvements is to *replace* D_n *by a better kernel* (with regard to uniform approximation). Before we pursue anything quite so delicate as an estimate on $E_n^T(f)$, though, let's investigate a simple (and useful) replacement for D_n.

Since the sequence of partial sums (s_n) need not converge to f, we might try looking at their arithmetic means (or Cesàro sums):

$$\sigma_n = \frac{s_0 + s_1 + \cdots + s_{n-1}}{n}.$$

(These averages typically have better convergence properties than the partial sums themselves. Consider σ_n in the (scalar) case $s_n = (-1)^n$, for example.) Specifically, we set

$$\begin{aligned}\sigma_n(f)(x) &= \frac{1}{n}\Big[s_0(f)(x) + \cdots + s_{n-1}(f)(x)\Big] \\ &= \frac{1}{\pi}\int_{-\pi}^{\pi} f(x+t)\,\left[\frac{1}{n}\sum_{k=0}^{n-1} D_k(t)\right]\,dt = \frac{1}{\pi}\int_{-\pi}^{\pi} f(x+t)\,K_n(t)\,dt,\end{aligned}$$

where $K_n = (D_0 + D_1 + \cdots + D_{n-1})/n$ is called *Fejér's kernel*. The same techniques we used earlier can be applied to find a closed form for $\sigma_n(f)$ which, of course, reduces to simplifying $(D_0 + D_1 + \cdots + D_{n-1})/n$. As before, we begin with a trig identity:

$$\begin{aligned}2\sin\theta \sum_{k=0}^{n-1} \sin\,(2k+1)\theta &= \sum_{k=0}^{n-1} \big[\,\cos 2k\theta - \cos\,(2k+2)\theta\,\big] \\ &= 1 - \cos 2n\theta = 2\sin^2 n\theta.\end{aligned}$$

Thus,

$$K_n(t) \;=\; \frac{1}{n}\sum_{k=0}^{n-1} \frac{\sin\,(2k+1)\,t/2}{2\sin\,(t/2)} \;=\; \frac{\sin^2(nt/2)}{2n\sin^2(t/2)}.$$

Please note that K_n is even, *nonnegative*, and $\frac{1}{\pi}\int_{-\pi}^{\pi} K_n(t)\,dt = 1$. Thus, $\sigma_n(f)$ is a positive, linear map from $C^{2\pi}$ onto $\mathcal{T}_n$ (but it's not a projection—why?), satisfying $\|\sigma_n(f)\|_2 \le \|f\|_2$ (why?).

Now the arithmetic mean operator $\sigma_n(f)$ is still a good approximation f in L_2 norm. Indeed,

$$\|f - \sigma_n(f)\|_2 = \frac{1}{n}\left\|\sum_{k=0}^{n-1}(f - s_k(f))\right\|_2 \le \frac{1}{n}\sum_{k=0}^{n-1}\|f - s_k(f)\|_2 \to 0$$

as $n \to \infty$ (since $\|f - s_k(f)\|_2 \to 0$). But, more to the point, $\sigma_n(f)$ is actually a good *uniform* approximation to f, a fact that we'll call *Fejér's theorem*:

Theorem. *If $f \in C^{2\pi}$, then $\sigma_n(f)$ converges uniformly to f as $n \to \infty$.*

Note that, since $\sigma_n(f) \in \mathcal{T}_n$, Fejér's theorem implies Weierstrass's second theorem. Curiously, Fejér was only 19 years old when he proved this result (about 1900) while Weierstrass was 75 at the time he proved his approximation theorems.

We'll give two proofs of Fejér's theorem; one with details, one without. But both follow from quite general considerations. First:

Theorem. *Suppose that $k_n \in C^{2\pi}$ satisfies*

(a) $k_n \ge 0$,

(b) $\displaystyle\frac{1}{\pi}\int_{-\pi}^{\pi} k_n(t)\,dt = 1$, *and*

(c) $\displaystyle\int_{\delta \le |t| \le \pi} k_n(t)\,dt \to 0$ *for every* $\delta > 0$.

Then, $\displaystyle\frac{1}{\pi}\int_{-\pi}^{\pi} f(x+t)\,k_n(t)\,dt \rightrightarrows f(x)$ *for each* $f \in C^{2\pi}$.

PROOF. Let $\varepsilon > 0$. Since f is uniformly continuous, we may choose $\delta > 0$ so that $|f(x) - f(x+t)| < \varepsilon$, for any x, whenever $|t| < \delta$. Next, we use the fact that k_n is nonnegative and integrates to 1 to write

$$\begin{aligned}\left|f(x) - \frac{1}{\pi}\int_{-\pi}^{\pi} f(x+t)\,k_n(t)\,dt\right| &= \frac{1}{\pi}\left|\int_{-\pi}^{\pi}\big[f(x) - f(x+t)\big]\,k_n(t)\,dt\right| \\ &\le \frac{1}{\pi}\int_{-\pi}^{\pi}\big|f(x) - f(x+t)\big|\,k_n(t)\,dt\end{aligned}$$

$$\leq \frac{\varepsilon}{\pi}\int_{|t|<\delta} k_n(t)\,dt \;+\; \frac{2\|f\|}{\pi}\int_{\delta\leq|t|\leq\pi} k_n(t)\,dt$$
$$< \varepsilon + \varepsilon = 2\varepsilon,$$

for n sufficiently large. $\square$

To see that Fejér's kernel satisfies the conditions of the Theorem is easy: In particular, (c) follows from the fact that $K_n(t) \rightrightarrows 0$ on the set $\delta \leq |t| \leq \pi$. Indeed, since $\sin(t/2)$ increases on $\delta \leq t \leq \pi$ we have

$$K_n(t) \;=\; \frac{\sin^2(nt/2)}{2n\sin^2(t/2)} \;\leq\; \frac{1}{2n\sin^2(\delta/2)} \;\to 0.$$

Our second proof, or sketch, really, is based on a variant of the Bohman-Korovkin theorem for $C^{2\pi}$, due to Korovkin. In this setting, the three "test cases" are

$$f_0(x) = 1, \qquad f_1(x) = \cos x, \qquad \text{and} \qquad f_2(x) = \sin x.$$

Theorem. *Let (L_n) be a sequence of positive, linear maps on $C^{2\pi}$. If $L_n(f) \rightrightarrows f$ for each of the three functions $f_0(x) = 1$, $f_1(x) = \cos x$, and $f_2(x) = \sin x$, then $L_n(f) \rightrightarrows f$ for every $f \in C^{2\pi}$.*

We won't prove this theorem; rather, we'll check that $\sigma_n(f) \rightrightarrows f$ in each of the three test cases. Since s_n is a projection, this is painfully simple!

$$\sigma_n(f_0) = \tfrac{1}{n}\,(f_0 + f_0 + \cdots + f_0) = f_0,$$
$$\sigma_n(f_1) = \tfrac{1}{n}\,(0 + f_1 + \cdots + f_1) = \tfrac{n-1}{n}\cdot f_1 \rightrightarrows f_1,$$
$$\sigma_n(f_2) = \tfrac{1}{n}\,(0 + f_2 + \cdots + f_2) = \tfrac{n-1}{n}\cdot f_2 \rightrightarrows f_2.$$

Kernel operators abound in analysis; for example, Landau's proof of the Weierstrass theorem uses the kernel $L_n(x) = c_n(1-x^2)^n$. And, in the next section, we'll encounter *Jackson's kernel* $J_n(t) = c_n \sin^4 nt/n^3 \sin^4 t$, which is essentially the square of Fejér's kernel. While we will have no need for a general theory of such operators, please note that the key to their utility is the fact that they're nonnegative!

Lastly, a word or two about Fourier series involving complex coefficients. Most modern textbooks consider the case of a 2π-periodic, integrable function $f : \mathbb{R} \to \mathbb{C}$ and define the Fourier series of f by

$$\sum_{k=-\infty}^{\infty} c_k e^{ikt},$$

where now we have only one formula for the c_k's:

$$c_k = \frac{1}{2\pi} \int_{-\pi}^{\pi} f(t)\, e^{-ikt}\, dt,$$

but, of course, the c_k's may well be complex. This somewhat simpler approach has other advantages; for one, the exponentials e^{ikt} are now an orthonormal set (relative to the normalizing constant $1/2\pi$). And, if we remain consistent with this choice and define the L_2 norm by

$$\|f\|_2 = \left(\frac{1}{2\pi} \int_{-\pi}^{\pi} |f(t)|^2\, dt \right)^{1/2},$$

then we have the simpler estimate $\|f\|_2 \le \|f\|$ for $f \in C^{2\pi}$.

The Dirichlet and Fejer kernels are essentially the same in this case, too, except that we would now write $s_n(f)(x) = \sum_{k=-n}^{n} c_k e^{ikx}$. Given this, the Dirichlet and Fejér kernels can be written

$$\begin{aligned} D_n(x) = \sum_{k=-n}^{n} e^{ikx} &= 1 + \sum_{k=1}^{n} (e^{ikx} + e^{-ikx}) \\ &= 1 + 2 \sum_{k=1}^{n} \cos kx \\ &= \frac{\sin\left(n + \frac{1}{2}\right) x}{\sin \frac{1}{2}\, x} \end{aligned}$$

and

$$\begin{aligned} K_n(x) &= \frac{1}{n} \sum_{m=0}^{n-1} \sum_{k=-m}^{m} e^{ikx} = \sum_{k=-n}^{n} \left(1 - \frac{|k|}{n} \right) e^{ikx} \\ &= \frac{1}{n} \sum_{m=0}^{n-1} \frac{\sin\left(m + \frac{1}{2}\right) x}{\sin \frac{1}{2}\, x} \\ &= \frac{\sin^2(nt/2)}{n \sin^2(t/2)}. \end{aligned}$$

In other words, each is twice its real coefficient counterpart. Since the choice of normalizing constant ($1/\pi$ versus $1/2\pi$, and sometimes even $1/\sqrt{\pi}$ or $1/\sqrt{2\pi}$) has a (small) effect on these formulas, you may find some variation in other textbooks.

$*$ **57.** Define $f(x) = (\pi - x)^2$ for $0 \le x \le 2\pi$, and extend f to a 2π-periodic continuous function on $\mathbb{R}$ in the obvious way. Check that the Fourier series for f is $\pi^2/3 + 4\sum_{n=1}^{\infty} \cos nx/n^2$. Since this series is uniformly convergent, it actually converges to f. In particular, note that setting $x = 0$ yields the familiar formula $\sum_{n=1}^{\infty} 1/n^2 = \pi^2/6$.

58. (a) Given $n \ge 1$ and $\varepsilon > 0$, show that there is a continuous function $f \in C^{2\pi}$ satisfying $\|f\| = 1$ and $\frac{1}{\pi}\int_{-\pi}^{\pi} |f(t) - \operatorname{sgn} D_n(t)|\, dt < \varepsilon/(n+1)$.

(b) Show that $s_n(f)(0) \ge \lambda_n - \varepsilon$ and, hence, that $\|s_n(f)\| \ge \lambda_n - \varepsilon$.

59. (a) If f, $k \in C^{2\pi}$, prove that $g(x) = \int_{-\pi}^{\pi} f(x+t)\, k(t)\, dt$ is also in $C^{2\pi}$.

(b) If we only assume that f is 2π-periodic and Riemann integrable on $[-\pi, \pi\,]$ (but still $k \in C^{2\pi}$), is g still continuous?

(c) If we simply assume that f and k are 2π-periodic and Riemann integrable on $[-\pi, \pi\,]$, is g still continuous?

60. Suppose that $k_n \in C^{2\pi}$ satisfies

$$k_n \ge 0, \quad \frac{1}{\pi}\int_{-\pi}^{\pi} k_n(t)\, dt = 1, \quad \text{and} \quad \int_{\delta \le |t| \le \pi} k_n(t)\, dt \to 0 \ (n \to \infty)$$

for every $\delta > 0$. If f is Riemann integrable, show that $\frac{1}{\pi}\int_{-\pi}^{\pi} f(x+t)\, k_n(t)\, dt \to f(x)$ pointwise, as $n \to \infty$, at each point of continuity of f. In particular, $\sigma_n(f)(x) \to f(x)$ at each point of continuity of f.

$*$ **61.** Given f, $g \in C^{2\pi}$, we define the *convolution* of f and g, written $f * g$, by

$$(f * g)(x) = \frac{1}{\pi}\int_{-\pi}^{\pi} f(t)\, g(x - t)\, dt.$$

(Compare this integral with that used in problem **59**.)

(a) Show that $f * g = g * f$ and that $f * g \in C^{2\pi}$.

(b) If one of f or g is a trig polynomial, show that $f * g$ is again a trig polynomial (of the same degree).

(c) If one of f or g is continuously differentiable, show that $f * g$ is likewise continuously differentiable and find an integral formula for $(f * g)'(x)$.

We continue our investigations of the "middle ground" between algebraic and trigonometric approximation by presenting several results due to the great American mathematician Dunham Jackson (from roughly 1911–1912). The first of these results will give us the best possible estimate of $E_n(f)$ in terms of ω_f and n.

Jackson's Theorem 1. *If $f \in C^{2\pi}$, then $E_n^T(f) \le 6\,\omega_f([-\pi,\pi\,]; \frac{1}{n})$.*

Theorem 1 should be viewed as an improvement over Bernstein's Theorem, which stated that $E_n(f) \le \frac{3}{2}\omega_f(\frac{1}{\sqrt{n}})$ for $f \in C[-1,1]$. As we'll see, the proof of Theorem 1 not only mimics the proof of Bernstein's result, but also uses some of the ideas we talked about in the last section. In particular, the proof we'll give involves integration against an "improved" Dirichlet kernel.

Before we dive into the proof, let's list several immediate and important Corollaries:

Corollary. *Weierstrass's second theorem* (since $\omega_f(\frac{1}{n}) \to 0$ for any $f \in C^{2\pi}$).

Corollary. *The Dini-Lipschitz theorem: If $\omega_f(\frac{1}{n})\log n \to 0$ as $n \to \infty$, then the Fourier series for f converges uniformly to f.*

PROOF. From Lebesgue's theorem,

$$\|f - s_n(f)\| \;\le\; (4 + \log n)\, E_n^T(f) \;\le\; 6\,(4 + \log n)\, \omega_f\left(\frac{1}{n}\right) \to 0. \qquad \square$$

Jackson's Theorem 2. *If $f \in C[-1,1]$, then $E_n(f) \le 6\,\omega_f([-1,1\,]; \frac{1}{n})$.*

PROOF. Let $\varphi(\theta) = f(\cos\theta)$. Then, as we've seen,

$$E_n(f) \;=\; E_n^T(\varphi) \;\le\; 6\,\omega_\varphi\left([-\pi,\pi\,]; \frac{1}{n}\right) \;\le 6\,\omega_f\left([-1,1\,]; \frac{1}{n}\right),$$

where the last inequality follows from the fact that

$$|\varphi(\alpha) - \varphi(\beta)| \;=\; |f(\cos\alpha) - f(\cos\beta)| \;\le\; \omega_f(|\cos\alpha - \cos\beta\,|) \;\le\; \omega_f(|\alpha - \beta\,|). \qquad \square$$

Corollary. *If* $f \in \mathrm{lip}_K\alpha$ *on* $[-1,1]$, *then* $E_n(f) \le 6Kn^{-\alpha}$. (Recall that Bernstein's theorem gives only $n^{-\alpha/2}$.)

Corollary. *If* $f \in C[-1,1]$ *has a bounded derivative, then* $E_n(f) \le \frac{6}{n}\,\|f'\|$.

Corollary. *If* $f \in C[-1,1]$ *has a continuous derivative, then* $E_n(f) \le \frac{6}{n}\,E_{n-1}(f')$.

PROOF. Let $p^* \in \mathcal{P}_{n-1}$ be the best uniform approximation to f' and consider $p(x) = \int_{-1}^x p^*(t)\,dt \in \mathcal{P}_n$. From the previous Corollary,

$$\begin{aligned} E_n(f) &= E_n(f-p) \qquad \text{(Why?)} \\ &\le \frac{6}{n}\,\|f'-p^*\| \;=\; \frac{6}{n}\,E_{n-1}(f'). \qquad \square \end{aligned}$$

Iterating this last inequality will give the following result:

Corollary. *If* $f \in C[-1,1]$ *is* k*-times continuously differentiable, then*

$$E_n(f) \;\le\; \frac{6^{k+1}}{n(n-1)\cdots(n-k+1)}\,\omega_k\left(\frac{1}{n-k}\right),$$

where ω_k *is the modulus of continuity of* $f^{(k)}$.

Well, enough corollaries. It's time we proved Jackson's Theorem 1. Now Jackson's approach was to show that

$$\frac{1}{\pi}\int_{-\pi}^{\pi} f(x+t)\,\cdot c_n\left(\frac{\sin nt}{\sin t}\right)^4 dt \;\rightrightarrows\; f(x),$$

where $J_n(t) = c_n(\sin nt/\sin t)^4$ is the "improved" kernel we alluded to earlier (it's essentially the square of Fejér's kernel). The approach we'll take, due to Korovkin, proves the existence of a suitable kernel without giving a tidy formula for it. On the other hand, it's relatively easy to outline the idea. The key here is that $J_n(t)$ should be an even, nonnegative, trig polynomial of degree n with $\frac{1}{\pi}\int_{-\pi}^{\pi} J_n(t)\,dt = 1$. In other words,

$$J_n(t) \;=\; \frac{1}{2} + \sum_{k=1}^{n} \rho_{k,n}\cos kt$$

(why is the first term $1/2$?), where $\rho_{1,n},\ldots,\rho_{n,n}$ must be chosen so that $J_n(t) \ge 0$. Assuming we can find such $\rho_{k,n}$'s, here's what we get:

Lemma. *If* $f \in C^{2\pi}$, *then*

$$\left| f(x) - \frac{1}{\pi} \int_{-\pi}^{\pi} f(x+t)\, J_n(t)\, dt \right| \;\le\; \omega_f\left(\frac{1}{n}\right) \cdot \left[1 + n\pi \sqrt{\frac{1-\rho_{1,n}}{2}}\, \right].$$

PROOF. We already know how the first several lines of the proof will go:

$$\begin{aligned} \left| f(x) - \frac{1}{\pi} \int_{-\pi}^{\pi} f(x+t)\, J_n(t)\, dt \right| &= \frac{1}{\pi} \left| \int_{-\pi}^{\pi} \big[f(x) - f(x+t) \big]\, J_n(t)\, dt \right| \\ &\le \frac{1}{\pi} \int_{-\pi}^{\pi} |f(x) - f(x+t)|\, J_n(t)\, dt \\ &\le \frac{1}{\pi} \int_{-\pi}^{\pi} \omega_f(\,|t|\,)\, J_n(t)\, dt. \end{aligned}$$

Next we borrow a trick from Bernstein. We replace $\omega_f(\,|t|\,)$ by

$$\omega_f(\,|t|\,) \;=\; \omega_f\left(n|t| \cdot \frac{1}{n} \right) \;\le\; (1 + n|t|\,)\, \omega_f\left(\frac{1}{n}\right),$$

and so the last integral on the right-hand side, above, is dominated by

$$\omega_f\left(\frac{1}{n}\right) \cdot \frac{1}{\pi} \int_{-\pi}^{\pi} (1 + n|t|\,)\, J_n(t)\, dt \;=\; \omega_f\left(\frac{1}{n}\right) \cdot \left[1 + \frac{n}{\pi} \int_{-\pi}^{\pi} |t|\, J_n(t)\, dt \right].$$

All that remains is to estimate $\int_{-\pi}^{\pi} |t|\, J_n(t)\, dt$, and for this we'll appeal to the Cauchy-Schwarz inequality (again, compare this to the proof of Bernstein's theorem).

$$\begin{aligned} \frac{1}{\pi} \int_{-\pi}^{\pi} |t|\, J_n(t)\, dt &= \frac{1}{\pi} \int_{-\pi}^{\pi} |t|\, J_n(t)^{1/2}\, J_n(t)^{1/2}\, dt \\ &\le \left(\frac{1}{\pi} \int_{-\pi}^{\pi} |t|^2\, J_n(t)\, dt \right)^{1/2} \left(\frac{1}{\pi} \int_{-\pi}^{\pi} J_n(t)\, dt \right)^{1/2} \\ &= \left(\frac{1}{\pi} \int_{-\pi}^{\pi} |t|^2\, J_n(t)\, dt \right)^{1/2}. \end{aligned}$$

But,

$$|t|^2 \;\le\; \left[\pi \sin\left(\frac{t}{2}\right) \right]^2 \;=\; \frac{\pi^2}{2}\,(1 - \cos t\,).$$

So,

$$\frac{1}{\pi} \int_{-\pi}^{\pi} |t|\, J_n(t)\, dt \;\le\; \left(\frac{\pi^2}{2} \cdot \frac{1}{\pi} \int_{-\pi}^{\pi} (1 - \cos t)\, J_n(t)\, dt \right)^{1/2} \;=\; \pi \sqrt{\frac{1-\rho_{1,n}}{2}}. \qquad \square$$

Now we still have to prove that we can actually find a suitable choice of scalars $\rho_{1,n}, \ldots, \rho_{n,n}$. We already know that we need to choose the $\rho_{k,n}$'s so that $J_n(t)$ will be nonnegative, but now it's clear that we also want $\rho_{1,n}$ to be very close to 1. To get us started, let's first see why it's easy to generate nonnegative cosine polynomials. Given *real numbers* $c_0, \ldots, c_n$, note that

$$\begin{aligned} 0 \;\le\; \left| \sum_{k=0}^{n} c_k e^{ikx} \right|^2 &= \left(\sum_{k=0}^{n} c_k e^{ikx} \right) \left(\sum_{j=0}^{n} c_j e^{-ijx} \right) \;=\; \sum_{k,j} c_k c_j e^{i(k-j)x} \\ &= \sum_{k=0}^{n} c_k^2 \;+\; \sum_{k>j} c_k c_j \left(e^{i(k-j)x} + e^{i(j-k)x} \right) \\ &= \sum_{k=0}^{n} c_k^2 \;+\; 2 \sum_{k>j} c_k c_j \cos(k-j)x \\ &= \sum_{k=0}^{n} c_k^2 \;+\; 2 \sum_{k=0}^{n-1} c_k c_{k+1} \cos x + \cdots + 2 c_0 c_n \cos nx. \qquad (*) \end{aligned}$$

In particular, we need to find $c_0, \ldots, c_n$ with

$$\sum_{k=0}^{n} c_k^2 \;=\; \frac{1}{2} \qquad \text{and} \qquad \rho_{1,n} = 2 \sum_{k=0}^{n-1} c_k c_{k+1} \;\approx\; 1.$$

What we'll do is find c_k's with $\sum_{k=0}^{n-1} c_k c_{k+1} \approx \sum_{k=0}^{n} c_k^2$, and then normalize. But, in fact, we won't actually *find* anything—we'll simply write down a choice of c_k's that happens to work! Consider:

$$\begin{aligned} \sum_{k=0}^{n} \sin\left(\frac{k+1}{n+2}\,\pi \right) \sin\left(\frac{k+2}{n+2}\,\pi \right) &= \sum_{k=0}^{n} \sin\left(\frac{k+1}{n+2}\,\pi \right) \sin\left(\frac{k}{n+2}\,\pi \right) \\ &= \frac{1}{2} \sum_{k=0}^{n} \left[\sin\left(\frac{k}{n+2}\,\pi \right) + \sin\left(\frac{k+2}{n+2}\,\pi \right) \right] \sin\left(\frac{k+1}{n+2}\,\pi \right). \end{aligned}$$

By changing the index of summation, it's easy to see that first two sums are equal and, hence, each is equal to the average of the two. Next we re-write this last sum, using the trig identity $\frac{1}{2}\left(\sin A + \sin B\right) = \cos\left(\frac{A-B}{2}\right) \sin\left(\frac{A+B}{2}\right)$, to get

$$\sum_{k=0}^{n} \sin\left(\frac{k+1}{n+2}\,\pi \right) \sin\left(\frac{k+2}{n+2}\,\pi \right) \;=\; \cos\left(\frac{\pi}{n+2} \right) \sum_{k=0}^{n} \sin^2\left(\frac{k+1}{n+2}\,\pi \right).$$

Since $\cos\left(\frac{\pi}{n+2}\right) \approx 1$ for large n, we've done it! If we define $c_k = c \cdot \sin\left(\frac{k+1}{n+2}\,\pi\right)$, where c is chosen so that $\sum_{k=0}^n c_k^2 = 1/2$, and if we define $J_n(x)$ using $(*)$, then $J_n(x) \geq 0$ and $\rho_{1,n} = \cos\left(\frac{\pi}{n+2}\right)$ (why?). The estimate needed in our Lemma becomes

$$\sqrt{\frac{1-\rho_{1,n}}{2}} \;=\; \sqrt{\frac{1-\cos\left(\frac{\pi}{n+2}\right)}{2}} \;=\; \sin\left(\frac{\pi}{2n+4}\right) \;\leq\; \frac{\pi}{2n},$$

and so we have

$$E_n^T(f) \;\leq\; \left(1+\frac{\pi^2}{2}\right)\omega_f\left(\frac{1}{n}\right) \;<\; 6\,\omega_f\left(\frac{1}{n}\right). \quad \square$$

Jackson's theorems are what we might call *direct theorems*. If we know something about f, then we can say something about $E_n(f)$. There is also the notion of an *inverse theorem*, meaning that if we know something about $E_n(f)$, we should be able to say something about f. In other words, we would expect an inverse theorem to be, more or less, the converse of some direct theorem. Now inverse theorems are typically much harder to prove than direct theorems, but in order to have some idea of what such theorems might tell us (and to see some of the techniques used in their proofs), we present one of the easier inverse theorems, due to Bernstein. This result gives the converse to one of our corollaries to Jackson's theorem (see the top of page 105).

Theorem. *If $f \in C^{2\pi}$ satisfies $E_n^T(f) \leq A\,n^{-\alpha}$, for some constants A and $0 < \alpha < 1$, then $f \in \mathrm{lip}_K\alpha$ for some constant K.*

PROOF. For each n, choose $U_n \in \mathcal{T}_n$ so that $\|f - U_n\| \leq A\,n^{-\alpha}$. Then, in particular, (U_n) converges uniformly to f. Now if we set $V_0 = U_1$ and $V_n = U_{2^n} - U_{2^{n-1}}$ for $n \geq 1$, then $V_n \in \mathcal{T}_{2^n}$ and $f = \sum_{n=0}^\infty V_n$. Indeed,

$$\|V_n\| \;\leq\; \|U_{2^n} - f\| + \|U_{2^{n-1}} - f\| \;\leq\; A\,(2^n)^{-\alpha} + A\,(2^{n-1})^{-\alpha} \;=\; B\cdot 2^{-n\alpha},$$

which is summable; thus, the (telescoping) series $\sum_{n=0}^\infty V_n$ converges uniformly to f. (Why?)

Next we estimate $|f(x) - f(y)|$ using finitely many of the V_n's, the precise number to be specified later. Using the mean value theorem and Bernstein's inequality we get

$$|f(x) - f(y)| \;\leq\; \sum_{n=0}^\infty |V_n(x) - V_n(y)|$$

$$
\begin{aligned}
&\le \sum_{n=0}^{m-1} |V_n(x) - V_n(y)| \ + \ 2\sum_{n=m}^{\infty} \|V_n\| \\
&= \sum_{n=0}^{m-1} |V_n'(\xi_n)|\,|x-y| \ + \ 2\sum_{n=m}^{\infty} \|V_n\| \\
&\le |x-y| \sum_{n=0}^{m-1} 2^n \|V_n\| \ + \ 2\sum_{n=m}^{\infty} \|V_n\| \\
&\le |x-y| \sum_{n=0}^{m-1} B\,2^{n(1-\alpha)} \ + \ 2\sum_{n=m}^{\infty} B\,2^{-n\alpha} \\
&\le C\left[\,|x-y|\cdot 2^{m(1-\alpha)} \ + \ 2^{-m\alpha}\right], \qquad (**)
\end{aligned}
$$

where we've used, in the fourth line, the fact that $V_n \in \mathcal{T}_{2^n}$ and, in the last line, standard estimates for geometric series. Now we want the right-hand side to be dominated by a constant times $|x-y|^\alpha$. In other words, if we set $|x-y| = \delta$, then we want

$$\delta \cdot 2^{m(1-\alpha)} \ + \ 2^{-m\alpha} \ \le \ D \cdot \delta^\alpha$$

or, equivalently,

$$(2^m\delta)^{(1-\alpha)} \ + \ (2^m\delta)^{-\alpha} \ \le \ D.$$

Thus, we should choose m so that $2^m\delta$ is both bounded above and bounded away from zero. For example, if $0 < \delta < 1$, we could choose m so that $1 \le 2^m\delta < 2$. □

In order to better explain the phrase "more or less the converse of some direct theorem," let's see how the previous result falls apart when $\alpha = 1$. Although we might hope that $E_n^T(f) \le A/n$ would imply that $f \in \mathrm{lip}_K 1$, it happens not to be true. The best result in this regard is due to Zygmund, who gave necessary and sufficient conditions on f so that $E_n^T(f) \le A/n$ (and these conditions do not characterize $\mathrm{lip}_K 1$ functions). Instead of pursuing Zygmund's result, we'll settle for simple "surgery" on our previous result, keeping an eye out for what goes wrong. This result is again due to Bernstein.

Theorem. *If $f \in C^{2\pi}$ satisfies $E_n^T(f) \le A/n$, then $\omega_f(\delta) \le K\delta|\log\delta\,|$ for some constant K and all δ sufficiently small.*

Proof. If we repeat the previous proof, setting $\alpha = 1$, only a few lines change. In particular, the conclusion of that long string of inequalites $(**)$ would now read

$$|f(x) - f(y)| \ \le \ C\left[\,|x-y|\cdot m \ + \ 2^{-m}\,\right] \ = \ C\left[\,m\delta \ + \ 2^{-m}\,\right].$$

Clearly, the right-hand side cannot be dominated by a constant times δ, as we might have hoped, for this would force m to be bounded (independent of δ), which in turn bounds δ away from zero. But, if we again think of $2^m\delta$ as the "variable" in this inequality, then the term $m\delta$ suggests that the correct order of magnitude of the right-hand side is $\delta|\log\delta\,|$. Thus, we would try to find a constant D so that

$$m\delta \;+\; 2^{-m} \;\le\; D\cdot\delta|\log\delta\,|$$

or

$$m(2^m\delta) \;+\; 1 \;\le\; D\cdot(2^m\delta)|\log\delta\,|.$$

Now if we take $0<\delta<1/2$, then $\log 2 < -\log\delta = |\log\delta\,|$. Hence, if we again choose $m\ge 1$ so that $1\le 2^m\delta<2$, we'll get

$$m\log 2+\log\delta<\log 2 \quad\Longrightarrow\quad m<\frac{\log 2-\log\delta}{\log 2}<\frac{2}{\log 2}\,|\log\delta\,|$$

and, finally,

$$m(2^m\delta)+1 \;\le\; 2m+1 \;\le\; 3m \;\le\; \frac{6}{\log 2}\,|\log\delta\,| \;\le\; \frac{6}{\log 2}\,(2^m\delta)\,|\log\delta\,|. \qquad \square$$

Given a positive (except possibly at finitely many points), Riemann integrable *weight function* $w(x)$ on $[\,a,b\,]$, the expression

$$\langle f,g\rangle \;=\; \int_a^b f(x)\,g(x)\,w(x)\,dx$$

defines an inner product on $C[\,a,b\,]$ and

$$\|f\|_2 \;=\; \left(\int_a^b f(x)^2\,w(x)\,dx\right)^{1/2} \;=\; \sqrt{\langle f,f\rangle}$$

defines a *strictly convex* norm on $C[\,a,b\,]$. Thus, given a finite dimensional subspace E of $C[\,a,b\,]$ and an element $f\in C[\,a,b\,]$, there is a unique $g\in E$ such that

$$\|f-g\|_2 \;=\; \min_{h\in E}\|f-h\|_2.$$

We say that g is the *least-squares approximation* to f out of E (relative to w).

Now if we apply the Gram-Schmidt procedure to the sequence $1,x,x^2,\ldots$, we will arrive at a sequence (Q_n) of *orthogonal polynomials* relative to the above inner product. In this special case, however, the Gram-Schmidt procedure simplifies substantially:

Theorem. *The following procedure defines a sequence (Q_n) of orthogonal polynomials (relative to w). Set:*

$$Q_0(x)=1,\quad Q_1(x)=x-a_0,\quad \text{and}\quad Q_{n+1}(x)=(x-a_n)Q_n(x)-b_nQ_{n-1}(x),$$

for $n\ge 1$, where

$$a_n=\langle\,x\,Q_n,Q_n\,\rangle/\langle\,Q_n,Q_n\,\rangle \quad\text{and}\quad b_n=\langle\,x\,Q_n,Q_{n-1}\,\rangle/\langle\,Q_{n-1},Q_{n-1}\,\rangle$$

(and where $x\,Q_n$ is shorthand for the polynomial $x\,Q_n(x)$).

PROOF. It's easy to see from these formulas that Q_n is a monic polynomial of degree exactly n. In particular, the Q_n's are linearly independent (and nonzero).

Now we checked in class that Q_0, Q_1, and Q_2 are mutually orthogonal, so let's use induction and check that Q_{n+1} is orthogonal to each Q_k, $k \le n$. First,

$$\langle\, Q_{n+1}, Q_n \,\rangle = \langle\, x\, Q_n, Q_n \,\rangle - a_n\langle\, Q_n, Q_n \,\rangle - b_n\langle\, Q_{n-1}, Q_n \,\rangle = 0$$

and

$$\langle\, Q_{n+1}, Q_{n-1} \,\rangle = \langle\, x\, Q_n, Q_{n-1} \,\rangle - a_n\langle\, Q_n, Q_{n-1} \,\rangle - b_n\langle\, Q_{n-1}, Q_{n-1} \,\rangle = 0,$$

since $\langle\, Q_{n-1}, Q_n \,\rangle = 0$. Next, we take $k < n-1$ and use the recurrence formula twice:

$$\begin{aligned}\langle\, Q_{n+1}, Q_k \,\rangle &= \langle\, x\, Q_n, Q_k \,\rangle - a_n\langle\, Q_n, Q_k \,\rangle - b_n\langle\, Q_{n-1}, Q_k \,\rangle \\ &= \langle\, x\, Q_n, Q_k \,\rangle = \langle\, Q_n, x\, Q_k \,\rangle \qquad \text{(Why?)} \\ &= \langle\, Q_n, Q_{k+1} + a_k Q_k + b_k Q_{k-1} \,\rangle = 0,\end{aligned}$$

since $k+1 < n$. $\square$

Observations

1. Using the same trick as above, we have

$$b_n = \langle\, x\, Q_n, Q_{n-1} \,\rangle \big/ \langle\, Q_{n-1}, Q_{n-1} \,\rangle = \langle\, Q_n, Q_n \,\rangle \big/ \langle\, Q_{n-1}, Q_{n-1} \,\rangle > 0.$$

2. Each $p \in \mathcal{P}_n$ can be uniquely written $p = \sum_{i=0}^n \alpha_i Q_i$, where $\alpha_i = \langle\, p, Q_i \,\rangle \big/ \langle\, Q_i, Q_i \,\rangle$.

3. If Q is any monic polynomial of degree exactly n, then $Q = Q_n + \sum_{i=0}^{n-1} \alpha_i Q_i$ (why?) and hence

$$\|Q\|_2^2 \;=\; \|Q_n\|_2^2 \;+\; \sum_{i=0}^{n-1} \alpha_i^2 \|Q_i\|_2^2 \;>\; \|Q_n\|_2^2,$$

unless $Q = Q_n$. That is, Q_n has the least $\|\cdot\|_2$ norm of all monic polynomials of degree n.

4. The Q_n's are *unique* in the following sense: If (P_n) is another sequence of orthogonal polynomials such that P_n has degree exactly n, then $P_n = \alpha_n Q_n$ for some $\alpha_n \neq 0$. (Why?) Consequently, there's no harm in referring to the Q_n's as *the* sequence of orthogonal polynomials relative to w.

5. For $n \ge 1$ note that $\int_a^b Q_n(t)\, w(t)\, dt = \langle Q_0, Q_n \rangle = 0$.

Examples

1. On $[-1,1]$, the Chebyshev polynomials of the first kind (T_n) are orthogonal relative to the weight $w(x) = 1/\sqrt{1-x^2}$.

$$\int_{-1}^{1} T_m(x)\, T_n(x)\, \frac{dx}{\sqrt{1-x^2}} = \int_0^\pi \cos m\theta\, \cos n\theta\, d\theta = \begin{cases} 0, & m \neq n \\ \pi, & m = n = 0 \\ \pi/2, & m = n \neq 0. \end{cases}$$

Since T_n has degree exactly n, this must be the right choice. Notice, too, that $\frac{1}{\sqrt{2}} T_0, T_1, T_2, \ldots$ are *orthonormal* relative to the weight $2/\pi\sqrt{1-x^2}$.

In terms of the inductive procedure given above, we must have $Q_0 = T_0 = 1$ and $Q_n = 2^{-n+1} T_n$ for $n \geq 1$. (Why?) From this it follows that $a_n = 0$, $b_1 = 1/2$, and $b_n = 1/4$ for $n \geq 2$. (Why?) That is, the recurrence formula given in our first Theorem reduces to the familar relationship $T_{n+1}(x) = 2x\, T_n(x) - T_{n-1}(x)$. Curiously, $Q_n = 2^{-n+1} T_n$ minimizes both

$$\max_{-1 \leq x \leq 1} |p(x)| \qquad \text{and} \qquad \left(\int_{-1}^{1} p(x)^2\, \frac{dx}{\sqrt{1-x^2}} \right)^{1/2}$$

over all monic polynomials of degree exactly n.

The Chebyshev polynomials also satisfy $(1-x^2)\, T_n''(x) - x\, T_n'(x) + n^2\, T_n(x) = 0$. Since this is a polynomial identity, it suffices to check it for all $x = \cos\theta$. In this case,

$$T_n'(x) = \frac{n \sin n\theta}{\sin\theta}$$

and

$$T_n''(x) = \frac{n^2 \cos n\theta \sin\theta - n \sin n\theta \cos\theta}{\sin^2\theta\, (-\sin\theta)}.$$

Hence,

$$\begin{aligned} &(1-x^2)\, T_n''(x) - x\, T_n'(x) + n^2\, T_n(x) \\ &\qquad = -n^2 \cos n\theta + n \sin n\theta \cot\theta - n \sin n\theta \cot\theta + n^2 \cos\theta = 0 \end{aligned}$$

2. On $[-1,1]$, the Chebyshev polynomials of the second kind (U_n) are orthogonal relative to the weight $w(x) = \sqrt{1-x^2}$.

$$\begin{aligned} &\int_{-1}^{1} U_m(x)\, U_n(x)\, (1-x^2)\, \frac{dx}{\sqrt{1-x^2}} \\ &\qquad = \int_0^\pi \frac{\sin(m+1)\theta}{\sin\theta} \cdot \frac{\sin(n+1)\theta}{\sin\theta} \cdot \sin^2\theta\, d\theta = \begin{cases} 0, & m \neq n \\ \pi/2, & m = n. \end{cases} \end{aligned}$$

While we're at it, notice that

$$T_n'(x) = \frac{n \sin n\theta}{\sin\theta} = n\, U_{n-1}(x).$$

As a rule, the derivatives of a sequence of orthogonal polynomials are again orthogonal polynomials, but relative to a different weight.

3. On $[-1,1]$ with weight $w(x) \equiv 1$, the sequence (P_n) of *Legendre polynomials* are orthogonal, and are typically normalized by $P_n(1) = 1$. The first few Legendre polynomials are $P_0(x) = 1$, $P_1(x) = x$, $P_2(x) = \frac{3}{2}\,x^2 - \frac{1}{2}$, and $P_3(x) = \frac{5}{2}\,x^3 - \frac{3}{2}\,x$. (Check this!) After we've seen a few more examples, we'll come back and give an explicit formula for P_n.

4. All of the examples we've seen so far are special cases of the following: On $[-1,1]$, consider the weight $w(x) = (1-x)^\alpha(1+x)^\beta$, where $\alpha, \beta > -1$. The corresponding orthogonal polynomials $(P_n^{(\alpha,\beta)})$ are called the *Jacobi polynomials* and are typically normalized by requiring that

$$P_n^{(\alpha,\beta)}(1) = \binom{n+\alpha}{\alpha} \equiv \frac{(\alpha+1)(\alpha+2)\cdots(\alpha+n)}{n!}.$$

It follows that $P_n^{(0,0)} = P_n$,

$$P_n^{(-1/2,-1/2)} = \frac{1\cdot 3\cdot 5\cdots(2n-1)}{2^n n!}\, T_n,$$

and

$$P_n^{(1/2,1/2)} = \frac{1\cdot 3\cdot 5\cdots(2n+1)}{2^n (n+1)!}\, U_n.$$

The polynomials $P_n^{(\alpha,\alpha)}$ are called *ultraspherical* polynomials.

5. There are also several classical examples of orthogonal polynomials on unbounded intervals. In particular,

$(0,\infty)$	$w(x) = e^{-x}$	Laguerre polynomials,
$(0,\infty)$	$w(x) = x^\alpha e^{-x}$	generalized Laguerre polynomials,
$(-\infty,\infty)$	$w(x) = e^{-x^2}$	Hermite polynomials.

Since Q_n is orthogonal to every element of $\mathcal{P}_{n-1}$, a fuller understanding of Q_n will follow from a characterization of the *orthogonal complement* of $\mathcal{P}_{n-1}$. We begin with an easy fact about least-squares approximations in inner product spaces.

Lemma. *Let E be a finite dimensional subspace of an inner product space X, and let $x \in X \setminus E$. Then, $y^* \in E$ is the least-squares approximation to x out of E (a.k.a. the nearest point to x in E) if and only if $\langle x - y^*, y \rangle = 0$ for every $y \in E$; that is, if and only if $(x - y^*) \perp E$.*

Proof. [We've taken E to be finite dimensional so that nearest points will exist; since X is an inner product space, nearest points must also be unique (see the exercises for a proof that every inner product norm is strictly convex).]

($\Longleftarrow$) First suppose that $(x - y^*) \perp E$. Then, given any $y \in E$, we have

$$\|x - y\|_2^2 \;=\; \|(x - y^*) + (y^* - y)\|_2^2 \;=\; \|x - y^*\|_2^2 + \|y^* - y\|_2^2,$$

because $y^* - y \in E$ and, hence, $(x - y^*) \perp (y^* - y)$. Thus, $\|x - y\| > \|x - y^*\|$ unless $y = y^*$; that is, y^* is the (unique) nearest point to x in E.

($\Longrightarrow$) Suppose that $x - y^*$ is *not* orthogonal to E. Then, there is some $y \in E$ with $\|y\| = 1$ such that $\alpha = \langle x - y^*, y \rangle \neq 0$. Now I claim that $y^* + \alpha y \in E$ is a better approximation to x than y^* (and $y^* + \alpha y \neq y^*$, of course); that is, y^* is *not* the least-squares approximation to x. To see this, we again compute:

$$\begin{aligned}
\|x - (y^* + \alpha y)\|_2^2 \;&=\; \|(x - y^*) - \alpha y\|_2^2 \;=\; \langle (x - y^*) - \alpha y, \,(x - y^*) - \alpha y \rangle \\
&=\; \|x - y^*\|_2^2 - 2\alpha \,\langle x - y^*, y \rangle + \alpha^2 \\
&=\; \|x - y^*\|_2^2 - \alpha^2 \;<\; \|x - y^*\|_2^2.
\end{aligned}$$

Thus, we must have $\langle x - y^*, y \rangle = 0$ for every $y \in E$. $\square$

Lemma 1. (Integration by-parts.)

$$\int_a^b u^{(n)} v \;=\; \sum_{k=1}^{n} (-1)^{k-1} u^{(n-k)} v^{(k-1)} \Big]_a^b \;+\; (-1)^n \int_a^b u v^{(n)}.$$

Now if v is a polynomial of degree $< n$, then $v^{(n)} = 0$ and we get:

Lemma 2. *$f \in C[\,a, b\,]$ satisfies $\displaystyle\int_a^b f(x)\, p(x)\, w(x)\, dx = 0$ for all polynomials $p \in \mathcal{P}_{n-1}$ if and only if there is an n-times differentiable function u on $[\,a, b\,]$ satisfying $fw = u^{(n)}$ and $u^{(k)}(a) = u^{(k)}(b) = 0$ for all $k = 0, 1, \ldots, n-1$.*

Proof. One direction is clear from Lemma 1: Given u as above, we would have $\int_a^b fpw = \int_a^b u^{(n)} p = (-1)^n \int_a^b u p^{(n)} = 0$.

So, suppose we have that $\int_a^b fpw = 0$ for all $p \in \mathcal{P}_{n-1}$. By integrating fw repeatedly, choosing constants appropriately, we may define a function u satisfying $fw = u^{(n)}$ and $u^{(k)}(a) = 0$ for all $k = 0, 1, \ldots, n-1$. We want to show that the hypotheses on f force $u^{(k)}(b) = 0$ for all $k = 0, 1, \ldots, n-1$.

Now Lemma 1 tells us that

$$0 \;=\; \int_a^b fpw \;=\; \sum_{k=1}^{n} (-1)^{k-1} u^{(n-k)}(b)\, p^{(k-1)}(b)$$

for all $p \in \mathcal{P}_{n-1}$. But the numbers $p(b), p'(b), \ldots, p^{(n-1)}(b)$ are completely arbitrary; that is (again by integrating repeatedly, choosing our constants as we please), we can find polynomials p_k of degree $k < n$ such that $p_k^{(k)}(b) \neq 0$ and $p_k^{(j)}(b) = 0$ for $j \neq k$. In fact, $p_k(x) = (x - b)^k$ works just fine! In any case, we must have $u^{(k)}(b) = 0$ for all $k = 0, 1, \ldots, n-1$. $\square$

Rolle's theorem tells us a bit more about the functions orthogonal to $\mathcal{P}_{n-1}$:

Lemma 3. *If $w(x) > 0$ in (a, b), and if $f \in C[\,a, b\,]$ satisfies $\displaystyle\int_a^b f(x)\, p(x)\, w(x)\, dx = 0$ for all polynomials $p \in \mathcal{P}_{n-1}$, then f has at least n distinct zeros in the open interval (a, b).*

Proof. Write $fw = u^{(n)}$, where $u^{(k)}(a) = u^{(k)}(b) = 0$ for all $k = 0, 1, \ldots, n-1$. In particular, since $u(a) = u(b) = 0$, Rolle's theorem tells us that u' would have at least one zero in (a, b). But then $u'(a) = u'(c) = u'(b) = 0$, and so u'' must have at least two zeros in (a, b). Continuing, we find that $fw = u^{(n)}$ must have at least n zeros in (a, b). Since $w > 0$, the result follows. $\square$

Corollary. *Let (Q_n) be the sequence of orthogonal polynomials associated to a given weight w with $w > 0$ in (a, b). Then, the roots of Q_n are real, simple, and lie in (a, b).*

Lemma 4. *If p^* is the least-squares approximation to $f \in C[\,a, b\,]$ out of $\mathcal{P}_{n-1}$, and if $w > 0$ in (a, b), then $f - p^*$ has at least n distinct zeros in (a, b).*

Proof. The least-squares approximation satisfies $\langle\, f - p^*, p\,\rangle = 0$ for all $p \in \mathcal{P}_{n-1}$. $\square$

The sheer volume of literature on orthogonal polynomials and other "special functions" is truly staggering. We'll content ourselves with the Legendre and the Chebyshev polynomials. In particular, let's return to the problem of finding an explicit formula for the Legendre polynomials. We could, as Rivlin does, use induction and a few observations that simplify the basic recurrence formula (you're encouraged to read this; see pp. 53–54). Instead we'll give a simple (but at first sight intimidating) formula that is of use in more general settings than ours.

Lemma 2 (with $w \equiv 1$ and $[a, b] = [-1, 1]$) says that if we want to find a *polynomial* f of degree n that is orthogonal to $\mathcal{P}_{n-1}$, then we'll need to take a *polynomial* for u, and this u will have to be divisible by $(x-1)^n(x+1)^n$. (Why?) That is, we must have $P_n(x) = c_n \cdot D^n\big[(x^2-1)^n\big]$, where D denotes differentiation, and where we find c_n by evaluating the right-hand side at $x = 1$.

Lemma 5. (Leibniz's formula) $D^n(fg) \;=\; \displaystyle\sum_{k=0}^{n} \binom{n}{k} D^k(f)\, D^{n-k}(g)$.

PROOF. Induction and the fact that $\binom{n-1}{k-1} + \binom{n-1}{k} = \binom{n}{k}$. $\square$

Consequently, $Q(x) = D^n\big[(x-1)^n(x+1)^n\big] = \sum_{k=0}^{n}\binom{n}{k} D^k(x-1)^n\, D^{n-k}(x+1)^n$ and it follows that $Q(1) = 2^n n!$ and $Q(-1) = (-1)^n 2^n n!$. This, finally, gives us the formula discovered by Rodrigues in 1814:

$$P_n(x) \;=\; \frac{1}{2^n n!}\, D^n\big[(x^2-1)^n\big].$$

The Rodrigues formula is quite useful (and easily generalizes to the Jacobi polynomials).

Observations

6. By Lemma 3, the roots of P_n are real, distinct, and lie in $(-1, 1)$.

7. $(x^2-1)^n = \sum_{k=0}^{n}(-1)^k\binom{n}{k}x^{2n-2k}$. If we apply $\frac{1}{2^n n!}D^n$ and simplify, we get another formula for the Legendre polynomials.

$$P_n(x) \;=\; \frac{1}{2^n}\sum_{k=0}^{[n/2]}(-1)^k\binom{n}{k}\binom{2n-2k}{n}x^{n-2k}.$$

In particular, if n is even (odd), then P_n is even (odd). Notice, too, that if we let $\widetilde{P}_n$ denote the polynomial given by the standard construction, then we must have $P_n = 2^{-n}\binom{2n}{n}\widetilde{P}_n$.

8. In terms of our standard recurrence formula, it follows that $a_n = 0$ (because $xP_n(x)^2$ is always odd). It remains to compute b_n. First, integrating by parts,

$$\int_{-1}^{1} P_n(x)^2\,dx \;=\; xP_n(x)^2\Big]_{-1}^{1} \;-\; \int_{-1}^{1} x\cdot 2P_n(x)\,P_n'(x)\,dx,$$

or $\langle\, P_n, P_n\,\rangle \;=\; 2 - 2\langle\, P_n, xP_n'\,\rangle$. But $xP_n' \;=\; nP_n$ + lower degree terms; hence, $\langle\, P_n, xP_n'\,\rangle \;=\; n\langle\, P_n, P_n\,\rangle$. Thus, $\langle\, P_n, P_n\,\rangle \;=\; 2/(2n+1)$. Using this and the fact that $P_n = 2^{-n}\binom{2n}{n}\widetilde{P}_n$, we'd find that $b_n = n^2/(4n^2-1)$. Thus,

$$\begin{aligned} P_{n+1} \;&=\; 2^{-n-1}\binom{2n+2}{n+1}\widetilde{P}_{n+1} \;=\; 2^{-n-1}\binom{2n+2}{n+1}\left[x\,\widetilde{P}_n \;-\; \frac{n^2}{(4n^2-1)}\,\widetilde{P}_{n-1}\right] \\ &=\; \frac{2n+1}{n+1}\,x\,P_n \;-\; \frac{n}{n+1}\,P_{n-1}. \end{aligned}$$

That is, the Legendre polynomials satisfy the recurrence formula

$$(n+1)\,P_{n+1}(x) \;=\; (2n+1)\,x\,P_n(x) - n\,P_{n-1}(x).$$

9. It follows from **8** that the sequence $\widehat{P}_n = \sqrt{\frac{2n+1}{2}}\,P_n$ is *orthonormal* on $[-1,1]$.

10. The Legendre polynomials satisfy $(1-x^2)\,P_n''(x) - 2x\,P_n'(x) + n\,(n+1)\,P_n(x) = 0$. If we set $u = (x^2-1)^n$; that is, if $u^{(n)} = 2^n n! P_n$, note that $u'(x^2-1) = 2nxu$. Now we apply D^{n+1} to both sides of this last equation (using Leibniz's formula) and simplify:

$$\begin{aligned} &u^{(n+2)}(x^2-1) + (n+1)\,u^{(n+1)}\,2x + \frac{(n+1)n}{2}\,u^{(n)}\,2 = 2n\left[u^{(n+1)}\,x + (n+1)\,u^{(n)}\right] \\ &\implies\; (1-x^2)\,u^{(n+2)} - 2x\,u^{(n+1)} + n\,(n+1)\,u^{(n)} = 0. \end{aligned}$$

11. Through a series of exercises, similar in spirit to **10**, Rivlin shows that $|P_n(x)| \le 1$ on $[-1,1]$. See pp. 63–64 of Rivlin for details.

Given an orthogonal sequence, it makes sense to consider "generalized Fourier series" relative to the sequence and to find analogues of the Dirichlet kernel, Lebesgue's theorem, and so on. In case of the Legendre polynomials we have the following:

Example. *The "Fourier-Legendre" series for $f \in C[-1,1]$ is given by $\sum_k \langle\, f, \widehat{P}_k\,\rangle\,\widehat{P}_k$, where*

$$\widehat{P}_k = \sqrt{\frac{2k+1}{2}}\,P_k \qquad \text{and} \qquad \langle\, f, \widehat{P}_k\,\rangle = \int_{-1}^{1} f(x)\,\widehat{P}_k(x)\,dx.$$

The partial sum operator $S_n(f) = \sum_{k=0}^{n} \langle\, f, \widehat{P}_k \,\rangle\, \widehat{P}_k$ *is a linear projection onto* $\mathcal{P}_n$ *and may be written as*

$$S_n(f)(x) = \int_{-1}^{1} f(t)\, K_n(t,x)\, dt,$$

where $K_n(t,x) = \sum_{k=0}^{n} \widehat{P}_k(t)\, \widehat{P}_k(x)$. (Why?)

Since the $\widehat{P}_k$'s are orthonormal, we have

$$\sum_{k=0}^{n} |\langle\, f, \widehat{P}_k \,\rangle|^2 \;=\; \|S_n(f)\|_2^2 \;\le\; \|f\|_2^2 \;=\; \sum_{k=0}^{\infty} |\langle\, f, \widehat{P}_k \,\rangle|^2,$$

and so the generalized Fourier coefficients $\langle\, f, \widehat{P}_k \,\rangle$ are square summable; in particular, $\langle\, f, \widehat{P}_k \,\rangle \to 0$ as $k \to \infty$. As in the case of Fourier series, the fact that the polynomials (i.e., the span of the $\widehat{P}_k$'s) are dense in $C[\,a,b\,]$ implies that $S_n(f)$ actually converges to f in the $\|\cdot\|_2$ norm. These same observations remain valid for any sequence of orthogonal polynomials. The real question remains, just as with Fourier series, whether $S_n(f)$ is a good uniform (or even pointwise) approximation to f.

If you're willing to swallow the fact that $|P_n(x)| \le 1$, then

$$|K_n(t,x)| \;\le\; \sum_{k=0}^{n} \sqrt{\frac{2k+1}{2}}\sqrt{\frac{2k+1}{2}} \;=\; \frac{1}{2}\sum_{k=0}^{n}(2k+1) \;=\; \frac{(n+1)^2}{2}.$$

Hence, $\|S_n(f)\| \le (n+1)^2 \|f\|$. That is, the "Lebesgue numbers" for this process are at most $(n+1)^2$. The analogue of Lebesgue's theorem in this case would then read:

$$\|f - S_n(f)\| \;\le\; Cn^2 E_n(f).$$

Thus, $S_n(f) \rightrightarrows f$ whenever $n^2 E_n(f) \to 0$, and Jackson's theorem tells us when this will happen: *If* f *is twice continuously differentiable, then the Fourier-Legendre series for* f *converges uniformly to* f *on* $[-1,1\,]$.

The Christoffel-Darboux Identity

It would also be of interest to have a closed form for $K_n(t,x)$. That this is indeed always possible, for any sequence of orthogonal polynomials, is a very important fact.

Using our original notation, let (Q_n) be the sequence of monic orthogonal polynomials corresponding to a given weight w, and let $(\widehat{Q}_n)$ be the *orthonormal* counterpart of (Q_n);

in other words, $Q_n = \lambda_n \widehat{Q}_n$, where $\lambda_n = \sqrt{\langle\, Q_n, Q_n \,\rangle}$. It will help things here if you recall (from Observation 1 on page 112) that $\lambda_n^2 = b_n \lambda_{n-1}^2$.

As with the Legendre polynomials, each $f \in C[\,a, b\,]$ is represented by the generalized Fourier series $\sum_k \langle\, f, \widehat{Q}_k \,\rangle\, \widehat{Q}_k$, with partial sum operator

$$S_n(f)(x) = \int_a^b f(t)\, K_n(t, x)\, w(t)\, dt,$$

where $K_n(t, x) = \sum_{k=0}^n \widehat{Q}_k(t)\, \widehat{Q}_k(x)$. As before, S_n is a projection onto $\mathcal{P}_n$; in particular, $S_n(1) = 1$ for every n.

Theorem. (Christoffel-Darboux) *The kernel $K_n(t, x)$ can be written*

$$\sum_{k=0}^n \widehat{Q}_k(t)\, \widehat{Q}_k(x) \;=\; \lambda_{n+1} \lambda_n^{-1}\, \frac{\widehat{Q}_{n+1}(t)\, \widehat{Q}_n(x) - \widehat{Q}_n(t)\, \widehat{Q}_{n+1}(x)}{t - x}.$$

Proof. We begin with the standard recurrence formulas

$$Q_{n+1}(t) = (t - a_n)\, Q_n(t) - b_n Q_{n-1}(t)$$
$$Q_{n+1}(x) = (x - a_n)\, Q_n(x) - b_n Q_{n-1}(x)$$

(where $b_0 = 0$). Multiplying the first by $Q_n(x)$, the second by $Q_n(t)$, and subtracting:

$$\begin{aligned} &Q_{n+1}(t)\, Q_n(x) - Q_n(t)\, Q_{n+1}(x) \\ &\qquad = \; (t - x)\, Q_n(t)\, Q_n(x) \; + \; b_n \big[\, Q_n(t)\, Q_{n-1}(x) \; - \; Q_n(x)\, Q_{n-1}(t) \,\big] \end{aligned}$$

(and again, $b_0 = 0$). If we divide both sides of this equation by λ_n^2 we get

$$\begin{aligned} &\lambda_n^{-2} \big[\, Q_{n+1}(t)\, Q_n(x) - Q_n(t)\, Q_{n+1}(x) \,\big] \\ &\qquad = \; (t - x)\, \widehat{Q}_n(t)\, \widehat{Q}_n(x) \; + \; \lambda_{n-1}^{-2} \big[\, Q_n(t)\, Q_{n-1}(x) \; - \; Q_n(x)\, Q_{n-1}(t) \,\big]. \end{aligned}$$

Thus, we may repeat the process; arriving finally at

$$\lambda_n^{-2} \big[\, Q_{n+1}(t)\, Q_n(x) - Q_n(t)\, Q_{n+1}(x) \,\big] \;=\; (t - x) \sum_{k=0}^n \widehat{Q}_n(t)\, \widehat{Q}_n(x).$$

The Christoffel-Darboux identity now follows by writing $Q_n = \lambda_n \widehat{Q}_n$, etc. $\square$

And we now have a version of the Dini-Lipschitz theorem:

Theorem. *Let $f \in C[\,a,b\,]$ and suppose that at some point x_0 in $[\,a,b\,]$ we have*

(i) *f is Lipschitz at x_0; that is, $|f(x_0) - f(x)| \le K|x_0 - x|$ for some constant K and all x in $[\,a,b\,]$; and*

(ii) *the sequence $(\widehat{Q}_n(x_0))$ is bounded.*

Then, the series $\sum_k \langle\, f, \widehat{Q}_k\,\rangle\, \widehat{Q}_k(x_0)$ converges to $f(x_0)$.

Proof. First note that the sequence $\lambda_{n+1}\lambda_n^{-1}$ is bounded: Indeed, by Cauchy-Schwarz,

$$\begin{aligned}\lambda_{n+1}^2 &= \langle\, Q_{n+1}, Q_{n+1}\,\rangle \;=\; \langle\, Q_{n+1}, x\,Q_n\,\rangle \\ &\le \|Q_{n+1}\|_2 \cdot \|\,x\,\| \cdot \|Q_n\|_2 \;=\; \max\{|a|,|b|\}\,\lambda_{n+1}\lambda_n.\end{aligned}$$

Thus, $\lambda_{n+1}\lambda_n^{-1} \le c = \max\{|a|,|b|\}$. Now, using the Christoffel-Darboux identity,

$$\begin{aligned}S_n(f)(x_0) - f(x_0) &= \int_a^b \big[\, f(t) - f(x_0)\,\big]\, K_n(t,x_0)\, w(t)\,dt \\ &= \lambda_{n+1}\lambda_n^{-1} \int_a^b \frac{f(t) - f(x_0)}{t - x_0}\,\big[\,\widehat{Q}_{n+1}(t)\,\widehat{Q}_n(x_0) - \widehat{Q}_n(t)\,\widehat{Q}_{n+1}(x_0)\,\big]\, w(t)\,dt \\ &= \lambda_{n+1}\lambda_n^{-1}\,\big[\,\langle\, h, \widehat{Q}_{n+1}\,\rangle\,\widehat{Q}_n(x_0) - \langle\, h, \widehat{Q}_n\,\rangle\,\widehat{Q}_{n+1}(x_0)\,\big],\end{aligned}$$

where $h(t) = (f(t) - f(x_0))/(t - x_0)$. But h is bounded (and continuous everywhere except, possibly, at x_0) by hypothesis (i), $\lambda_{n+1}\lambda_n^{-1}$ is bounded, and $\widehat{Q}_n(x_0)$ is bounded by hypothesis (ii). All that remains is to notice that the numbers $\langle\, h, \widehat{Q}_n\,\rangle$ are the generalized Fourier coefficients of the bounded, Riemann integrable function h, and so must tend to zero (since, in fact, they're even square summable). $\square$

We end this section with a negative result, due to Nikolaev:

Theorem. *There is no weight w such that every $f \in C[\,a,b\,]$ has a uniformly convergent expansion in terms of orthogonal polynomials. In fact, given any w, there is always some f for which $\|f - S_n(f)\|$ is unbounded.*

Throughout, w denotes a fixed, positive (except possibly at finitely many points), Riemann integrable weight function on $[\,a,b\,]$, and we consider the inner product on $C[\,a,b\,]$ defined by

$$\langle f,g\rangle \;=\; \int_a^b f(x)\,g(x)\,w(x)\,dx,$$

and the associated (strictly convex) norm

$$\|f\|_2 \;=\; \sqrt{\langle f,f\rangle} \;=\; \left(\int_a^b |f(x)|^2\,w(x)\,dx\right)^{1/2}.$$

62. Prove that every inner product norm is strictly convex. Specifically, let $\langle\cdot,\cdot\rangle$ be an inner product on a vector space X, and let $\|x\| = \sqrt{\langle x,x\rangle}$ be the associated norm. Show that:

(a) $\|x+y\|^2 + \|x-y\|^2 = 2\,(\|x\|^2 + \|y\|^2)$ for all x, $y \in X$ (the parallelogram identity).

(b) If $\|x\| = r = \|y\|$ and if $\|x-y\| = \delta$, then $\left\|\frac{x+y}{2}\right\|^2 = r^2 - (\delta/2)^2$. In particular, $\left\|\frac{x+y}{2}\right\| < r$ whenever $x \neq y$.

We define a sequence of polynomials (Q_n) which are mutually orthogonal, relative to w, by setting $Q_0(x) = 1$, $Q_1(x) = x - a_0$, and

$$Q_{n+1}(x) = (x-a_n)Q_n(x) - b_nQ_{n-1}(x), \quad \text{for } n \geq 1, \text{ where}$$

$$a_n = \langle\, x\,Q_n, Q_n\,\rangle / \langle\, Q_n, Q_n\,\rangle \quad \text{and} \quad b_n = \langle\, x\,Q_n, Q_{n-1}\,\rangle / \langle\, Q_{n-1}, Q_{n-1}\,\rangle$$

(and where $x\,Q_n$ is shorthand for the polynomial $x\,Q_n(x)$).

63. Check that Q_n is a monic polynomial of degree exactly n.

64. If (P_n) is another sequence of orthogonal polynomials such that P_n has degree exactly n, for each n, show that $P_n = \alpha_nQ_n$ for some $\alpha_n \neq 0$. In particular, if P_n is a monic polynomial, then $P_n = Q_n$. [Hint: Choose α_n so that $P_n - \alpha_nQ_n \in \mathcal{P}_{n-1}$ and note that $(P_n - \alpha_nQ_n) \perp \mathcal{P}_{n-1}$. Conclude that $P_n - \alpha_nQ_n = 0$.]

65. Check that $\langle\, x\,Q_n, Q_{n-1}\,\rangle = \langle\, Q_n, Q_n\,\rangle$, and conclude that $b_n > 0$ for each n.

66. Given $f \in C[\,a,b\,]$ and $n \ge 0$, prove that $q_n^* \in \mathcal{P}_n$ is the least-squares approximation to f out of $\mathcal{P}_n$ (with respect to w) if and only if

$$\langle f - q_n^*\,,\,p\,\rangle = \int_a^b (f(x) - q_n^*(x))\,p(x)\,w(x)\,dx = 0$$

for every $p \in \mathcal{P}_n$; that is, if and only if $(f - q_n^*) \perp \mathcal{P}_n$.

67. If $f \in C[\,a,b\,]$ but $f \notin \mathcal{P}_n$, show that $f - q_n^*$ changes sign at $n+1$ (or more) points in (a,b). [Hint: If not, show that there is a polynomial $p \in \mathcal{P}_n$ such that $(f - q_n^*)\,p \ge 0$ (but $(f - q_n^*)\,p \ne 0$) in (a,b). Now appeal to the result in problem **66** to arrive at a contradiction.]

68. Show that the least-squares approximation to $f(x) = x^n$ out of $\mathcal{P}_{n-1}$ (relative to w) is $q_{n-1}^*(x) = x^n - Q_n(x)$.

69. Show that Q_n has n distinct, simple zeros in (a,b). [Hint: Combine **67** and **68**.]

70. Given $f \in C[\,a,b\,]$, let p_n^* denote the best *uniform* approximation to f out of $\mathcal{P}_n$ and let q_n^* denote the least-squares approximation to f out of $\mathcal{P}_n$. Show that $\|f - q_n^*\|_2 \le \|f - p_n^*\|_2$ and conclude that $\|f - q_n^*\|_2 \to 0$ as $n \to \infty$.

71. Show that the Chebyshev polynomials of the first kind, (T_n), and of the second kind, (U_n), satisfy the identities

$$T_n(x) = U_n(x) - x\,U_{n-1}(x) \quad \text{and} \quad (1 - x^2)\,U_{n-1}(x) = x\,T_n(x) - T_{n+1}(x).$$

72. Show that the Chebyshev polynomials of the second kind, (U_n), satisfy the recurrence relation

$$U_{n+1}(x) = 2x\,U_n(x) - U_{n-1}(x), \quad n \ge 1,$$

where $U_0(x) = 1$ and $U_1(x) = 2x$. [Please compare this with the recurrence relation satisfied by the T_n's!]

Numerical integration, or *quadrature*, is the process of approximating the value of a definite integral $\int_a^b f(x)\,w(x)\,dx$ based only on a finite number of values or "samples" of f (much like a Riemann sum). A *linear quadrature formula* takes the form

$$\int_a^b f(x)\,w(x)\,dx \;\approx\; \sum_{k=1}^n A_k f(x_k),$$

where the *nodes* (x_k) and the *weights* (A_k) are at our disposal. (Note that both sides of the formula are linear in f.)

Example. Consider the quadrature formula

$$I(f) \;=\; \int_{-1}^1 f(x)\,dx \;\approx\; \frac{1}{n}\sum_{k=-n}^{n-1} f\left(\frac{2k+1}{2n}\right) \;=\; I_n(f).$$

If f is continuous, then we clearly have $I_n(f) \to \int_{-1}^1 f$ as $n \to \infty$. (Why?) But in the particular case $f(x) = x^2$ we have (after some simplification)

$$I_n(f) \;=\; \frac{1}{n}\sum_{k=-n}^{n-1}\left(\frac{2k+1}{2n}\right)^2 \;=\; \frac{1}{2n^3}\sum_{k=0}^{n-1}(2k+1)^2 \;=\; \frac{2}{3} - \frac{1}{6n^2}.$$

That is, $|\,I_n(f) - I(f)\,| = 1/6n^2$. In particular, we would need to take $n \ge 130$ to get $1/6n^2 \le 10^{-5}$, for example, and this would require that we perform over 250 evaluations of f. We'd like a method that converges a bit faster! In other words, there's no shortage of quadrature formulas—we just want *faster* ones.

One reasonable requirement for our proposed quadrature formula is that it be *exact for polynomials* of low degree. As it happens, this is easy to come by.

Lemma 1. *Given $w(x)$ on $[\,a,b\,]$ and nodes $a \le x_1 < \cdots < x_n \le b$, there exist unique weights $A_1,\ldots,A_n$ such that*

$$\int_a^b p(x)\,w(x)\,dx \;=\; \sum_{i=1}^n A_i\,p(x_i)$$

for all polynomials $p \in \mathcal{P}_{n-1}$.

PROOF. Let $\ell_1, \ldots, \ell_n$ be the Lagrange interpolating polynomials of degree $n-1$ associated to the nodes $x_1, \ldots, x_n$, and recall that we have $p = \sum_{i=1}^n p(x_i)\, \ell_i$ for all $p \in \mathcal{P}_{n-1}$. Hence,

$$\int_a^b p(x)\, w(x)\, dx \;=\; \sum_{i=1}^n p(x_i) \int_a^b \ell_i(x)\, w(x)\, dx.$$

That is, $A_i = \int_a^b \ell_i(x)\, w(x)\, dx$ works. To see that this is the only choice, suppose that

$$\int_a^b p(x)\, w(x)\, dx \;=\; \sum_{i=1}^n B_i\, p(x_i)$$

is exact for all $p \in \mathcal{P}_{n-1}$, and set $p = \ell_j$:

$$A_j \;=\; \int_a^b \ell_j(x)\, w(x)\, dx \;=\; \sum_{i=1}^n B_i\, \ell_j(x_i) \;=\; B_j. \qquad \square$$

The point here is that $\ell_1, \ldots, \ell_n$ form a basis for $\mathcal{P}_{n-1}$ and integration is linear; thus, integration is completely determined by its action on the basis—that is, by the n values $A_i = I(\ell_i)$, $i = 1, \ldots, n$.

Said another way, the n point evaluations $\delta_i(p) = p(x_i)$ satisfy $\mathcal{P}_{n-1} \cap \left(\bigcap_{i=1}^n \ker \delta_i\right) = \{0\}$, and it follows that every linear, real-valued function on $\mathcal{P}_{n-1}$ must be a linear combination of the δ_i's. Here's why: Since the x_i's are distinct, $\mathcal{P}_{n-1}$ may be identified with $\mathbb{R}^n$ by way of the *isomorphism* $p \mapsto (p(x_1), \ldots, p(x_n))$. A linear, real-valued function on $\mathcal{P}_{n-1}$ must, then, correspond to some linear, real-valued function on $\mathbb{R}^n$. In other words, it's given by inner product against some fixed vector $(A_1, \ldots, A_n)$; in particular, we must have $I(p) = \sum_{i=1}^n A_i\, p(x_i)$.

In any case, we now have our quadrature formula: For $f \in C[\,a, b\,]$ we define $I_n(f) = \sum_{i=1}^n A_i\, f(x_i)$, where $A_i = \int_a^b \ell_i(x)\, w(x)\, dx$. But notice that the proof of our last result suggests an alternate way of writing our quadrature formula. Indeed, if $L_{n-1}(f)(x) = \sum_{i=1}^n f(x_i)\ell_i(x)$ is the Lagrange interpolating polynomial for f of degree $n-1$ based on the nodes $x_1, \ldots, x_n$, then

$$\int_a^b (L_{n-1}(f))(x)\, w(x)\, dx \;=\; \sum_{i=1}^n f(x_i) \int_a^b \ell_i(x)\, w(x)\, dx \;=\; \sum_{i=1}^n A_i\, f(x_i).$$

In summary, $I_n(f) = I(L_{n-1}(f)) \approx I(f)$; that is,

$$I_n(f) \;=\; \sum_{i=1}^n A_i\, f(x_i) \;=\; \int_a^b (L_{n-1}(f))(x)\, w(x)\, dx \;\approx\; \int_a^b f(x)\, w(x)\, dx \;=\; I(f),$$

where L_{n-1} is the Lagrange interpolating polynomial of degree $n-1$ based on the nodes $x_1, \ldots, x_n$. This formula is obviously exact for $f \in \mathcal{P}_{n-1}$.

It's easy to give a bound on $|I_n(f)|$ in terms of $\|f\|$; indeed,

$$|I_n(f)| \;\le\; \sum_{i=1}^n |A_i|\, |f(x_i)| \;\le \|f\| \left(\sum_{i=1}^n |A_i| \right).$$

By considering a norm one continuous function f satisfying $f(x_i) = \mathrm{sgn} A_i$ for each $i = 1, \ldots, n$, it's easy to see that $\sum_{i=1}^n |A_i|$ is the smallest constant that works in this inequality. In other words, $\lambda_n = \sum_{i=1}^n |A_i|$, $n = 1, 2, \ldots$, are the "Lebesgue numbers" for this process. As with all previous settings, we want these numbers to be uniformly bounded.

If $w(x) \equiv 1$ and if f is n-times continuously differentiable, we even have an error estimate for our quadrature formula:

$$\left| \int_a^b f - \int_a^b L_{n-1}(f) \right| \;\le\; \int_a^b |f - L_{n-1}(f)| \le\; \frac{1}{n!}\, \|f^{(n)}\| \int_a^b \prod_{i=1}^n |x - x_i|\, dx$$

(recall the Theorem on page 72 of "A Brief Introduction to Interpolation"). As it happens, the integral on the right is minimized when the x_i's are taken to be the zeros of the Chebyshev polynomial U_n (see Rivlin, page 72).

The fact that a quadrature formula is exact for polynomials of low degree does not by itself guarantee that the formula is highly accurate. The problem is that $\sum_{i=1}^n A_i\, f(x_i)$ may be estimating a very small quantity through the cancellation of very large quantities. So, for example, a positive function may yield a negative result in this approximate integral. This wouldn't happen if the A_i's were all positive—and we've already seen how useful positivity can be. Our goal here is to further improve our quadrature formula to have this property. But we have yet to take advantage of the fact that the x_i's are at our disposal. We'll let Gauss show us the way!

Theorem. (Gauss) *Fix a weight $w(x)$ on $[\,a, b\,]$, and let (Q_n) be the canonical sequence of orthogonal polynomials relative to w. Given n, let $x_1, \ldots, x_n$ be the zeros of Q_n (these all* *lie in (a,b)), and choose $A_1,\ldots,A_n$ so that the formula $\sum_{i=1}^n A_i f(x_i) \approx \int_a^b f(x)\,w(x)\,dx$ is exact for polynomials of degree less than n. Then, in fact, the formula is exact for all polynomials of degree less than $2n$.*

Proof. Given a polynomial P of degree less than $2n$, we may divide: $P = Q_n R + S$, where R and S are polynomials of degree less than n. Then,

$$\begin{aligned}\int_a^b P(x)\,w(x)\,dx &= \int_a^b Q_n(x)\,R(x)\,w(x)\,dx + \int_a^b S(x)\,w(x)\,dx \\ &= \int_a^b S(x)\,w(x)\,dx, \quad \text{since deg } R < n \\ &= \sum_{i=1}^n A_i S(x_i), \quad \text{since deg } S < n.\end{aligned}$$

But $P(x_i) = Q_n(x_i)\,R(x_i) + S(x_i) = S(x_i)$, since $Q_n(x_i) = 0$. Hence, $\int_a^b P(x)\,w(x)\,dx = \sum_{i=1}^n A_i P(x_i)$ for all polynomials P of degree less than $2n$. $\square$

Amazing! But, well, not really: $\mathcal{P}_{2n-1}$ is of dimension $2n$, and we had $2n$ numbers $x_1,\ldots,x_n$ and $A_1,\ldots,A_n$ to choose as we saw fit. Said another way, the division algorithm tells us that $\mathcal{P}_{2n-1} = Q_n\mathcal{P}_{n-1} \oplus \mathcal{P}_{n-1}$. Since $Q_n\mathcal{P}_{n-1} \subset \ker(I_n)$, the action of I_n on $\mathcal{P}_{2n-1}$ is the same as its action on a "copy" of $\mathcal{P}_{n-1}$.

In still other words, since any polynomial that vanishes at all the x_i's must be divisible by Q_n (and conversely), we have $Q_n\mathcal{P}_{n-1} = \mathcal{P}_{2n-1} \cap (\bigcap_{i=1}^n \ker \delta_i) = \ker(I_n\,|_{\mathcal{P}_{2n-1}})$. Thus, I_n "factors through" the quotient space $\mathcal{P}_{2n-1}/Q_n\mathcal{P}_{n-1} = \mathcal{P}_{n-1}$.

Also not surprising is that this particular choice of x_i's is unique.

Lemma 2. *Suppose that $a \le x_1 < \cdots < x_n \le b$ and $A_1,\ldots,A_n$ are given so that the equation $\int_a^b P(x)\,w(x)\,dx = \sum_{i=1}^n A_i P(x_i)$ is satisfied for all polynomials P of degree less than $2n$. Then, $x_1,\ldots,x_n$ are the zeros of Q_n.*

Proof. Let $Q(x) = \prod_{i=1}^n (x - x_i)$. Then, for $k < n$, the polynomial $Q \cdot Q_k$ has degree $n + k < 2n$. Hence,

$$\int_a^b Q(x)\,Q_k(x)\,w(x)\,dx = \sum_{i=1}^n A_i Q(x_i)\,Q_k(x_i) = 0.$$

Since Q is a monic polynomial of degree n which is orthogonal to each Q_k, $k < n$, we must have $Q = Q_n$. Thus, the x_i's are actually the zeros of Q_n. $\square$

According to Rivlin, the phrase *Gaussian quadrature* is usually reserved for the specific quadrature formula whereby $\int_{-1}^{1} f(x)\,dx$ is approximated by $\int_{-1}^{1}(L_{n-1}(f))(x)\,dx$, where $L_{n-1}(f)$ is the Lagrange interpolating polynomial to f using the zeros of the n-th Legendre polynomial as nodes. (What a mouthful!) What is actually being described in our version of Gauss's theorem is *Gaussian-type quadrature*.

Before computers, Gaussian quadrature was little more than a curiosity; the roots of Q_n are typically irrational, and certainly not easy to come by. By now, though, it's considered a standard quadrature technique. In any case, we still can't judge the quality of Gauss's method without a bit more information.

Gaussian-type Quadrature

First, let's summarize our rather cumbersome notation.

orthogonal polynomial	zeros	weights	approximate integral
Q_1	$x_1^{(1)}$	$A_1^{(1)}$	I_1
Q_2	$x_1^{(2)}, x_2^{(2)}$	$A_1^{(2)}, A_2^{(2)}$	I_2
Q_3	$x_1^{(3)}, x_2^{(3)}, x_3^{(3)}$	$A_1^{(3)}, A_2^{(3)}, A_3^{(3)}$	I_3
$\vdots$	$\vdots$	$\vdots$	$\vdots$

Hidden here is the Lagrange interpolation formula $L_{n-1}(f) = \sum_{i=1}^{n} f(x_i^{(n)})\,\ell_i^{(n-1)}$, where $\ell_i^{(n-1)}$ denote the Lagrange polynomials of degree $n-1$ based on $x_1^{(n)}, \ldots, x_n^{(n)}$. The n-th quadrature formula is then

$$I_n(f) \;=\; \int_a^b L_{n-1}(f)(x)\,w(x)\,dx \;=\; \sum_{i=1}^{n} A_i^{(n)} f(x_i^{(n)}) \;\approx\; \int_a^b f(x)\,w(x)\,dx,$$

which is exact for polynomials of degree less than $2n$.

By way of one example, Hermite showed that $A_k^{(n)} = \pi/n$ for the Chebyshev weight $w(x) = (1-x^2)^{-1/2}$ on $[-1,1]$. Remarkably, $A_k^{(n)}$ doesn't depend on k! The quadrature formula in this case reads:

$$\int_{-1}^{1} \frac{f(x)\,dx}{\sqrt{1-x^2}} \;\approx\; \frac{\pi}{n} \sum_{k=1}^{n} f\left(\cos\,\frac{2k-1}{2n}\,\pi\right).$$

Or, if you prefer,

$$\int_{-1}^{1} f(x)\,dx \;\approx\; \frac{\pi}{n} \sum_{k=1}^{n} f\left(\cos\,\frac{2k-1}{2n}\,\pi\right) \sin\,\frac{2k-1}{2n}\,\pi.$$

(Why?) You can find full details in Natanson's *Constructive Function Theory*, Vol. III.

The key result, due to Stieltjes, is that I_n is *positive*:

Lemma 3. *$A_1^{(n)}, \ldots, A_n^{(n)} > 0$ and $\sum_{i=1}^n A_i^{(n)} = \int_a^b w(x)\,dx$.*

Proof. The second assertion is obvious (since $I_n(1) = I(1)$). For the first, fix $1 \le j \le n$ and notice that $(\ell_j^{(n-1)})^2$ is of degree $2(n-1) < 2n$. Thus,

$$0 \;<\; \langle\, \ell_j^{(n-1)}, \ell_j^{(n-1)} \,\rangle \;=\; \int_a^b \left[\ell_j^{(n-1)}(x)\right]^2 w(x)\,dx \;=\; \sum_{i=1}^{n} A_i^{(n)} \left[\ell_j^{(n-1)}(x_i^{(n)})\right]^2 \;=\; A_j^{(n)},$$

because $\ell_j^{(n-1)}(x_i^{(n)}) = \delta_{i,j}$. □

Now our last calculation is quite curious; what we've shown is that

$$A_j^{(n)} \;=\; \int_a^b \ell_j^{(n-1)}(x)\,w(x)\,dx \;=\; \int_a^b \left[\ell_j^{(n-1)}(x)\right]^2 w(x)\,dx.$$

The same calculation as above also proves

Corollary. *$\langle\, \ell_i^{(n-1)}, \ell_j^{(n-1)} \,\rangle = 0$ for $i \ne j$.*

Since $A_1^{(n)}, \ldots, A_n^{(n)} > 0$, it follows that $I_n(f)$ is positive; that is, $I_n(f) \ge 0$ whenever $f \ge 0$. The second assertion in Lemma 3 tells us that the I_n's are *uniformly bounded*:

$$|I_n(f)| \;\le\; \|f\| \sum_{i=1}^{n} A_i^{(n)} \;=\; \|f\| \int_a^b w(x)\,dx,$$

and this is the same bound that holds for $I(f) = \int_a^b f(x)\,w(x)\,dx$ itself. Given all of this, proving that $I_n(f) \to I(f)$ is a piece of cake. The following result is again due to Stieltjes (á la Lebesgue).

Theorem. *In the above notation,* $|I_n(f) - I(f)| \le 2\left(\int_a^b w(x)\,dx\right) E_{2n-1}(f)$. *In particular,* $I_n(f) \to I(f)$ *for evey* $f \in C[\,a,b\,]$.

Proof. Let p^* be the best uniform approximation to f out of $\mathcal{P}_{2n-1}$. Then, since $I_n(p^*) = I(p^*)$, we have

$$\begin{aligned} |I(f) - I_n(f)| &\le |I(f-p^*)| \;+\; |I_n(f-p^*)| \\ &\le \|f-p^*\| \int_a^b w(x)\,dx \;+\; \|f-p^*\| \sum_{i=1}^n A_i^{(n)} \\ &= 2\,\|f-p^*\| \int_a^b w(x)\,dx \;=\; 2E_{2n-1}(f) \int_a^b w(x)\,dx. \qquad \square \end{aligned}$$

Computational Considerations

You've probably been asking yourself: "How do I find the A_i's without integrating?" Well, first let's recall the definition: In the case of Gaussian-type quadrature we have

$$A_i^{(n)} \;=\; \int_a^b \ell_i^{(n-1)}(x)\,w(x)\,dx \;=\; \int_a^b \frac{Q_n(x)}{(x - x_i^{(n)})\,Q_n'(x_i^{(n)})}\,w(x)\,dx$$

(because "W" is the same as Q_n here—the x_i's are the zeros of Q_n). Next, consider the function

$$\varphi_n(x) \;=\; \int_a^b \frac{Q_n(t) - Q_n(x)}{t - x}\,w(t)\,dt.$$

Since $t - x$ divides $Q_n(t) - Q_n(x)$, note that φ_n is actually a *polynomial* (of degree at most $n-1$) and that

$$\varphi_n(x_i^{(n)}) \;=\; \int_a^b \frac{Q_n(t)}{t - x_i^{(n)}}\,w(t)\,dt \;=\; A_i^{(n)} Q_n'(x_i^{(n)}).$$

Now $Q_n'(x_i^{(n)})$ is readily available; we just need to compute $\varphi_n(x_i^{(n)})$.

Claim. *The* φ_n*'s satisfy the same recurrence formula as the* Q_n*'s*

$$\varphi_{n+1}(x) = (x - a_n)\varphi_n(x) - b_n\varphi_{n-1}(x), \quad n \ge 1,$$

but with different starting values

$$\varphi_0(x) \equiv 0, \qquad \text{and} \qquad \varphi_1(x) \equiv \int_a^b w(x)\,dx.$$

PROOF. The formulas for φ_0 and φ_1 are obviously correct, since $Q_0(x) \equiv 1$ and $Q_1(x) = x - a_0$. We only need to check the recurrence formula itself.

$$\begin{aligned}
\varphi_{n+1}(x) &= \int_a^b \frac{Q_{n+1}(t) - Q_{n+1}(x)}{t - x}\, w(t)\, dt \\
&= \int_a^b \frac{(t - a_n)\, Q_n(t) - b_n Q_{n-1}(t) - (x - a_n)\, Q_n(x) + b_n Q_{n-1}(x)}{t - x}\, w(t)\, dt \\
&= (x - a_n) \int_a^b \frac{Q_n(t) - Q_n(x)}{t - x}\, w(t)\, dt \;-\; b_n \int_a^b \frac{Q_{n-1}(t) - Q_{n-1}(x)}{t - x}\, w(t)\, dt \\
&= (x - a_n)\, \varphi_n(x) - b_n\, \varphi_{n-1}(x),
\end{aligned}$$

since $\int_a^b Q_n(t)\, w(t)\, dt = 0$. $\square$

Of course, the derivatives Q_n' satisfy a recurrence relation of sorts, too:

$$Q_{n+1}'(x) \;=\; Q_n(x) + (x - a_n)\, Q_n'(x) - b_n\, Q_{n-1}'(x).$$

But $Q_n'(x_i^{(n)})$ can be computed without knowing $Q_n'(x)$. Indeed, $Q_n(x) = \prod_{i=1}^n (x - x_i^{(n)})$, so we have $Q_n'(x_i^{(n)}) = \prod_{j \neq i} (x_i^{(n)} - x_j^{(n)})$.

The weights $A_i^{(n)}$, or *Christoffel numbers*, together with the zeros of Q_n are tabulated in a variety of standard cases. See, for example, *Handbook of Mathematical Functions with Formulas, Graphs, and Tables*, by Abramowitz and Stegun, eds. In practice, of course, it's enough to tabulate data for the case $[\,a, b\,] = [-1, 1\,]$.

Applications to Interpolation

Although $L_n(f)$ isn't typically a good uniform approximation to f, if we interpolate at the zeros of an orthogonal polynomial Q_{n+1}, then $L_n(f)$ will be a good approximation in the $\|\cdot\|_1$ or $\|\cdot\|_2$ norm generated by the corresponding weight w. Specifically, by rewording our earlier results, it's easy to get estimates for each of the errors $\int_a^b |f - L_n(f)|\, w$ and $\int_a^b |f - L_n(f)|^2\, w$. We use essentially the same notation as before, except now we take

$$L_n(f) \;=\; \sum_{i=1}^{n+1} f\big(x_i^{(n+1)}\big)\, \ell_i^{(n)},$$

where $x_1^{(n+1)}, \ldots, x_{n+1}^{(n+1)}$ are the roots of Q_{n+1} and $\ell_i^{(n)}$ is of degree n. This leads to a quadrature formula that's exact on polynomials of degree less than $2(n+1)$.

As we've already seen, $\ell_1^{(n)}, \ldots, \ell_{n+1}^{(n)}$ are *orthogonal* and so $\|L_n(f)\|_2$ may be computed exactly.

Lemma. $\|L_n(f)\|_2 \;\le\; \|f\| \left(\int_a^b w(x)\,dx\right)^{1/2}$.

PROOF. Since $L_n(f)^2$ is a polynomial of degree $\le 2n < 2(n+1)$, we have

$$\begin{aligned}
\|L_n(f)\|_2^2 &= \int_a^b [L_n(f)]^2\, w(x)\,dx \\
&= \sum_{j=1}^{n+1} A_j^{(n+1)} \left[\sum_{i=1}^{n+1} f\big(x_i^{(n+1)}\big)\,\ell_i^{(n)}\big(x_j^{(n+1)}\big)\right]^2 \\
&= \sum_{j=1}^{n+1} A_j^{(n+1)} \left[f\big(x_j^{(n+1)}\big)\right]^2 \\
&\le \|f\|^2 \sum_{j=1}^{n+1} A_j^{(n+1)} \;=\; \|f\|^2 \int_a^b w(x)\,dx. \quad \square
\end{aligned}$$

Please note that we also have $\|f\|_2 \le \|f\| \left(\int_a^b w(x)\,dx\right)^{1/2}$; that is, this same estimate holds for $\|f\|_2$ itself.

As usual, once we have an estimate for the norm of an operator, we also have an analogue of Lebesgue's theorem.

Theorem. $\|f - L_n(f)\|_2 \;\le\; 2\left(\int_a^b w(x)\,dx\right)^{1/2} E_n(f)$.

PROOF. Here we go again! Let p^* be the best uniform approximation to f out of $\mathcal{P}_n$ and use the fact that $L_n(p^*) = p^*$ to see that:

$$\begin{aligned}
\|f - L_n(f)\|_2 &\le \|f - p^*\|_2 + \|L_n(f - p^*)\|_2 \\
&\le \|f - p^*\| \left(\int_a^b w(x)\,dx\right)^{1/2} + \|f - p^*\| \left(\int_a^b w(x)\,dx\right)^{1/2} \\
&= 2E_n(f) \left(\int_a^b w(x)\,dx\right)^{1/2}. \quad \square
\end{aligned}$$

Hence, if we interpolate $f \in C[\,a,b\,]$ at the zeros of (Q_n), then $L_n(f) \to f$ in $\|\cdot\|_2$ norm. The analogous result for the $\|\cdot\|_1$ norm is now easy:

Corollary. $\int_a^b |f(x) - L_n(f)(x)|\, w(x)\, dx \;\le\; 2\left(\int_a^b w(x)\, dx\right) E_n(f).$

PROOF. We apply the Cauchy-Schwarz inequality:

$$\begin{aligned}
\int_a^b |f(x) - L_n(f)(x)|\, w(x)\, dx &= \int_a^b |f(x) - L_n(f)(x)|\, \sqrt{w(x)}\, \sqrt{w(x)}\, dx \\
&\le \left(\int_a^b |f(x) - L_n(f)(x)|^2\, w(x)\, dx\right)^{1/2} \left(\int_a^b w(x)\, dx\right)^{1/2} \\
&\le 2E_n(f) \int_a^b w(x)\, dx. \qquad \square
\end{aligned}$$

Essentially the same device allows an estimate of $\int_a^b f(x)\, dx$ in terms of $\int_a^b f(x)\, w(x)\, dx$ (which may be easier to compute).

Corollary. *If $\int_a^b w(x)^{-1}\, dx$ is finite, then*

$$\begin{aligned}
\int_a^b |f(x) - L_n(f)(x)|\, dx &= \int_a^b |f(x) - L_n(f)(x)|\, \sqrt{w(x)}\, \frac{1}{\sqrt{w(x)}}\, dx \\
&\le \left(\int_a^b |f(x) - L_n(f)(x)|^2\, w(x)\, dx\right)^{1/2} \left(\int_a^b \frac{1}{w(x)}\, dx\right)^{1/2} \\
&\le 2E_n(f) \left(\int_a^b w(x)\, dx\right)^{1/2} \left(\int_a^b \frac{1}{w(x)}\, dx\right)^{1/2}. \qquad \square
\end{aligned}$$

In particular, the Chebyshev weight satisfies

$$\int_{-1}^1 \frac{dx}{\sqrt{1-x^2}} = \pi \qquad \text{and} \qquad \int_{-1}^1 \sqrt{1-x^2}\, dx = \frac{\pi}{2}.$$

Thus, interpolation at the zeros of the Chebyshev polynomials (of the first kind) would provide good, *simultaneous* approximation in each of the norms $\|\cdot\|_1$, $\|\cdot\|_2$, and $\|\cdot\|$.

The Moment Problem

Given a positive, continuous weight function $w(x)$ on $[\,a, b\,]$, the number

$$\mu_k \;=\; \int_a^b x^k\, w(x)\, dx$$

is called the k-th *moment* of w. In physical terms, if we think of $w(x)$ as the density of a thin rod placed on the interval $[\,a, b\,]$, then μ_0 is the mass of the rod, μ_1/μ_0 is its center of mass, μ_2 is its moment of inertia (about 0), and so on. In probabilistic terms, if $\mu_0 = 1$, then w is the probability density function for some random variable, μ_1 is the expected value, or mean, of this random variable, and $\mu_2 - \mu_1^2$ is its variance. The moment problem (or problems, really) concern the inverse procedure. What can be measured in real life are the moments—can the moments be used to find the density function?

Questions: Do the moments determine w? Do different weights have different moment sequences? If we knew the sequence (μ_k), could we find w? How do we tell if a given sequence (μ_k) is the moment sequence for some positive weight? Do "special" weights give rise to "special" moment sequences?

Now we've already answered one of these questions: The Weierstrass theorem tells us that different weights have different moment sequences. Said another way, if

$$\int_a^b x^k \, w(x) \, dx = 0 \qquad \text{for all } k = 0, 1, 2, \ldots,$$

then $w \equiv 0$. Indeed, by linearity, this says that $\int_a^b p(x) \, w(x) \, dx = 0$ for all polynomials p which, in turn, tells us that $\int_a^b w(x)^2 \, dx = 0$. (Why?) The remaining questions are harder to answer. We'll settle for simply stating a few pertinent results.

Given a sequence of numbers (μ_k), we define the n-th difference sequence $(\Delta^n \mu_k)$ by

$$\begin{aligned} \Delta^0 \mu_k &= \mu_k \\ \Delta^1 \mu_k &= \mu_k - \mu_{k+1} \\ \Delta^n \mu_k &= \Delta^{n-1} \mu_k - \Delta^{n-1} \mu_{k+1}, \quad n \geq 1. \end{aligned}$$

For example, $\Delta^2 \mu_k = \mu_k - 2\mu_{k+1} + \mu_{k+2}$. More generally, induction will show that

$$\Delta^n \mu_k = \sum_{i=0}^{n} (-1)^i \binom{n}{i} \mu_{k+i}.$$

In the case of a weight w on the interval $[\,0, 1\,]$, this sum is easy to recognize as an integral. Indeed,

$$\int_0^1 x^k (1-x)^n \, w(x) \, dx = \sum_{i=0}^{n} (-1)^i \binom{n}{i} \int_0^1 x^{k+i} \, w(x) \, dx = \sum_{i=0}^{n} (-1)^i \binom{n}{i} \mu_{k+i}.$$

In particular, if w is nonnegative, then we must have $\Delta^n \mu_k \geq 0$ for every n and k. This observation serves as motivation for

Theorem. *The following are equivalent:*

(a) (μ_k) *is the moment sequence of some nonnegative weight function* w *on* $[0,1]$.

(b) $\Delta^n \mu_k \ge 0$ *for every* n *and* k.

(c) $a_0\mu_0 + a_1\mu_1 + \cdots + a_n\mu_n \ge 0$ *whenever* $a_0 + a_1 x + \cdots + a_n x^n \ge 0$ *for all* $0 \le x \le 1$.

The equivalence of (a) and (b) is due to Hausdorff. A real sequence satisfying (b) or (c) is sometimes said to be *positive definite*.

Now dozens of mathematicians worked on various aspects of the moment problem: Chebyshev, Markov, Stieltjes, Cauchy, Riesz, Fréchet, and on and on. And several of them, in particular Cauchy and Stieltjes, noticed the importance of the integral $\int_a^b \frac{w(t)}{x-t}\,dt$ in attacking the problem. (Compare this expression to Cauchy's integral formula.) It was Stieltjes, however, who gave the first complete solution to such a problem—developing his own integral (by considering $\int_a^b \frac{dW(t)}{x-t}$), his own variety of continued fractions, and planting the seeds for the study of orthogonal polynomials while he was at it! We will attempt to at least sketch a few of these connections.

To begin, let's fix our notation: To simplifiy things, we suppose that we're given a nonnegative weight $w(x)$ on a *symmetric* interval $[-a,a]$, and that all of the moments of w are finite. We will otherwise stick to our usual notations for (Q_n), the Gaussian-type quadrature formulas, and so on. Next, we consider *the moment-generating function*:

Lemma. *If* $x \notin [-a,a]$, *then* $\displaystyle\int_{-a}^{a} \frac{w(t)}{x-t}\,dt \;=\; \sum_{k=0}^{\infty} \frac{\mu_k}{x^{k+1}}.$

Proof. $\displaystyle\frac{1}{x-t} = \frac{1}{x}\cdot\frac{1}{1-(t/x)} = \sum_{k=0}^{\infty}\frac{t^k}{x^{k+1}}$, and the sum converges uniformly because $|t/x| \le a/|x| < 1$. Now just multiply by $w(t)$ and integrate. $\square$

By way of an example, consider the Chebyshev weight $w(x) = (1-x^2)^{-1/2}$ on $[-1,1]$. For $x > 1$ we have

$$\int_{-1}^{1} \frac{dt}{(x-t)\sqrt{1-t^2}} \;=\; \frac{\pi}{\sqrt{x^2-1}} \qquad \left(\text{set } t = 2u/(1+u^2)\right)$$

$$= \frac{\pi}{x}\left(1-\frac{1}{x^2}\right)^{-1/2}$$

$$= \frac{\pi}{x}\left[1 + \frac{1}{2}\cdot\frac{1}{x^2} + \frac{1\cdot 3}{2\cdot 2}\cdot\frac{1}{2!}\cdot\frac{1}{x^4} + \cdots\right],$$

using the binomial formula. Thus, we've found all the moments:

$$\begin{aligned}
\mu_0 &= \int_{-1}^{1} \frac{dt}{\sqrt{1-t^2}} = \pi \\
\mu_{2n-1} &= \int_{-1}^{1} \frac{t^{2n-1}dt}{\sqrt{1-t^2}} = 0 \\
\mu_{2n} &= \int_{-1}^{1} \frac{t^{2n}dt}{\sqrt{1-t^2}} = \frac{1\cdot 3\cdot 5\cdots(2n-1)}{2^n n!}\,\pi.
\end{aligned}$$

Stieltjes proved much more: The integral $\int_{-a}^{a} \frac{w(t)}{x-t}\,dt$ is actually an *analytic* function of x in $\mathbb{C}\setminus[-a,a]$. In any case, since $x \notin [-a,a]$, we know that $\frac{1}{x-t}$ is continuous on $[-a,a]$. In particular, we can apply our quadrature formulas (and Stieltjes theorem, p. 132) to write

$$\int_{-a}^{a} \frac{w(t)}{x-t}\,dt = \lim_{n\to\infty} \sum_{i=1}^{n} \frac{A_i^{(n)}}{x - x_i^{(n)}},$$

and these sums are recognizable:

Lemma. $\displaystyle\sum_{i=1}^{n} \frac{A_i^{(n)}}{x - x_i^{(n)}} = \frac{\varphi_n(x)}{Q_n(x)}.$

Proof. Since φ_n has degree $< n$ and $\varphi_n(x_i^{(n)}) \neq 0$ for any i, we may appeal to partial-fractions to write

$$\frac{\varphi_n(x)}{Q_n(x)} = \frac{\varphi_n(x)}{(x - x_1^{(n)})\cdots(x - x_n^{(n)})} = \sum_{i=1}^{n} \frac{c_i}{x - x_i^{(n)}}$$

where c_i is given by

$$c_i = \left.\frac{\varphi_n(x)}{Q_n(x)}\,(x - x_i^{(n)})\right]_{x = x_i^{(n)}} = \frac{\varphi_n(x_i^{(n)})}{Q_n'(x_i^{(n)})} = A_i^{(n)}. \qquad \square$$

Now here's where the continued fractions come in: Stieltjes recognized the fact that

$$\frac{\varphi_{n+1}(x)}{Q_{n+1}(x)} = \cfrac{b_0}{(x-a_0) - \cfrac{b_1}{(x-a_1) - \ddots - \cfrac{b_n}{(x-a_n)}}}$$

(which can be proved by induction), where $b_0 = \int_a^b w(t)\,dt$. More generally, induction will show that the n-th *convergent* of a continued fraction can be written as

$$\frac{A_n}{B_n} = \cfrac{p_1}{q_1 - \cfrac{p_2}{q_2 - \ddots - \cfrac{p_n}{q_n}}}$$

by means of the *recurrence formulas*

$$\begin{aligned} A_0 &= 0 & B_0 &= 1 \\ A_1 &= p_1 & B_1 &= q_1 \\ A_n &= q_n A_{n-1} + p_n A_{n-2} & B_n &= q_n B_{n-1} + p_n B_{n-2} \end{aligned}$$

where $n = 2, 3, 4, \ldots$. Please note that A_n and B_n satisfy the same recurrence formula, but with different starting values (as is the case with φ_n and Q_n).

Again using the Chebyshev weight as an example, for $x > 1$ we have

$$\frac{\pi}{\sqrt{x^2 - 1}} = \int_{-1}^{1} \frac{dt}{(x - t)\sqrt{1 - t^2}} = \cfrac{\pi}{x - \cfrac{1/2}{x - \cfrac{1/4}{x - \cfrac{1/4}{\ddots}}}}$$

since $a_n = 0$ for all n, $b_1 = 1/2$, and $b_n = 1/4$ for $n \geq 2$. In other words, we've just found a continued fraction expansion for $(x^2 - 1)^{-1/2}$.

Appendix

Finally, here is a brief review of some of the fancier bits of linear algebra used in this chapter. To begin, we discuss sums and quotients of vector spaces.

Each subspace M of a finite-dimensional X induces an equivalence relation on X by

$$x \sim y \iff x - y \in M.$$

Standard arguments show that the equivalence classes under this relation are the cosets (translates) $x + M$, $x \in X$. That is,

$$x + M = y + M \iff x - y \in M \iff x \sim y.$$

Equally standard is the induced vector arithmetic

$$(x+M)+(y+M)=(x+y)+M \qquad \text{and} \qquad \alpha(x+M)=(\alpha x)+M,$$

where x, $y \in X$ and $\alpha \in \mathbb{R}$. The collection of cosets (or equivalence classes) is a vector space under these operations; it's denoted X/M and called the *quotient* of X by M. Please note the the zero vector in X/M is simply M itself.

Associated to the quotient space X/M is the *quotient map* $q(x)=x+M$. It's easy to check that $q : X \to X/M$ is a vector space homomorphism with kernel M. (Why?)

Next we recall *the isomorphism theorem.*

Theorem. *Let $T : X \to Y$ be a linear map between finite-dimensional vector spaces, and let $q : X \to X/\ker T$ be the quotient map. Then, there exists a (unique, into) isomorphism $S : X/\ker T \to Y$ satisfying $S(q(x)) = T(x)$ for every $x \in X$.*

Proof. Since q maps onto $X/\ker T$, it's "legal" to define a map $S : X/\ker T \to Y$ by setting $S(q(x)) = T(x)$ for $x \in X$. Please note that S is well-defined since

$$\begin{aligned} T(x) = T(y) \iff T(x-y) = 0 \iff x - y \in \ker T \\ \iff q(x-y) = 0 \iff q(x) = q(y). \end{aligned}$$

It's easy to see that S is linear and so precisely the same argument as above shows that S is one-to-one. $\square$

Corollary. *Let $T : X \to Y$ be a linear map between finite-dimensional vector spaces. Then, the range of T is isomorphic to $X/\ker T$.*

The Müntz Theorems

For several weeks now we've taken advantage of the fact that the monomials $1, x, x^2, \ldots$ have *dense linear span* in $C[\,0,1\,]$. What, if anything, is so special about these particular powers? How about if we consider polynomials of the form $\sum_{k=0}^{n} a_k x^{k^2}$; are they dense, too? More generally, what can be said about the span of a sequence of monomials (x^{λ_n}), where $\lambda_0 < \lambda_1 < \lambda_2 < \cdots$? Of course, we'll have to assume that $\lambda_0 \geq 0$, but it's not hard to see that we will actually need $\lambda_0 = 0$, for otherwise each of the polynomials $\sum_{k=0}^{n} a_k x^{\lambda_k}$ vanishes at $x = 0$ (and so has distance at least 1 from the constant 1 function, for example). If the λ_n's are integers, it's also clear that we'll have to have $\lambda_n \to \infty$ as $n \to \infty$. But what else is needed? The answer comes to us from Müntz in 1914. (You sometimes see the name Otto Szász associated with Müntz's theorem, because Szász proved a similar theorem at nearly the same time (1916).)

Theorem. *Let $0 \leq \lambda_0 < \lambda_1 < \lambda_2 < \cdots$. Then, the functions (x^{λ_n}) have dense linear span in $C[\,0,1\,]$ if and only if $\lambda_0 = 0$ and $\sum_{n=1}^{\infty} \lambda_n^{-1} = \infty$.*

What Müntz is trying to tell us here is that the λ_n's can't get big too quickly. In particular, the polynomials of the form $\sum_{k=0}^{n} a_k x^{k^2}$ are evidently *not* dense in $C[\,0,1\,]$. On the other hand, the λ_n's don't have to be unbounded; indeed, Müntz's theorem implies an earlier result of Bernstein from 1912: If $0 < \alpha_1 < \alpha_2 < \cdots < K$ (some constant), then $1, x^{\alpha_1}, x^{\alpha_2}, \ldots$ have dense linear span in $C[\,0,1\,]$.

Before we give the proof of Müntz's theorem, let's invent a bit of notation: We write

$$X_n \;=\; \left\{ \sum_{k=0}^{n} a_k x^{\lambda_k} \,:\, a_0, \ldots, a_n \in \mathbb{R} \right\}$$

and, given $f \in C[\,0,1\,]$, we write $\operatorname{dist}(f, X_n)$ to denote the distance from f to the space spanned by $1, x^{\lambda_1}, \ldots, x^{\lambda_n}$. Let's also write $X = \bigcup_{n=0}^{\infty} X_n$. That is, X is the linear span of the entire sequence $(x^{\lambda_n})_{n=0}^{\infty}$. The question here is whether X is dense, and we'll address the problem by determining whether $\operatorname{dist}(f, X_n) \to 0$, as $n \to \infty$, for every $f \in C[\,0,1\,]$.

If we can show that each (fixed) power x^m can be uniformly approximated by a linear combination of x^{λ_n}'s, then the Weierstrass theorem will tell us that X is dense in $C[\,0,1\,]$.

(How?) Surprisingly, the numbers $\operatorname{dist}(x^m, X_n)$ can be estimated. Our proof won't give the best estimate, but it will show how the condition $\sum_{n=1}^\infty \lambda_n^{-1} = \infty$ comes into the picture.

Lemma. *Let $m > 0$. Then,* $\operatorname{dist}(x^m, X_n) \;\le\; \prod_{k=1}^n \left| 1 - \frac{m}{\lambda_k} \right|$.

Proof. We may certainly assume that $m \ne \lambda_n$ for any n. Given this, we inductively define a sequence of functions by setting $P_0(x) = x^m$ and

$$P_n(x) \;=\; (\lambda_n - m)\, x^{\lambda_n} \int_x^1 t^{-1-\lambda_n}\, P_{n-1}(t)\, dt$$

for $n \ge 1$. For example,

$$P_1(x) \;=\; (\lambda_1 - m)\, x^{\lambda_1} \int_x^1 t^{-1-\lambda_1}\, t^m\, dt \;=\; -x^{\lambda_1}\, t^{m-\lambda_1} \Big]_x^1 \;=\; x^m - x^{\lambda_1}.$$

By induction, each P_n is of the form $x^m - \sum_{k=0}^n a_k x^{\lambda_k}$ for some scalars (a_k):

$$\begin{aligned} P_n(x) \;&=\; (\lambda_n - m)\, x^{\lambda_n} \int_x^1 t^{-1-\lambda_n}\, P_{n-1}(t)\, dt \\ &=\; (\lambda_n - m)\, x^{\lambda_n} \int_x^1 t^{-1-\lambda_n} \left[t^m - \sum_{k=0}^{n-1} a_k t^{\lambda_k} \right] dt \\ &=\; x^m - x^{\lambda_n} + (\lambda_n - m) \sum_{k=0}^{n-1} \frac{a_k}{\lambda_n - \lambda_k} \left(x^{\lambda_k} - x^{\lambda_n} \right). \end{aligned}$$

Finally, $\|P_0\| = 1$ and $\|P_n\| \le |1 - \frac{m}{\lambda_n}|\, \|P_{n-1}\|$, because

$$|\lambda_n - m|\, x^{\lambda_n} \int_x^1 t^{-1-\lambda_n}\, dt \;=\; \frac{|\lambda_n - m|}{\lambda_n} \left(1 - x^{\lambda_n} \right) \;\le\; \left| 1 - \frac{m}{\lambda_n} \right|.$$

Thus,

$$\operatorname{dist}(x^m, X_n) \;\le\; \|P_n\| \;\le\; \prod_{k=1}^n \left| 1 - \frac{m}{\lambda_k} \right|. \qquad \square$$

The preceding result is due to v. Golitschek. A slightly better estimate, also due to v. Golitschek (1970), is $\operatorname{dist}(x^m, X_n) \le \prod_{k=1}^n \frac{|m - \lambda_k|}{m + \lambda_k}$.

Now a well-known fact about infinite products is that, for positive a_k's, the product $\prod_{k=1}^\infty \left| 1 - a_k \right|$ diverges (to 0) if and only if the series $\sum_{k=1}^\infty a_k$ diverges (to ∞) if and only if the product $\prod_{k=1}^{\infty}\left|1+a_k\right|$ diverges (to ∞). In particular, $\prod_{k=1}^{n}\left|1-\frac{m}{\lambda_k}\right| \to 0$ if and only if $\sum_{k=1}^{n}\frac{1}{\lambda_k} \to \infty$. That is, $\operatorname{dist}(x^m, X_n) \to 0$ if and only if $\sum_{k=1}^{\infty}\frac{1}{\lambda_k} = \infty$. This proves the "backward" direction of Müntz's theorem.

We'll prove the "forward" direction of Müntz's theorem by proving a version of Müntz's theorem for the space $L_2[\,0,1\,]$. For our purposes, $L_2[\,0,1\,]$ denotes the space $C[\,0,1\,]$ endowed with the norm

$$\|f\|_2 \;=\; \left(\int_0^1 |f(x)|^2\,dx\right)^{1/2},$$

although our results are equally valid in the "real" space $L_2[\,0,1\,]$ (consisting of square-integrable, Lebegue measurable functions). In the latter case, we no longer need to assume that $\lambda_0 = 0$, but we do need to assume that each $\lambda_n > -1/2$ (in order that $x^{2\lambda_n}$ be integrable on $[\,0,1\,]$).

Remarkably, the distance from f to the span of $x^{\lambda_0}, x^{\lambda_1}, \ldots, x^{\lambda_n}$ can be computed *exactly* in the L_2 norm. For this we'll need some more notation: Given linearly independent vectors $f_1, \ldots, f_n$ in an inner product space, we call

$$G(f_1,\ldots,f_n) \;=\; \begin{vmatrix} \langle\, f_1, f_1\,\rangle & \cdots & \langle\, f_1, f_n\,\rangle \\ \vdots & \ddots & \vdots \\ \langle\, f_n, f_1\,\rangle & \cdots & \langle\, f_n, f_n\,\rangle \end{vmatrix} \;=\; \det\big[\,\langle\, f_i, f_j\,\rangle\,\big]_{i,j}$$

the *Gram determinant* of the f_k's.

Lemma. (Gram) *Let F be a finite dimensional subspace of an inner product space V, and let $g \in V \setminus F$. Then, the distance d from g to F is given by*

$$d^{\,2} \;=\; \frac{G(g, f_1,\ldots,f_n)}{G(f_1,\ldots,f_n)},$$

where $f_1,\ldots,f_n$ is any basis for F.

Proof. Let $f = \sum_{i=1}^{n} a_i f_i$ be the best approximation to g out of F. Then, since $g - f$ is orthogonal to F, we have, in particular, $\langle\, f_j, g\,\rangle = \langle\, f_j, g\,\rangle$ for all j; that is,

$$\sum_{i=1}^{n} a_i\langle\, f_j, f_i\,\rangle \;=\; \langle\, f_j, g\,\rangle, \quad j = 1,\ldots,n. \tag{$*$}$$

Since this system of equations always has a unique solution $a_1,\ldots,a_n$, we must have $G(f_1,\ldots,f_n) \neq 0$ (and so the formula in our Lemma at least makes sense).

Next, notice that

$$d^2 = \langle g-f, g-f \rangle = \langle g-f, g \rangle = \langle g, g \rangle - \langle g, f \rangle;$$

in other words,

$$d^2 + \sum_{i=1}^{n} a_i \langle g, f_i \rangle = \langle g, g \rangle. \tag{$**$}$$

Now consider $(*)$ and $(**)$ as a system of $n+1$ equations in the $n+1$ unknowns $a_1, \ldots, a_n$, and d^2; in matrix form we have

$$\begin{bmatrix} 1 & \langle g, f_1 \rangle & \cdots & \langle g, f_n \rangle \\ 0 & \langle f_1, f_1 \rangle & \cdots & \langle f_1, f_n \rangle \\ \vdots & \ddots & \vdots & \vdots \\ 0 & \langle f_n, f_1 \rangle & \cdots & \langle f_n, f_n \rangle \end{bmatrix} \begin{bmatrix} d^2 \\ a_1 \\ \vdots \\ a_n \end{bmatrix} = \begin{bmatrix} \langle g, g \rangle \\ \langle f_1, g \rangle \\ \vdots \\ \langle f_n, g \rangle \end{bmatrix}.$$

Solving for d^2 using Cramer's rule gives the desired result; expanding along the first column shows that the matrix of coefficients has determinant $G(f_1, \ldots, f_n)$, while the matrix obtained by replacing the "d column" by the right-hand side has determinant $G(g, f_1, \ldots, f_n)$. □

Note: By our last Lemma and induction, every Gram determinant is positive!

In what follows, we will still use X_n to denote the span of $x^{\lambda_0}, \ldots, x^{\lambda_n}$, but now we'll write $\operatorname{dist}_2(f, X_n)$ to denote the distance from f to X_n in the L_2 norm.

Theorem. *Let m, $\lambda_k > -1/2$ for $k = 0, 1, 2, \ldots$. Then,*

$$\operatorname{dist}_2(x^m, X_n) = \frac{1}{\sqrt{2m+1}} \prod_{k=0}^{n} \frac{|m - \lambda_k|}{m + \lambda_k + 1}.$$

PROOF. The proof is based on a determinant formula due to Cauchy:

$$\prod_{i,j} (a_i + b_j) \begin{vmatrix} \frac{1}{a_1+b_1} & \cdots & \frac{1}{a_1+b_n} \\ \vdots & \ddots & \vdots \\ \frac{1}{a_n+b_1} & \cdots & \frac{1}{a_n+b_n} \end{vmatrix} = \prod_{i>j} (a_i - a_j)(b_i - b_j).$$

If we consider each of the a_i's and b_j's as "variables," then each side of the equation is a *polynomial* in $a_1, \ldots, a_n, b_1, \ldots, b_n$. (Why?) Now the right-hand side clearly vanishes if

$a_i = a_j$ or $b_i = b_j$ for some $i \neq j$, but the left-hand side also vanishes in any of these cases. Thus, the right-hand side divides the left-hand side. But both polynomials have degree $n-1$ in each of the a_i's and b_j's. (Why?) Thus, the left-hand side is a constant multiple of the right-hand side. To show that the constant must be 1, write the left-hand side as

$$\prod_{i \neq j}(a_i + b_j) \begin{vmatrix} 1 & \frac{a_1+b_1}{a_1+b_2} & \cdots & \frac{a_1+b_1}{a_1+b_n} \\ \frac{a_2+b_2}{a_2+b_1} & 1 & \cdots & \frac{a_2+b_2}{a_2+b_n} \\ \vdots & & \ddots & \vdots \\ \frac{a_n+b_n}{a_n+b_1} & \cdots & \frac{a_n+b_n}{a_n+b_{n-1}} & 1 \end{vmatrix}$$

and now take the limit as $b_1 \to -a_1$, $b_2 \to -a_2$, etc. The expression above tends to $\prod_{i \neq j}(a_i - a_j)$, as does the right-hand side of Cauchy's formula.

Now, $\langle\, x^p, x^q \,\rangle = \int_0^1 x^{p+q}\, dx = \frac{1}{p+q+1}$ for p, $q > -1/2$, so

$$G(x^{\lambda_0}, \ldots, x^{\lambda_n}) \;=\; \det\left(\left[\frac{1}{\lambda_i + \lambda_j + 1}\right]_{i,j}\right) \;=\; \frac{\prod_{i>j}(\lambda_i - \lambda_j)^2}{\prod_{i,j}(\lambda_i + \lambda_j + 1)},$$

with a similar formula holding for $G(x^m, x^{\lambda_0}, \ldots, x^{\lambda_n})$. Substituting these expressions into our distance formula and taking square roots finishes the proof. $\square$

Now we can determine exactly when X is dense in $L_2[\,0,1\,]$. For easier comparison to the $C[\,0,1\,]$ case, we suppose that the λ_n's are nonnegative.

Theorem. *Let $0 \le \lambda_0 < \lambda_1 < \lambda_2 < \cdots$. Then, the functions (x^{λ_n}) have dense linear span in $L_2[\,0,1\,]$ if and only if $\sum_{n=1}^{\infty} \lambda_n^{-1} = \infty$.*

PROOF. If $\sum_{n=1}^{\infty} \frac{1}{\lambda_n} < \infty$, then each of the products $\prod_{k=1}^{n} \left|1 - \frac{m}{\lambda_k}\right|$ and $\prod_{k=1}^{n} \left|1 + \frac{(m+1)}{\lambda_k}\right|$ converges to some nonzero limit for any m not equal to any λ_k. Thus, $\operatorname{dist}_2(x^m, X_n) \not\to 0$, as $n \to \infty$, for any $m \neq \lambda_k$, $k = 0, 1, 2, \ldots$. In particular, the functions (x^{λ_n}) cannot have dense linear span in $L_2[\,0,1\,]$.

Conversely, if $\sum_{n=1}^{\infty} \frac{1}{\lambda_n} = \infty$, then $\prod_{k=1}^{n} \left|1 - \frac{m}{\lambda_k}\right|$ diverges to 0 while $\prod_{k=1}^{n} \left|1 + \frac{(m+1)}{\lambda_k}\right|$ diverges to $+\infty$. Thus, $\operatorname{dist}_2(x^m, X_n) \to 0$, as $n \to \infty$, for every $m > -1/2$. Since the polynomials are dense in $L_2[\,0,1\,]$, this finishes the proof. $\square$

Finally, we can finish the proof of Müntz's theorem in the case of $C[\,0,1\,]$. Suppose that the functions (x^{λ_n}) have dense linear span in $C[\,0,1\,]$. Then, since $\|f\|_2 \le \|f\|$, it follows that the functions (x^{λ_n}) must also have dense linear span in $L_2[\,0,1\,]$. (Why?) Hence, $\sum_{n=1}^{\infty} \frac{1}{\lambda_n} = \infty$. $\square$

Just for good measure, here's a second proof of the "backward" direction for $C[\,0,1\,]$ based on the $L_2[\,0,1\,]$ version. Suppose that $\sum_{n=1}^{\infty} \frac{1}{\lambda_n} = \infty$, and let $m \geq 1$. Then,

$$\begin{aligned}
\left| x^m - \sum_{k=0}^{n} a_k x^{\lambda_k} \right| &= \left| \frac{1}{m} \int_0^x t^{m-1}\,dt - \sum_{k=0}^{n} \frac{a_k}{\lambda_k} \int_0^x t^{\lambda_k - 1}\,dt \right| \\
&\leq \int_0^1 \left| \frac{1}{m}\, t^{m-1} - \sum_{k=0}^{n} \frac{a_k}{\lambda_k}\, t^{\lambda_k - 1} \right| dt \\
&\leq \left(\int_0^1 \left| \frac{1}{m}\, t^{m-1} - \sum_{k=0}^{n} \frac{a_k}{\lambda_k}\, t^{\lambda_k - 1} \right|^2 dt \right)^{1/2}.
\end{aligned}$$

Now the functions $(x^{\lambda_k - 1})$ have dense linear span in $L_2[\,0,1\,]$ because $\sum_{\lambda_n > 1} \frac{1}{\lambda_n - 1} = \infty$. Thus, we can find a_k's so that the right-hand side of this inequality is less than some ε. Since this estimate is independent of x, we've shown that

$$\max_{0 \leq x \leq 1} \left| x^m - \sum_{k=0}^{n} a_k x^{\lambda_k} \right| < \varepsilon. \qquad \square$$

Application. *Let $0 = \lambda_0 < \lambda_1 < \lambda_2 < \cdots$ with $\sum_{n=1}^{\infty} \lambda_n^{-1} = \infty$, and let f be a continuous function on $[\,0,\infty)$ for which $c = \lim_{t\to\infty} f(t)$ exists. Then, f can be uniformly approximated by finite linear combinations of the exponentials $(e^{-\lambda_n t})_{n=0}^{\infty}$.*

Proof. The function $g(x) = f(-\log x)$, for $0 < x \leq 1$, and $g(0) = c$, is continuous on $[\,0,1\,]$. In other words, $g(e^{-t}) = f(t)$ for each $0 \leq t < \infty$. Thus, given $\varepsilon > 0$, we can find n and $a_0, \ldots, a_n$ such that

$$\max_{0 \leq x \leq 1} \left| g(x) - \sum_{k=0}^{n} a_k x^{\lambda_k} \right| = \max_{0 \leq t < \infty} \left| f(t) - \sum_{k=0}^{n} a_k e^{-\lambda_k t} \right| < \varepsilon. \qquad \square$$

The Stone-Weierstrass Theorem

To begin, an *algebra* is a vector space A on which there is a multiplication $(f, g) \mapsto fg$ (from $A \times A$ into A) satisfying

(i) $(fg)h = f(gh)$, for all f, g, $h \in A$;

(ii) $f(g+h) = fg + fh$ and $(f+g)h = fg + gh$, for all f, g, $h \in A$;

(iii) $\alpha(fg) = (\alpha f)g = f(\alpha g)$, for all scalars α and all f, $g \in A$.

In other words, an algebra is a *ring* under vector addition and multiplication, together with a compatible scalar multiplication. The algebra is *commutative* if

(iv) $fg = gf$, for all f, $g \in A$.

And we say that A has an *identity* element if there is a vector $e \in A$ such that

(v) $fe = ef = f$, for all $f \in A$.

In case A is a normed vector space, we also require that the norm satisfy

(vi) $\|fg\| \le \|f\|\,\|g\|$

(this simplifies things a bit), and in this case we refer to A as a *normed algebra*. If a normed algebra is complete, we refer to it as a *Banach algebra*. Finally, a subset B of an algebra A is called a *subalgebra* (of A) if B is itself an algebra (under the same operations); that is, if B is a (vector) subspace of A which is closed under multiplication.

If A is a normed algebra, then all of the various operations on A (or $A \times A$) are continuous. For example, since

$$\|fg - hk\| \;=\; \|fg - fk + fk - hk\| \;\le\; \|f\|\,\|g-k\| + \|k\|\,\|f-h\|$$

it follows that multiplication is continuous. (How?) In particular, if B is a subspace (or subalgebra) of A, then $\overline{B}$, the closure of B, is also a subspace (or subalgebra) of A.

Examples

1. If we define multiplication of vectors "coordinatewise," then $\mathbb{R}^n$ is a commutative Banach algebra with identity (the vector $(1, \ldots, 1)$) when equipped with the norm $\|x\|_\infty = \max\limits_{1 \le i \le n} |x_i|$.

2. It's not hard to identify the subalgebras of $\mathbb{R}^n$ among its subspaces. For example, the subalgebras of $\mathbb{R}^2$ are $\{(x,0) : x \in \mathbb{R}\}$, $\{(0,y) : y \in \mathbb{R}\}$, and $\{(x,x) : x \in \mathbb{R}\}$, along with $\{(0,0)\}$ and $\mathbb{R}^2$.

3. Given a set X, we write $B(X)$ for the space of all bounded, real-valued functions on X. If we endow $B(X)$ with the sup norm, and if we define arithmetic with functions pointwise, then $B(X)$ is a commutative Banach algebra with identity (the constant 1 function). The constant functions in $B(X)$ form a subalgebra isomorphic (in every sense of the word) to $\mathbb{R}$.

4. If X is a metric (or topological) space, then we may consider $C(X)$, the space of all continuous, real-valued functions on X. If we again define arithmetic with functions pointwise, then $C(X)$ is a commutative algebra with identity (the constant 1 function). The bounded, continuous functions on X, written $C_b(X) = C(X) \cap B(X)$, form a closed subalgebra of $B(X)$. If X is compact, then $C_b(X) = C(X)$. In other words, if X is compact, then $C(X)$ is itself a closed subalgebra of $B(X)$ and, in particular, $C(X)$ is a Banach algebra with identity.

5. The polynomials form a dense subalgebra of $C[\,a,b\,]$. The trig polynomials form a dense subalgebra of $C^{2\pi}$. These two sentences summarize Weierstrass's two classical theorems in modern parlance and form the basis for Stone's version of the theorem.

Using this new language, we may restate the classical Weierstrass theorem to read: *If a subalgebra A of $C[\,a,b\,]$ contains the functions $e(x) = 1$ and $f(x) = x$, then A is dense in $C[\,a,b\,]$.* Any subalgebra of $C[\,a,b\,]$ containing 1 and x actually contains all the polynomials; thus, our restatement of Weierstrass's theorem amounts to the observation that any subalgebra containing a dense set is itself dense in $C[\,a,b\,]$.

Our goal in this section is to prove an analogue of this new version of the Weierstrass theorem for subalgebras of $C(X)$, where X is a compact metric space. In particular, we will want to extract the essence of the functions 1 and x from this statement. That is, we seek conditions on a subalgebra A of $C(X)$ that will force A to be dense in $C(X)$. The key role played by 1 and x, in the case of $C[\,a,b\,]$, is that a subalgebra containing these two functions must actually contain a much larger set of functions. But since we can't be assured of anything remotely like polynomials living in the more general $C(X)$ spaces, we might want to change our point of view. What we really need is some requirement on a subalgebra A of $C(X)$ that will allow us to *construct* a wide variety of functions in A. And, if A contains a sufficiently rich variety of functions, it might just be possible to show that A is dense.

Since the two replacement conditions we have in mind make sense in any collection of real-valued functions, we state them in some generality.

Let A be a collection of real-valued functions on some set X. We say that A *separates points in* X if, given $x \neq y \in X$, there is some $f \in A$ such that $f(x) \neq f(y)$. We say that A *vanishes at no point of* X if, given $x \in X$, there is some $f \in A$ such that $f(x) \neq 0$.

Examples

6. The single function $f(x) = x$ clearly separates points in $[\,a, b\,]$, and the function $e(x) = 1$ obviously vanishes at no point in $[\,a, b\,]$. Any subalgebra A of $C[\,a, b\,]$ containing these two functions will likewise separate points and vanish at no point in $[\,a, b\,]$.

7. The set E of even functions in $C[-1, 1\,]$ fails to separate points in $[-1, 1\,]$; indeed, $f(x) = f(-x)$ for any even function. However, since the constant functions are even, E vanishes at no point of $[-1, 1\,]$. It's not hard to see that E is a proper closed subalgebra of $C[-1, 1\,]$. The set of odd functions will separate points (since $f(x) = x$ is odd), but the odd functions all vanish at 0. The set of odd functions is a proper closed subspace of $C[-1, 1\,]$, although not a subalgebra.

8. The set of all functions $f \in C[-1, 1\,]$ for which $f(0) = 0$ is a proper closed subalgebra of $C[-1, 1\,]$. In fact, this set is a maximal (in the sense of containment) proper closed subalgebra of $C[-1, 1\,]$. Note, however, that this set of functions does separate points in $[-1, 1\,]$ (again, because it contains $f(x) = x$).

9. It's easy to construct examples of non-trivial closed subalgebras of $C(X)$. Indeed, given any closed subset X_0 of X, the set $A(X_0) = \{f \in C(X) : f \text{ vanishes on } X_0\}$ is a non-empty, proper subalgebra of $C(X)$. It's closed in any reasonable topology on $C(X)$ because it's closed under pointwise limits. Subalgebras of the type $A(X_0)$ are of interest because they're actually *ideals* in the ring $C(X)$. That is, if $f \in C(X)$, and if $g \in A(X_0)$, then $fg \in A(X_0)$.

As these few examples illustrate, neither of our new conditions, taken separately, is enough to force a subalgebra of $C(X)$ to be dense. But both conditions together turn out to be sufficient. In order to better appreciate the utility of these new conditions, let's isolate the key computational tool that they permit within an algebra of functions.

Lemma. *Let A be an algebra of real-valued functions on some set X, and suppose that A separates points in X and vanishes at no point of X. Then, given $x \neq y \in X$ and a, $b \in \mathbb{R}$, we can find an $f \in A$ with $f(x) = a$ and $f(y) = b$.*

PROOF. Given any pair of distinct points $x \neq y \in X$, the set $\widetilde{A} = \{(f(x), f(y)) : f \in A\}$ is a subalgebra of $\mathbb{R}^2$. If A separates points in X, then $\widetilde{A}$ is evidently neither $\{(0,0)\}$ nor $\{(x,x) : x \in \mathbb{R}\}$. If A vanishes at no point, then $\{(x,0) : x \in \mathbb{R}\}$ and $\{(0,y) : y \in \mathbb{R}\}$ are both excluded. Thus $\widetilde{A} = \mathbb{R}^2$. That is, for any a, $b \in \mathbb{R}$, there is some $f \in A$ for which $(f(x), f(y)) = (a, b)$. $\square$

Now we can state Stone's version of the Weierstrass theorem (for compact metric spaces). It should be pointed out that the theorem, as stated, also holds in $C(X)$ when X is a compact Hausdorff topological space (with the same proof), but does not hold for algebras of complex-valued functions over $\mathbb{C}$. More on this later.

Stone-Weierstrass Theorem. (real scalars) *Let X be a compact metric space, and let A be a subalgebra of $C(X)$. If A separates points in X and vanishes at no point of X, then A is dense in $C(X)$.*

What Cheney calls an "embryonic" version of this theorem appeared in 1937, as a small part of a massive 106 page paper! Later versions, appearing in 1948 and 1962, benefitted from the work of the great Japanese mathematician Kakutani and were somewhat more palatable to the general mathematical public. But, no matter which version you consult, you'll find them difficult to read. For more details, I would recommend you first consult Folland's *Real Analysis*, or Simmons's *Topology and Modern Analysis*.

As a first step in attacking the proof of Stone's theorem, notice that if A satisfies the conditions of the theorem, then so does its closure $\overline{A}$. (Why?) Thus, we may assume that A is actually a *closed* subalgebra of $C(X)$ and prove, instead, that $A = C(X)$. Now the closed subalgebras of $C(X)$ inherit more structure than you might first imagine.

Theorem. *If A is a subalgebra of $C(X)$, and if $f \in A$, then $|f| \in \overline{A}$. Consequently, $\overline{A}$ is a sublattice of $C(X)$.*

PROOF. Let $\varepsilon > 0$, and consider the function $|t|$ on the interval $\big[-\|f\|, \|f\|\,\big]$. By the Weierstrass theorem, there is a polynomial $p(t) = \sum_{k=0}^{n} a_k t^k$ such that $\big|\,|t| - p(t)\,\big| < \varepsilon$ for all $|t| \le \|f\|$. In particular, notice that $|p(0)| = |a_0| < \varepsilon$.

Now, since $|f(x)| \le \|f\|$ for all $x \in X$, it follows that $\big|\,|f(x)| - p(f(x))\,\big| < \varepsilon$ for all $x \in X$. But $p(f(x)) = (p(f))(x)$, where $p(f) = a_0\mathbf{1} + a_1 f + \cdots + a_n f^n$, and the function $g = a_1 f + \cdots + a_n f^n \in A$, since A is an algebra. Thus, $\big|\,|f(x)| - g(x)\,\big| \le |a_0| + \varepsilon < 2\varepsilon$ for all $x \in X$. In other words, for each $\varepsilon > 0$, we can supply an element $g \in A$ such that $\|\,|f| - g\| < 2\varepsilon$. That is, $|f| \in \overline{A}$.

The statement that $\overline{A}$ is a sublattice of $C(X)$ means that if we're given $f,\, g \in \overline{A}$, then $\max\{f, g\} \in \overline{A}$ and $\min\{f, g\} \in \overline{A}$, too. But this is actually just a statement about real numbers. Indeed, since

$$2\max\{a, b\} = a + b + |a - b| \qquad \text{and} \qquad 2\min\{a, b\} = a + b - |a - b|$$

it follows that a *subspace* of $C(X)$ is a sublattice precisely when it contains the absolute values of all its elements. $\square$

The point to our last result is that if we're given a closed subalgebra A of $C(X)$, then A is "closed" in every sense of the word: Sums, products, absolute values, max's, and min's of elements from A, and even limits of sequences of these, are all back in A. This is precisely the sort of freedom we'll need if we hope to show that $A = C(X)$.

Please notice that we could have avoided our appeal to the Weierstrass theorem in this last result. Indeed, we really only need to supply polynomial approximations for the single function $|x|$ on $[-1, 1\,]$, and this can be done directly. For example, we could appeal instead to the binomial theorem, using $|x| = \sqrt{1 - (1 - x^2)}$. The resulting series can be shown to converge uniformly on $[-1, 1\,]$. By side-stepping the classical Weierstrass theorem, it becomes a corollary to Stone's version (rather than the other way around).

Now we're ready for the proof of the Stone-Weierstrass theorem. As we've already pointed out, we may assume that we're given a closed subalgebra (subspace, and sublattice)

A of $C(X)$ and we want to show that $A = C(X)$. We'll break the remainder of the proof into two steps:

<u>Step 1</u>: Given $f \in C(X)$, $x \in X$, and $\varepsilon > 0$, there is an element $g_x \in A$ with $g_x(x) = f(x)$ and $g_x(y) > f(y) - \varepsilon$ for all $y \in X$.

From our "computational" Lemma, we know that for each $y \in X$, $y \neq x$, we can find an $h_y \in A$ so that $h_y(x) = f(x)$ and $h_y(y) = f(y)$. Since $h_y - f$ is continuous and vanishes at both x and y, the set $U_y = \{t \in X : h_y(t) > f(t) - \varepsilon\}$ is open and contains both x and y. Thus, the sets $(U_y)_{y \neq x}$ form an open cover for X. Since X is compact, finitely many U_y's suffice, say $X = U_{y_1} \cup \cdots \cup U_{y_n}$. Now set $g_x = \max\{h_{y_1}, \ldots, h_{y_n}\}$. Because A is a lattice, we have $g_x \in A$. Note that $g_x(x) = f(x)$ since each h_{y_i} agrees with f at x. And $g_x > f - \varepsilon$ since, given $y \neq x$, we have $y \in U_{y_i}$ for some i, and hence $g_x(y) \geq h_{y_i}(y) > f(y) - \varepsilon$.

<u>Step 2</u>: Given $f \in C(X)$ and $\varepsilon > 0$, there is an $h \in A$ with $\|f - h\| < \varepsilon$.

From Step 1, for each $x \in X$ we can find some $g_x \in A$ such that $g_x(x) = f(x)$ and $g_x(y) > f(y) - \varepsilon$ for all $y \in X$. And now we reverse the process used in Step 1: For each x, the set $V_x = \{y \in X : g_x(y) < f(y) + \varepsilon\}$ is open and contains x. Again, since X is compact, $X = V_{x_1} \cup \cdots V_{x_m}$ for some $x_1, \ldots, x_m$. This time, set $h = \min\{g_{x_1}, \ldots, g_{x_m}\} \in A$. As before, $h(y) > f(y) - \varepsilon$ for all y, since each g_{x_i} does so, and $h(y) < f(y) + \varepsilon$ for all y, since at least one g_{x_i} does so.

The conclusion of Step 2 is that A is dense in $C(X)$; but, since A is closed, this means that $A = C(X)$. $\square$

Corollary. *If X and Y are compact metric spaces, then the subspace of $C(X \times Y)$ spanned by the functions of the form $f(x, y) = g(x)\,h(y)$, $g \in C(X)$, $h \in C(Y)$, is dense in $C(X \times Y)$.*

Corollary. *If K is a compact subset of $\mathbb{R}^n$, then the polynomials (in n-variables) are dense in $C(K)$.*

Applications to $C^{2\pi}$

In many texts, the Stone-Weierstrass theorem is used to show that the trig polynomials are dense in $C^{2\pi}$. One approach here might be to identify $C^{2\pi}$ with the closed subalgebra of $C[\,0, 2\pi\,]$ consisting of those functions f satisfying $f(0) = f(2\pi)$. Probably easier, though, is to identify $C^{2\pi}$ with the continuous functions on the unit circle $\mathbb{T} = \{e^{i\theta} : \theta \in \mathbb{R}\} = \{z \in \mathbb{C} : |z| = 1\}$ in the complex plane using the identification

$$f \in C^{2\pi} \quad \longleftrightarrow \quad g \in C(\mathbb{T}), \ \text{ where } \ g(e^{it}) = f(t).$$

Under this correspondence, the trig polynomials in $C^{2\pi}$ match up with (certain) polynomials in $z = e^{it}$ and $\overline{z} = e^{-it}$. But, as we've seen, even if we start with real-valued trig polynomials, we'll end up with polynomials in z and $\overline{z}$ having complex coefficients.

Given this, it might make more sense to consider the complex-valued continuous functions on $\mathbb{T}$. We'll write $C_{\mathbb{C}}(\mathbb{T})$ to denote the complex-valued continuous functions on $\mathbb{T}$, and $C_{\mathbb{R}}(\mathbb{T})$ to denote the real-valued continuous functions on $\mathbb{T}$. Similarly, $C^{2\pi}_{\mathbb{C}}$ is the space of complex-valued, 2π-periodic functions on $\mathbb{R}$, while $C^{2\pi}_{\mathbb{R}}$ stands for the real-valued, 2π-periodic functions on $\mathbb{R}$. Now, under the identification we made earlier, we have $C_{\mathbb{C}}(\mathbb{T}) = C^{2\pi}_{\mathbb{C}}$ and $C_{\mathbb{R}}(\mathbb{T}) = C^{2\pi}_{\mathbb{R}}$. The complex-valued trig polynomials in $C^{2\pi}_{\mathbb{C}}$ now match up with the full set of polynomials, with complex coefficients, in $z = e^{it}$ and $\overline{z} = e^{-it}$. We'll use the Stone-Weierstrass theorem to show that these polynomials are dense in $C_{\mathbb{C}}(\mathbb{T})$.

Now the polynomials in z obviously separate points in $\mathbb{T}$ and vanish at no point of $\mathbb{T}$. Nevertheless, the polynomials in z alone are not dense in $C_{\mathbb{C}}(\mathbb{T})$. To see this, here's a proof that $f(z) = \overline{z}$ cannot be uniformly approximated by polynomials in z. First, suppose that we're given some polynomial $p(z) = \sum_{k=0}^n c_k z^k$. Then,

$$\int_0^{2\pi} \overline{f(e^{it})}\, p(e^{it})\, dt \;=\; \int_0^{2\pi} e^{it}\, p(e^{it})\, dt \;=\; \sum_{k=0}^n c_k \int_0^{2\pi} e^{i(k+1)t}\, dt \;=\; 0,$$

and so

$$2\pi \;=\; \int_0^{2\pi} \overline{f(e^{it})}\, f(e^{it})\, dt \;=\; \int_0^{2\pi} \overline{f(e^{it})}\, \big[f(e^{it}) - p(e^{it})\big]\, dt,$$

because $\overline{f(z)}\, f(z) = |f(z)|^2 = 1$. Now, taking absolute values, we get

$$2\pi \;\le\; \int_0^{2\pi} \big|f(e^{it}) - p(e^{it})\big|\, dt \;\le\; 2\pi \|f - p\|.$$

That is, $\|f - p\| \ge 1$ for any polynomial p.

We might as well proceed in some generality: Given a compact metric space X, we'll write $C_{\mathbb{C}}(X)$ for the set of all continuous, complex-valued functions $f : X \to \mathbb{C}$, and we norm $C_{\mathbb{C}}(X)$ by $\|f\| = \max_{x\in X} |f(x)|$ (where $|f(x)|$ is the modulus of the complex number $f(x)$, of course). $C_{\mathbb{C}}(X)$ is a Banach algebra over $\mathbb{C}$. In order to make it clear which field of scalars is involved, we'll write $C_{\mathbb{R}}(X)$ for the real-valued members of $C_{\mathbb{C}}(X)$. Notice, though, that $C_{\mathbb{R}}(X)$ is nothing other than $C(X)$ with a new name.

More generally, we'll write $A_{\mathbb{C}}$ to denote an algebra, over $\mathbb{C}$, of complex-valued functions and $A_{\mathbb{R}}$ to denote the real-valued members of $A_{\mathbb{C}}$. It's not hard to see that $A_{\mathbb{R}}$ is then an algebra, over $\mathbb{R}$, of real-valued functions.

Now if f is in $C_{\mathbb{C}}(X)$, then so is the function $\bar{f}(x) = \overline{f(x)}$ (the complex-conjugate of $f(x)$). This puts

$$\mathrm{Re}f = \frac{1}{2}(f + \bar{f}\,) \quad \text{and} \quad \mathrm{Im}f = \frac{1}{2i}(f - \bar{f}\,),$$

the real and imaginary parts of f, in $C_{\mathbb{R}}(X)$ too. Conversely, if g, $h \in C_{\mathbb{R}}(X)$, then $g + ih \in C_{\mathbb{C}}(X)$.

This simple observation gives us a hint as to how we might apply the Stone-Weierstrass theorem to subalgebras of $C_{\mathbb{C}}(X)$. Given a subalgebra $A_{\mathbb{C}}$ of $C_{\mathbb{C}}(X)$, suppose that we could prove that $A_{\mathbb{R}}$ is dense in $C_{\mathbb{R}}(X)$. Then, given any $f \in C_{\mathbb{C}}(X)$, we could approximate $\mathrm{Re}f$ and $\mathrm{Im}f$ by elements g, $h \in A_{\mathbb{R}}$. But since $A_{\mathbb{R}} \subset A_{\mathbb{C}}$, this means that $g + ih \in A_{\mathbb{C}}$, and $g+ih$ approximates f. That is, $A_{\mathbb{C}}$ is dense in $C_{\mathbb{C}}(X)$. Great! And what did we really use here? Well, we need $A_{\mathbb{R}}$ to contain the real and imaginary parts of "most" functions in $C_{\mathbb{C}}(X)$. If we insist that $A_{\mathbb{C}}$ separate points and vanish at no point, then $A_{\mathbb{R}}$ will contain "most" of $C_{\mathbb{R}}(X)$. And, to be sure that we get both the real and imaginary parts of each element of $A_{\mathbb{C}}$, we'll insist that $A_{\mathbb{C}}$ contain the conjugates of each of its members: $\bar{f} \in A_{\mathbb{C}}$ whenever $f \in A_{\mathbb{C}}$. That is, we'll require that $A_{\mathbb{C}}$ be *self-conjugate* (or, as some authors say, *self-adjoint*).

Stone-Weierstrass Theorem. (complex scalars) *Let X be a compact metric space, and let $A_{\mathbb{C}}$ be a subalgebra, over $\mathbb{C}$, of $C_{\mathbb{C}}(X)$. If $A_{\mathbb{C}}$ separates points in X, vanishes at no point of X, and is self-conjugate, then $A_{\mathbb{C}}$ is dense in $C_{\mathbb{C}}(X)$.*

PROOF. Again, write $A_{\mathbb{R}}$ for the set of real-valued members of $A_{\mathbb{C}}$. Since $A_{\mathbb{C}}$ is self-conjugate, $A_{\mathbb{R}}$ contains the real and imaginary parts of every $f \in A_{\mathbb{C}}$;

$$\mathrm{Re}f = \frac{1}{2}(f + \bar{f}\,) \in A_{\mathbb{R}} \quad \text{and} \quad \mathrm{Im}f = \frac{1}{2i}(f - \bar{f}\,) \in A_{\mathbb{R}}.$$

Moreover, $A_{\mathbb{R}}$ is a subalgebra, over $\mathbb{R}$, of $C_{\mathbb{R}}(X)$. In addition, $A_{\mathbb{R}}$ separates points in X and vanishes at no point of X. Indeed, given $x \neq y \in X$ and $f \in A_{\mathbb{C}}$ with $f(x) \neq f(y)$, we must have at least one of $\mathrm{Re} f(x) \neq \mathrm{Re} f(y)$ or $\mathrm{Im} f(x) \neq \mathrm{Im} f(y)$. Similarly, $f(x) \neq 0$ means that at least one of $\mathrm{Re} f(x) \neq 0$ or $\mathrm{Im} f(x) \neq 0$ holds. That is, $A_{\mathbb{R}}$ satisfies the hypotheses of the real-scalar version of the Stone-Weierstrass theorem. Consequently, $A_{\mathbb{R}}$ is dense in $C_{\mathbb{R}}(X)$.

Now, given $f \in C_{\mathbb{C}}(X)$ and $\varepsilon > 0$, take g, $h \in A_{\mathbb{R}}$ with $\|g - \mathrm{Re} f\| < \varepsilon/2$ and $\|h - \mathrm{Im} f\| < \varepsilon/2$. Then, $g + ih \in A_{\mathbb{C}}$ and $\|f - (g + ih)\| < \varepsilon$. Thus, $A_{\mathbb{C}}$ is dense in $C_{\mathbb{C}}(X)$. □

Corollary. *The polynomials, with complex coefficients, in z and $\overline{z}$ are dense in $C_{\mathbb{C}}(\mathbb{T})$. In other words, the complex trig polynomials are dense in $C^{2\pi}_{\mathbb{C}}$.*

Note that it follows from the complex-scalar proof that the real parts of the polynomials in z and $\overline{z}$, that is, the real trig polynomials, are dense in $C_{\mathbb{R}}(\mathbb{T}) = C^{2\pi}_{\mathbb{R}}$.

Corollary. *The real trig polynomials are dense in $C^{2\pi}_{\mathbb{R}}$.*

Application: Lipschitz Functions

In most Real Analysis courses, the classical Weierstrass theorem is used to prove that $C[\,a, b\,]$ is separable. Likewise, the Stone-Weierstrass theorem can be used to show that $C(X)$ is separable, where X is a compact metric space. While we won't have anything quite so convenient as polynomials at our disposal, we do, at least, have a familiar collection of functions to work with.

Given a metric space $(X, d\,)$, and $0 \le K < \infty$, we'll write $\mathrm{lip}_K(X)$ to denote the collection of all real-valued Lipschitz functions on X with constant at most K; that is, $f : X \to \mathbb{R}$ is in $\mathrm{lip}_K(X)$ if $|f(x) - f(y)| \le K d(x, y)$ for all x, $y \in X$. And we'll write $\mathrm{lip}(X)$ to denote the set of functions that are in $\mathrm{lip}_K(X)$ for some K; in other words, $\mathrm{lip}(X) = \bigcup_{K=1}^{\infty} \mathrm{lip}_K(X)$. It's easy to see that $\mathrm{lip}(X)$ is a subspace of $C(X)$; in fact, if X is compact, then $\mathrm{lip}(X)$ is even a subalgebra of $C(X)$. Indeed, given $f \in \mathrm{lip}_K(X)$ and $g \in \mathrm{lip}_M(X)$, we have

$$\begin{aligned} |f(x)g(x) - f(y)g(y)| &\le |f(x)g(x) - f(y)g(x)| + |f(y)g(x) - f(y)g(y)| \\ &\le K\|g\|\,|x - y| + M\|f\|\,|x - y|. \end{aligned}$$

Lemma. *If X is a compact metric space, then* $\mathrm{lip}(X)$ *is dense in* $C(X)$.

PROOF. Clearly, $\mathrm{lip}(X)$ contains the constant functions and so vanishes at no point of X. To see that $\mathrm{lip}(X)$ separates point in X, we use the fact that the metric d is Lipschitz: Given $x_0 \neq y_0 \in X$, the function $f(x) = d(x, y_0)$ satisfies $f(x_0) > 0 = f(y_0)$; moreover, $f \in \mathrm{lip}_1(X)$ since

$$|f(x) - f(y)| \ = \ |d(x, y_0) - d(y, y_0)| \ \leq \ d(x, y).$$

Thus, by the Stone-Weierstrass Theorem, $\mathrm{lip}(X)$ is dense in $C(X)$. □

Theorem. *If X is a compact metric space, then $C(X)$ is separable.*

PROOF. It suffices to show that $\mathrm{lip}(X)$ is separable. (Why?) To see this, first notice that $\mathrm{lip}(X) = \bigcup_{K=1}^{\infty} E_K$, where

$$E_K = \{f \in C(X) : \|f\| \leq K \text{ and } f \in \mathrm{lip}_K(X)\}.$$

(Why?) The sets E_K are (uniformly) bounded and equicontinuous. Hence, by the Arzelà-Ascoli theorem, each E_K is compact in $C(X)$. Since compact sets are separable, as are countable unions of compact sets, it follows that $\mathrm{lip}(X)$ is separable. □

As it happens, the converse is also true (which is why this is interesting); see Folland's *Real Analysis* for more details.

Theorem. *If $C(X)$ is separable, where X is a compact Hausdorff topological space, then X is metrizable.*

A Short List of References

Books

Abramowitz, M. and Stegun, I., eds., *Handbook of Mathematical Functions with Formulas, Graphs, and Mathematical Tables,* Dover, 1965.

Birkhoff, G., *A Source Book in Classical Analysis,* Harvard, 1973.

Buck, R. C., ed., *Studies in Modern Analysis,* MAA, 1962.

Carothers, N. L., *Real Analysis,* Cambridge, 2000.

Cheney, E. W., *Introduction to Approximation Theory,* Chelsea, 1982.

Davis, P. J., *Interpolation and Approximation,* Dover, 1975.

DeVore, R. A. and Lorentz, G. G., *Constructive Approximation,* Springer-Verlag, 1993.

Dudley, R. M., *Real Analysis and Probability,* Wadsworth & Brooks/Cole, 1989.

Folland, G. B., *Real Analysis: modern techniques and their applications,* Wiley, 1984.

Fox, L. and Parker, I. B., *Chebyshev Polynomials,* Oxford University Press, 1968.

Hoffman, K., *Analysis in Euclidean Space,* Prentice-Hall, 1975.

Jackson, D., *Theory of Approximation,* AMS, 1930.

Jackson, D., *Fourier Series and Orthogonal Polynomials,* MAA, 1941.

Körner, T. W., *Fourier Analysis,* Cambridge, 1988.

Korovkin, P. P., *Linear Operators and Approximation Theory,* Hindustan Publishing, 1960.

La Vallée Poussin, Ch.-J. de, *Leçons sur l'Approximation des Fonctions d'une Variable Réele,* Gauthier-Villars, 1919.

Lorentz, G. G., *Bernstein Polynomials,* Chelsea, 1986.

Lorentz, G. G., *Approximation of Functions,* Chelsea, 1986.

Natanson, I., *Constructive Function Theory,* 3 vols., Ungar, 1964–1965.

Powell, M. J. D., *Approximation Theory and Methods,* Cambridge, 1981.

Rivlin, T. J., *An Introduction to the Approximation of Functions,* Dover, 1981.

Rivlin, T. J., *The Chebyshev Polynomials,* Wiley, 1974.

Rudin, W., *Principles of Mathematical Analysis,* 3rd. ed., McGraw-Hill, 1976.

Simmons, G. F., *Introduction to Topology and Modern Analysis,* McGraw-Hill, 1963; reprinted by Robert E. Krieger Publishing, 1986.

Articles

Boas, R. P., "Inequalities for the derivatives of polynomials," *Mathematics Magazine,* **42** (1969), 165–174.

Fisher, S., "Quantitative approximation theory," *The American Mathematical Monthly,* **85** (1978), 318–332.

HEDRICK, E. R., "The significance of Weierstrass's theorem," *The American Mathematical Monthly,* **20** (1927), 211–213.

JACKSON, D., "The general theory of approximation by polynomials and trigonometric sums," *Bulletin of the American Mathematical Society,* **27** (1920–1921), 415–431.

LEBESGUE, H., "Sur l'approximation des fonctions," *Bulletin des Sciences Mathématique,* **22** (1898), 278–287.

SHIELDS, A., "Polynomial approximation," *The Mathematical Intelligencer,* **9** (1987), No. 3, 5–7.

SHOHAT, J. A., "On the development of functions in series of orthogonal polynomials," *Bulletin of the American Mathematical Society,* **41** (1935), 49–82.

STONE, M. H., "Applications of the theory of Boolean rings to general topology," *Transactions of the American Mathematical Society,* **41** (1937), 375–481.

STONE, M. H., "A generalized Weierstrass theorem," in *Studies in Modern Analysis*, R. C. Buck, ed., MAA, 1962.

VAN VLECK, E. B., "The influence of Fourier's series upon the development of mathematics," *Science,* **39** (1914), 113–124.

WEIERSTRASS, K., "Über die analytische Darstellbarkeit sogenannter willkürlicher Functionen einer reellen Veränderlichen," *Sitzungsberichte der Königlich Preussischen Akademie der Wissenshcaften zu Berlin,* (1885), 633–639, 789–805.

WEIERSTRASS, K., "Sur la possibilité d'une représentation analytique des fonctions dites arbitraires d'une variable réele," *Journal de Mathématiques Pures et Appliquées,* **2** (1886), 105–138.