


Lecture 21: Minimax Approximation


In many applications, the L^2 -norm has a physical interpretation – often associated with some measure of energy – and this makes continuous least-squares approximation particularly appealing (e.g., the best approximation minimizes the energy of the error). Moreover, that optimal polynomial approximation in this norm can be computed at the expense of a few inner products (integrals). However, like discrete least-squares, this approach suffers from a potential problem: it minimizes the influence of outlying data, i.e., points where the function f varies wildly over a small portion of the interval $[a, b]$. Such an example is shown below.[†]



For this function the L^2 -norm of the error,

$$\|f - p_*\|_{L^2} = \left(\int_a^b (f(x) - p_*(x))^2 dx \right)^{1/2},$$

averages out the discrepancy $f - p_*$ over all $x \in [a, b]$, so it is possible to have a large error $f(x) - p_*(x)$ on some narrow range of x values that makes a negligible contribution to the integral. Below on the left, we compare the function shown above to its degree-5 least-squares approximation; on the right, we show the error $f - p_*$, which is small throughout $[-1, 1]$ except for a large spike.




[†]The function in question is $f(x) = \sin(\pi x) + 3 \exp(-50(x - \frac{1}{2})^2)$. Despite the nasty appearance of the plot, this function is perfectly smooth: $f(x) \in C^\infty[-1, 1]$.

Functions of this sort may seem pathological, but they highlight the fact that L^2 -optimization does not always generate a polynomial p_* that is close to f throughout the interval $[a, b]$. Indeed, in a number of settings the L^2 -norm is the wrong way to measure error: we really want to minimize $\max_{x \in [a, b]} |f(x) - p(x)|$.

3.4 Minimax approximation. The goal of minimizing the maximum deviation of a polynomial p from our function f is called *minimax* (or *uniform*, or L^∞) approximation, since

$$\min_{p \in \mathcal{P}_k} \max_{x \in [a, b]} |f(x) - p(x)| = \min_{p \in \mathcal{P}_k} \|f - p\|_{L^\infty}.$$


A simple example. Suppose we seek the constant that best approximates $f(x) = e^x$ over the interval $[0, 1]$, shown below.



Since $f(x)$ is monotonically increasing for $x \in [0, 1]$, the optimal constant approximation $p_* = c_0$ must fall somewhere between $f(0) = 1$ and $f(1) = e$, i.e., $1 \leq c_0 \leq e$. Moreover, since f is monotonic and p_* is a constant, the function $f - p_*$ is also monotonic, so the maximum error $\max_{x \in [a, b]} |f(x) - p_*(x)|$ must be attained at one of the end points, $x = 0$ or $x = 1$. Thus,

$$\|f - p_*\|_{L^\infty} = \max\{|e^0 - c_0|, |e^1 - c_0|\}.$$

The following figure shows $|e^0 - c_0|$ (broken line) and $|e^1 - c_0|$ (dotted line) for $c_0 \in [1, e]$.




The optimal value for c_0 will be the point at which the larger of these two lines is minimal. The figure above clearly reveals that this happens when the errors are equal, at $c_0 = (1 + e)/2$. We conclude that the optimal minimax constant polynomial approximation to e^x on $x \in [0, 1]$ is $p_*(x) = c_0 = (1 + e)/2$.

The plots below compare f to the optimal polynomial p_* (left), and show the error $f - p_*$ (right). We picked c_0 to be the point at which the error was equal in magnitude at the end points $x = 0$ and $x = 1$; in fact, it is equal in magnitude, but opposite in sign,

$$e^0 - c_0 = -(e^1 - c_0),$$

as seen in the illustration on the right below. It turns out that this property—maximal error attained at various points in the interval with alternating sign—is a key feature of minimax approximation.



3.4.1. Oscillation Theorem. As hinted in the previous example, the points at which the error $f - p_*$ attains its maximum magnitude play a central role in the theory of minimax approximation. The Theorem of de la Vallée Poussin is a first step toward such a result. We include its proof (from Süli and Mayers, §8.3) to give a general impression of how such results are established.

Theorem (de la Vallée Poussin Theorem). Let $f \in C[a, b]$ and suppose $r \in \mathcal{P}_n$ is some polynomial for which there exist $n + 2$ points $\{x_j\}_{j=0}^{n+1}$ with $a \leq x_0 < x_1 < \dots < x_{n+1} \leq b$ at which the error $f(x) - r(x)$ oscillates signs, i.e.,


$$\text{sign}(f(x_j) - r(x_j)) = -\text{sign}(f(x_{j+1}) - r(x_{j+1}))$$

for $j = 0, \dots, n$. Then

$$\min_{p \in \mathcal{P}_n} \|f - p\|_{L^\infty} \geq \min_{0 \leq j \leq n+1} |f(x_j) - r(x_j)|.$$

Before proving this result, we provide a numerical illustration. Here we are approximating $f(x) = e^x$ with a quintic polynomial, $r \in \mathcal{P}_5$ (i.e., $n = 5$). *This polynomial is not necessarily the minimax approximation to f over the interval $[0, 1]$.* However, in the plot below we can see that for this r , we can find $n + 2 = 7$ points at which the sign of the error $f(x) - r(x)$ oscillates. The broken

line shows the error curve for the optimal minimax polynomial p_* (whose computation is discussed below). Here is the point of the de la Vallée Poussin theorem: *Since the error $f(x) - r(x)$ oscillates sign $n + 2$ times, there must be some $x \in [0, 1]$ at which the minimax error $\pm \|f - p_*\|_{L^\infty}$ (denoted by the horizontal broken lines) exceeds $|f(x) - r(x)|$ at one of the points that give the oscillating sign.* In other words, the de la Vallée Poussin theorem provides a mechanism for developing *lower bounds* on $\|f - p_*\|_{L^\infty}$.



Proof. Suppose we have $n + 2$ ordered points, $\{x_j\}_{j=0}^{n+1} \subset [a, b]$, such that $f(x_j) - r(x_j)$ alternates sign at consecutive points, and let p_* denote the minimax polynomial,

$$\|f - p_*\|_{L^\infty} = \min_{p \in \mathcal{P}_n} \|f - p\|_{L^\infty}.$$

We will prove the result by contradiction. Thus suppose

$$\|f - p_*\|_{L^\infty} < |f(x_j) - r(x_j)|, \quad \text{for all } j = 0, \dots, n + 1. \quad (21.1)$$

As the left hand side is the maximum difference of $f - p_*$ over all $x \in [a, b]$, that difference can be no larger at $x_j \in [a, b]$, and so:

$$|f(x_j) - p_*(x_j)| < |f(x_j) - r(x_j)|, \quad \text{for all } j = 0, \dots, n + 1. \quad (21.2)$$

Now consider $p_*(x) - r(x) = (f(x) - r(x)) - (f(x) - p_*(x))$, which is a degree n polynomial, since $p_*, r \in \mathcal{P}_n$. Equation (21.2) states that $f(x_j) - r(x_j)$ always has larger magnitude than $f(x_j) - p_*(x_j)$. Thus, regardless of the sign of $f(x_j) - p_*(x_j)$, the magnitude $|f(x_j) - p_*(x_j)|$ will never be large enough to overcome $|f(x_j) - r(x_j)|$, and hence $p_*(x_j) - r(x_j)$ will always have the same sign as $f(x_j) - r(x_j)$. We know from the hypothesis that $f(x) - r(x)$ must change sign at least $n + 1$ times (at least once in each interval (x_j, x_{j+1}) for $j = 0, \dots, n$), and thus $p_*(x) - r(x) \in \mathcal{P}_n$ must do the same. But $n + 1$ sign changes implies $n + 1$ roots; the only degree n polynomial with $n + 1$ roots is the zero polynomial, i.e., $p_* = r$. However, this contradicts the strict inequality in equation (21.1). Hence, there must be at least one j for which

$$\|f - p_*\|_{L^\infty} \geq |f(x_j) - r(x_j)|. \quad \blacksquare$$

The following result has the same flavor, but it is considerably more precise (with a more intricate proof, which we omit).[‡]

Theorem (Oscillation Theorem). Suppose $f \in C[a, b]$. Then $p_* \in \mathcal{P}_n$ is a minimax approximation to f from \mathcal{P}_n on $[a, b]$ if and only if there exist $n + 2$ points $x_0 < x_1 < \cdots < x_{n+1}$ such that

$$|f(x_j) - p_*(x_j)| = \|f - p_*\|_{L^\infty}, \quad j = 0, \dots, n + 1$$

and

$$f(x_j) - p_*(x_j) = -(f(x_{j+1}) - p_*(x_{j+1})), \quad j = 0, \dots, n.$$

In words, this means that the optimal error, $f - p_*$, attains its maximum at $n + 2$ points, with the error alternating sign between consecutive points.

Note that this result is *if and only if*: the oscillation property exactly *characterizes* the minimax approximation. If you can present some polynomial $p_* \in \mathcal{P}_n$ such that $f - p_*$ satisfies the oscillation property, then this p_* must be the unique minimax approximation!

Theorem (Uniqueness of minimax approximant). The minimax approximant $p_* \in \mathcal{P}_n$ of $f \in C[a, b]$ over the interval $[a, b]$ is unique.


The proof is a straightforward application of the Oscillation Theorem. One can show that any two potential minimax polynomials must have the same $n + 2$ critical oscillation points. Any two degree- n polynomials that agree at $n + 2$ points must be identical. See Süli and Mayers, Theorem 8.5, for details.

This oscillation property forms the basis of algorithms that find the minimax approximation: iteratively adjust an approximating polynomial until it satisfies the oscillation property. The most famous algorithm for computing the minimax approximation is called the *Remez exchange algorithm*, essentially a specialized linear programming procedure. In exact arithmetic, this algorithm is guaranteed to terminate with the correct answer in finitely many operations.

The oscillation property is demonstrated in the previous example, where we approximated $f(x) = e^x$ with a constant. Indeed, the maximum error is attained at two points (that is, $n + 2$, since $n = 0$), and the error differs in sign at those points. The pictures below show the errors $f(x) - p_*(x)$ for minimax approximations p_* of increasing degree.[§] The oscillation property becomes increasingly apparent as the polynomial degree increases. In each case, there are $n + 2$ extreme points of the error, where n is the degree of the approximating polynomial.

[‡]For a proof, see Süli and Mayers, §8.3. Another excellent resource is G. W. Stewart, *Afternotes Goes to Graduate School*, SIAM, 1998; see Stewart's Lecture 3.

[§]These examples were computed using the COCA package, software written by Bernd Fischer and Jan Modersitski that even solves minimax approximation problems when the interval $[a, b]$ is replaced by a region of the complex plane.



Example: e^x revisited. Now we shall use the Oscillation Theorem to compute the optimal linear minimax approximation to $f(x) = e^x$ on $[0, 1]$. Assume that the minimax polynomial $p_* \in \mathcal{P}_1$ has the form $p_*(x) = \alpha + \beta x$. Since f is convex, a quick sketch of the situation suggests the maximal error will be attained at the end points of the interval, $x_0 = 0$ and $x_2 = 1$. We assume this to be true, and seek some third point $x_1 \in (0, 1)$ that attains the same maximal error, δ , but with opposite sign. If we can find such a point, then by the Oscillation Theorem, we are guaranteed that the resulting polynomial is optimal, confirming our assumption that the maximal error was attained at the ends of the interval.

This scenario suggests the following three equations:

$$\begin{aligned} f(x_0) - p_*(x_0) &= \delta \\ f(x_1) - p_*(x_1) &= -\delta \\ f(x_2) - p_*(x_2) &= \delta. \end{aligned}$$

Substituting our values for x_0 , x_2 , and $p_*(x) = \alpha + \beta x$, these equations become

$$\begin{aligned}1 - \alpha &= \delta \\ e^{x_1} - \alpha - \beta x_1 &= -\delta \\ e - \alpha - \beta &= \delta.\end{aligned}$$

The first and third equation together imply $\beta = e - 1$. We also deduce that $2\alpha = e^{x_1} - x_1(e - 1) + 1$. There are a variety of choices for x_1 that will satisfy these conditions, but in those cases δ will not be the *maximal error*. It is key that

$$|\delta| = \max_{x \in [a, b]} |f(x) - p_*(x)|.$$


To make sure this happens, we can require that *the derivative of error* be zero at x_1 , reflecting that the error $f - p_*$ attains a local minimum/maximum at x_1 . The pictures on the previous page confirm that this is reasonable.[¶] Imposing the condition that $f'(x_1) - p'_*(x_1) = 0$ yields

$$e^{x_1} - \beta = 0.$$

Now we can explicitly solve the equations to obtain

$$\begin{aligned}\alpha &= \frac{1}{2}(e - (e - 1) \log(e - 1)) = 0.89406\dots \\ \beta &= e - 1 = 1.71828\dots \\ x_1 &= \log(e - 1) = 0.54132\dots \\ \delta &= \frac{1}{2}(2 - e + (e - 1) \log(e - 1)) = 0.10593\dots\end{aligned}$$

An illustration of the optimal linear approximation we have just computed, along with the associated error, is shown below. Compare this approximation to the L^2 -optimal linear polynomial computed at the beginning of our study of continuous least-squares minimization.



[¶]This requirement need not hold at the points x_0 and x_2 , since these points are on the ends of the interval $[a, b]$; it is only required at the interior points where the extreme error is attained, $x_j \in (a, b)$.