Approximation Theory

Annie Cuyt

1 Introduction

Approximation theory is an area of mathematics that has become indispensable to the computational sciences. The approximation of magnitudes and functions describing some physical behavior is an integral part of scientific computing, queueing problems, neural networks, graphics, robotics, network traffic, financial trading, antenna design, floating-point arithmetic, image processing, speech analysis, and video signal filtering, to name just a few areas.

The idea of seeking a simple mathematical function that describes some behavior approximately has two sources of motivation. The exact behavior that one is studying may not be able to be expressed in a closed mathematical formula. But even if an exact description is available it may be far too complicated for practical use. In both cases a best and simple approximation is required. What is meant by best and simple depends on the application at hand.

The approximation is often used in a computer implementation, and therefore its evaluation needs to be efficient. The simplest and fastest functions for implementation are polynomials, because they use only the fast hardware operations of addition and multiplication. Next come rational functions, which also need the hardware operation division, one or more depending on their representation as a quotient of polynomials or as a continued fraction. Rational functions offer the clear advantage that they can reproduce asymptotic behavior (vertical, horizontal, slant), which is something polynomials are incapable of doing. For periodic phenomena, linear combinations of trigonometric functions make good candidates. For growth models or decaying magnitudes, linear combinations of exponentials can be used.

In approximation theory one distinguishes between interpolation and so-called least-squares problems. In the former one wants the approximate model to take exactly the same values as prescribed by data given at precise argument values. In the latter a set of data (not necessarily discrete) is regarded as a trend and approximated by a simple model in one or other best sense. The difference is formalized in the following sections.

Besides constructing a good and efficient mathematical model, one should also take the following two issues into account.

- What can be said about the convergence of the selected mathematical model? In more practical terms: does the model improve when one adds more data?
- How sensitive is the mathematical model to perturbations in the input data? Data errors are usually unavoidable, and one wishes to know how much they can be magnified in the approximation process.

In the following sections we comment on both issues where appropriate. We do not aim to discuss convergence or undertake a sensitivity analysis for every technique.

Despite the need for and interest in multidimensional models and simulations, we restrict ourselves here mostly to one-dimensional approximation problems. In the penultimate section we include some brief remarks on multivariate interpolation and approximation and its additional complexity.

2 Numerical Interpolation

Let data f_i be given at points $x_i \in [a,b]$, where $i = 0,\ldots,n$. We assume that if some of the points x_i are repeated, then it is not only the value of some underlying function $f(x)$ that is given (or measured) at x_i but also as many higher derivatives $f^{(j)}(x_i)$ as there are copies of the point x_i. The interpolation problem is to find a function of a specified form that matches all the data at the points. In this section we deal with the two extreme cases: the one in which all the points x_i are mutually distinct and no derivative information is available, and the one in which the value of the function and that of the first n derivatives are all given at one single point x_0. Of course, intermediate situations can also be dealt with. The approximating functions that we consider are polynomials, piecewise polynomials (splines), and rational functions, each of which has particular advantages. Finally, we present the connection between exponential models and sparse interpolation on the one hand, and exponential models and Padé approximation on the other.

2.1 Polynomial Interpolation

For $n+1$ given values $f_i = f(x_i)$ at mutually distinct points x_i, the polynomial interpolation problem

of degree n,

$$p_n(x) = \sum_{j=0}^{n} a_j x^j, \qquad p_n(x_i) = f_i, \quad i = 0, \ldots, n, \quad (1)$$

has a unique solution for the coefficients a_j. Now let us turn to the computation of $p_n(x)$. Essentially, two approaches can be used, depending on the intended subsequent use of the polynomial interpolant. If one is interested in easily updating the polynomial interpolant by adding an extra data point and consequently increasing the degree of $p_n(x)$, then Newton's formula for the interpolating polynomial is very suitable. If one wants to use the interpolant for several sets of values f_i while keeping the points x_i fixed, then Lagrange's formula is most appropriate. A simple rearrangement of the Lagrange form as in (2) below results in the barycentric form, which combines the advantages of both approaches.

In the Newton form one writes the interpolating polynomial $p_n(x)$ as

$$p_n(x) = b_0 + b_1(x - x_0) + b_2(x - x_0)(x - x_1) + \cdots + b_n(x - x_0) \cdots (x - x_{n-1}).$$

The coefficients b_j then equal the *divided differences* $b_j = f[0, \ldots, j]$ obtained from the recursive scheme

$$\begin{aligned} f[j] &= f_j, \quad j = 0, \ldots, n, \\ f[0, j] &= \frac{f_j - f_0}{x_j - x_0}, \quad j = 1, \ldots, n, \\ f[0, 1, \ldots, k-1, k, j] &= \frac{f[0, 1, \ldots, k-1, j] - f[0, 1, \ldots, k-1, k]}{x_j - x_k}, \\ & \qquad k, j = 2, \ldots, n. \end{aligned}$$

Newton's form for the interpolating polynomial is very handy when one wants to update the interpolation with an additional point (x_{n+1}, f_{n+1}). It suffices to add the term

$$b_{n+1}(x - x_0) \cdots (x - x_n)$$

to $p_n(x)$ (which does not destroy the previous interpolation conditions since it evaluates to zero at all the previous x_i) and to complement the recursive scheme for the computation of the divided differences with the computation of the

$$f[0, 1, \ldots, k, n+1], \quad k = 0, \ldots, n.$$

In the Lagrange form, which is especially suitable if the interpolation needs to be repeated for different sets of f_i at the same points x_i, another form for $p_n(x)$ is used. We write

$$p_n(x) = \sum_{j=0}^{n} c_j \beta_j(x), \quad \beta_j(x) = \prod_{\substack{k=0 \\ k \neq j}}^{n} \frac{(x - x_k)}{(x_j - x_k)}.$$

The basis functions $\beta_j(x)$ satisfy a simple interpolation condition themselves, namely,

$$\beta_j(x_i) = \begin{cases} 0 & \text{for } j \neq i, \\ 1 & \text{for } j = i. \end{cases}$$

The choice $c_j = f_j$ for the coefficients therefore solves the interpolation problem. So when altering the f_i, without touching the x_i that make up the basis functions $\beta_j(x)$, it takes no computation at all to get the new coefficients c_j.

The *barycentric form* of the interpolation polynomial,

$$\left.\begin{aligned} p_n(x) &= (x - x_0) \cdots (x - x_n) \sum_{j=0}^{n} \frac{w_j}{x - x_j} f_j, \\ w_j &= \Big(\prod_{\substack{k=0 \\ k \neq j}}^{n} (x_j - x_k) \Big)^{-1}, \end{aligned}\right\} \quad (2)$$

is easy to update and is BACKWARD STABLE [**??**] for evaluation of $p_n(x)$.

The sensitivity of polynomial interpolation expressed in the Lagrange form is measured by the value

$$L_n = \max_{a \leqslant x \leqslant b} \sum_{j=0}^{n} |\beta_j(x)|, \quad (3)$$

which is also known as the *Lebesgue constant*. The growth rate of L_n with n is only logarithmic when the interpolation points are as in (5) below. This is the slowest possible growth for polynomial interpolation.

Despite the simplicity and elegance of polynomial interpolation, the technique has a significant drawback, as we discuss next: it may not converge for data $f_i = f(x_i)$ given at arbitrary points x_i, even if $f(x)$ is continuous on $[a, b]$.

2.2 The Runge Phenomenon

What happens if we continue updating the interpolation problem with new data? In other words, what happens if we let the degree n of the interpolating polynomial $p_n(x)$ increase? Will the interpolating polynomial of degree n become better and better? The answer is no, at least not for freely chosen points x_i. To see what can go wrong, consider

$$f(x) = \frac{1}{1 + 25x^2}, \quad -1 \leqslant x \leqslant 1, \quad (4)$$

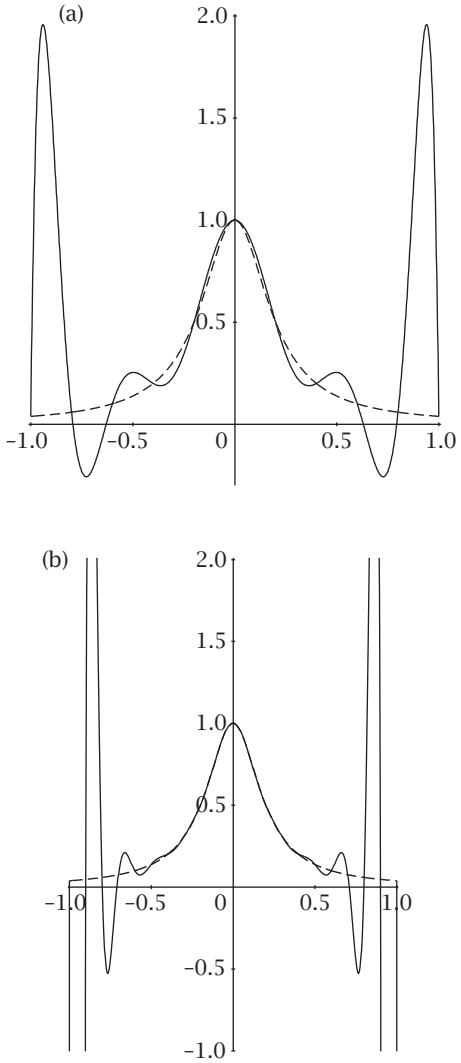

Figure 1 (a) Degree-10 and (b) degree-20 equidistant interpolation (solid lines) for function f in (4) (dashed lines).

and take equidistant interpolation points $x_i = -1 + 2i/n$, $i = 0, \ldots, n$. The error $(f-p)(x)$ toward the endpoints of the interval then increases dramatically with n. Take a look at the bell-shaped $f(x)$ and the interpolating polynomial $p_n(x)$ for $n = 10$ and $n = 20$ in figure 1.

This phenomenon is called *Runge's phenomenon*, after Carl Runge, who described this behavior for real-valued interpolation in 1901. An explanation for it can be found in the fundamental theorem of algebra, which states that a polynomial has as many zeros as its degree. Each of these zeros can be real or complex. So if n is large and the zeros are all real, the polynomial under consideration displays rather oscillatory behavior.

On the other hand, under certain simple conditions for $f(x)$ besides continuity in $[a,b]$, it can be proved that if the interpolation points x_i equal

$$x_i = \frac{a+b}{2} + \frac{b-a}{2} \cos\left(\frac{(2i+1)\pi}{2(n+1)}\right), \qquad i = 0, \ldots, n, \quad (5)$$

where the values $\cos((2i+1)/(n+1)(\pi/2))$ are the zeros of the Chebyshev polynomial of the first kind of degree $n+1$ (defined in section 3.3), then

$$\lim_{n\to\infty} \|f - p_n\|_\infty = \lim_{n\to\infty} \max_{x\in[-1,1]} |(f-p_n)(x)| = 0.$$

The effect of this choice of interpolation points is illustrated in figure 2(a). A similar result holds if the zeros of the Chebyshev polynomial of degree $n+1$ are replaced by the extrema $\cos(i\pi/n)$, $i = 0, \ldots, n$, of the Chebyshev polynomial of degree n.

In order to make use of this result in real-world applications, where the interpolation points x_i cannot usually be chosen arbitrarily, interpolation at the Chebyshev zeros is mimicked, for instance by selecting a proper subset of interpolation points $\tilde{x}_i$ from a fine equidistant grid, with $\tilde{x}_i \approx x_i$ from (5). The grid is considered to be sufficiently fine when the distance between the points ensures that a grid point nearest to a Chebyshev zero x_i is never repeated. In a coarse grid, the same grid point may be the closest one to more than one Chebyshev zero, especially toward the ends of the interval $[-1,1]$.

This technique is called *mock-Chebyshev interpolation*. For comparison, in figure 2(b) we display the degree-20 mock-Chebyshev interpolant with the interpolation points selected from an equispaced grid with gap $1/155$.

If a lot of accurate data points have to be used in an interpolation scheme, then splines, which are discussed in the next section, offer a better alternative than a monolithic high-degree polynomial interpolant.

2.3 Spline Interpolation

In order to avoid the Runge phenomenon when interpolating large data sets, piecewise polynomials, also called splines, can be used. To this end we divide the data set of $n+1$ points into smaller sets each containing two data points. Rather than interpolating the full data

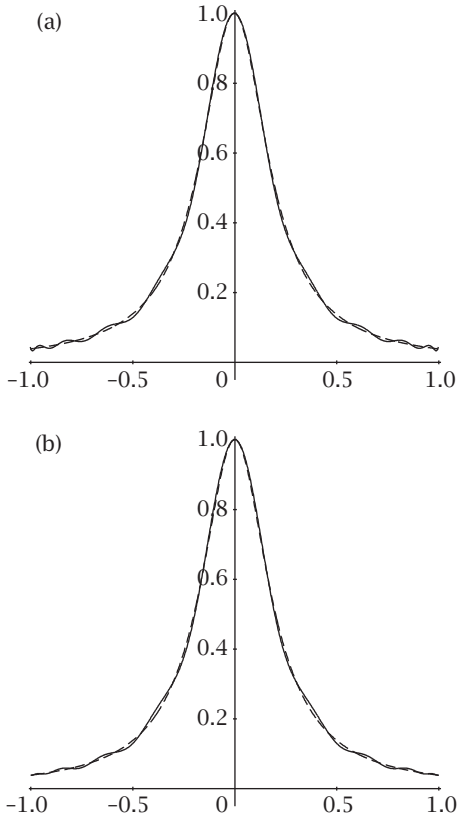

Figure 2 Degree-20 (a) Chebyshev and (b) mock-Chebyshev interpolation (solid lines) for function f in (4) (dashed lines).

set by one polynomial of degree n, we interpolate each of the smaller data sets by a low-degree polynomial. These separate polynomial functions are then pieced together in such a way that the resulting function is as continuously differentiable as possible.

Take, for instance, the data set (x_i, f_i) and consider linear polynomials interpolating every two consecutive (x_i, f_i) and (x_{i+1}, f_{i+1}). These linear polynomial pieces can be joined together at the data points (x_i, f_i) to produce a piecewise-linear continuous function or polygonal curve. Note that this function is continuous but not differentiable at the interpolation points since it is polygonal.

If we introduce two parameters, Δ and D, to respectively denote the degree of the polynomial pieces and the differentiability of the overall function, where obviously $D \leqslant \Delta$ (even $D < \Delta$ to avoid an overdetermined system of defining equations, as explained below), then, for the polygonal curve, $\Delta = 1$ and $D = 0$. With $\Delta = 2$ and $D = 1$, a piecewise-quadratic and smooth (meaning continuously differentiable in the entire interval $[x_0, \ldots, x_n]$) function is constructed. The slope of a smooth function is a continuous quantity. With $\Delta = 3$ and $D = 2$, a piecewise-cubic and twice continuously differentiable function is obtained. Twice continuously differentiable functions also enjoy continuous curvature. Can the naked eye distinguish between continuous and discontinuous curvature in a function? The untrained eye certainly cannot! As an example we take the cubic polynomial pieces

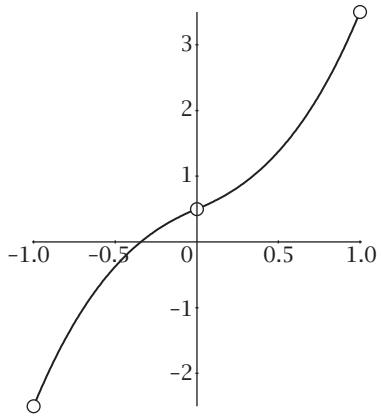

Figure 3 A piecewise-cubic function that is not twice continuously differentiable.

$$c_1(x) = x^3 - x^2 + x + 0.5, \quad x \in [-1, 0],$$
$$c_2(x) = x^3 + x^2 + x + 0.5, \quad x \in [0, 1],$$

and join these together at $x = 0$ to obtain a new piecewise-cubic function $c(x)$ on $[-1, 1]$. The result is a function that is continuous and differentiable at the origin, but for the second derivatives at the origin we have $\lim_{x \to 0-} c^{(2)}(0) = -2$ and $\lim_{x \to 0+} c^{(2)}(0) = 2$. Nevertheless, the result of the gluing procedure shown in figure 3 is a very pleasing function that at first sight looks fine. But while Δ equals 3, D is only 1.

Since a trained eye can spot these discontinuities, the most popular choice for piecewise-polynomial interpolation in industrial applications is $\Delta = 3$ and $D = 2$. Indeed, for manufacturing the continuity of the curvature is important.

Let us take a look at the general situation where $\Delta = m$ and $D = m - 1$, for which the resulting piecewise polynomial is called a *spline*. Assume we are given the interpolation points $x_0, \ldots, x_n$. With these $n + 1$ points we can construct n intervals $[x_i, x_{i+1}]$. The points x_0 and x_n are the endpoints and the other $n - 1$ interpo-

lation points are called the internal points. If $\Delta = m$, then for every interval $[x_i, x_{i+1}]$ we have to determine $m+1$ coefficients, because the explicit formula for the spline on $[x_i, x_{i+1}]$ is a polynomial of degree m:

$$S(x) = s_i(x), \quad x \in [x_i, x_{i+1}],\ i = 0, \ldots, n-1,$$
$$s_i(x) = \sum_{j=0}^{m} a_j^{(i)} x^j.$$

So, in total, $n(m+1)$ unknown coefficients $a_j^{(i)}$ have to be computed. From which conditions? There are the $n+1$ interpolation conditions $S(x_i) = f_i$, and we have the smoothness or continuity requirements at the internal points, meaning that a number of derivatives of $s_{i-1}(x)$ evaluated at the right endpoint of the domain $[x_{i-1}, x_i]$ should coincide with the derivatives of $s_i(x)$ when evaluated at the left endpoint of the domain $[x_i, x_{i+1}]$:

$$s_{i-1}^{(k)}(x_i) = s_i^{(k)}(x_i), \quad i = 1, \ldots, n-1,\ k = 0, \ldots, m-1.$$

The latter requirements add another $(n-1)m$ continuity conditions. This brings us to a total of $n+1+(n-1)m = n(m+1)-m+1$ conditions for $n(m+1)$ unknowns. In other words, we lack $m-1$ conditions to determine the degree-m piecewise-polynomial interpolant with overall smoothness of order $m-1$. When $m = 1$, which is the case for the piecewise-linear spline or the polygonal curve, no conditions are lacking. When $m = 2$, a value for $s_0'(x_0)$ is usually given as an additional piece of information. When $m = 3$, which is the case for the widely used cubic spline, values for $s_0''(x_0)$ and $s_{n-1}''(x_n)$ are often provided (the cubic spline with clamped end conditions) or they are set to zero (the natural cubic spline).

The natural cubic spline interpolant has a very elegant property, namely, that it avoids oscillatory behavior between interpolation points. More precisely, for every twice continuously differentiable function $f(x)$ defined on $[a, b]$ and satisfying $f(x_i) = f_i$ for all i, we have

$$\int_a^b S''(x)^2 \, \mathrm{d}x \leqslant \int_a^b f''(x)^2 \, \mathrm{d}x.$$

A simple illustration is given in figure 4 for $n = 6$ with $x_i = i$, $i = 1, \ldots, 6$.

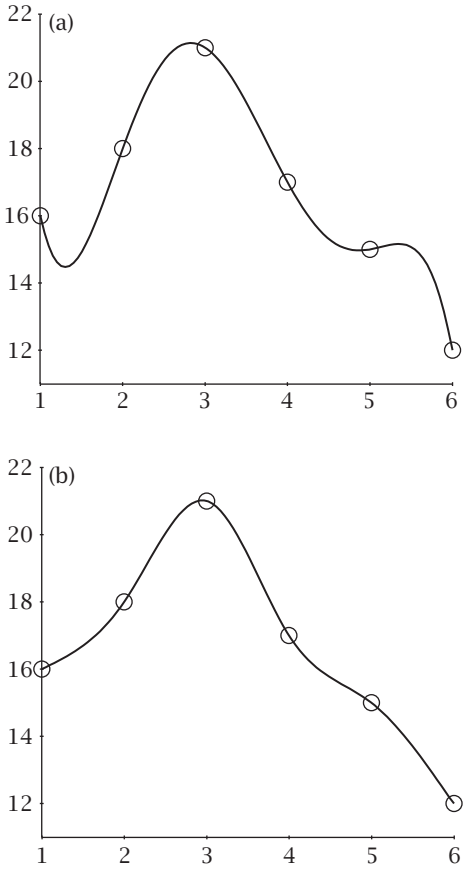

Figure 4 (a) The polynomial interpolant and (b) the natural cubic spline.

2.4 Padé Approximation

The rational equivalent of the Taylor series partial sum is the irreducible rational function $r_{k,\ell}(x) = p_{k,\ell}(x)/q_{k,\ell}(x)$ with numerator of degree at most k and denominator of degree at most ℓ that satisfies

$$r_{k,\ell}^{(i)}(x_0) = f^{(i)}(x_0), \quad i = 0, 1, \ldots, n, \tag{6}$$

with n as large as possible. It is also called the $[k/l]$ *Padé approximant*. The aim is to have $n = k + \ell$. Note that we are imposing one fewer condition than the total number $k + \ell + 2$ of coefficients in $r_{k,\ell}$. The reason is that one degree of freedom is lost because multiplying $p_{k,\ell}$ and $q_{k,\ell}$ by a scalar does not change $r_{k,\ell}$.

A key question is whether n can be less than $k + \ell$. The answer to this question requires some analysis. Computing the numerator and denominator coefficients of $r_{k,\ell}(x)$ from (6) gives rise to a nonlinear system of equations. So let us explore whether the Padé approximant can also be obtained from the linearized approximation conditions

$$(f q_{k,\ell} - p_{k,\ell})^{(i)}(x_0) = 0, \quad i = 0, 1, \ldots, k+\ell. \tag{7}$$

We denote $f^{(j)}(x_0)/j!$ by d_j, where $d_j = 0$ for $j < 0$. The linearized conditions (7) always have at least one nontrivial solution for the numerator coefficients

$a_0, \ldots, a_k$ and the denominator coefficients $b_0, \ldots, b_\ell$ because they form a homogeneous linear system of $k + \ell + 1$ conditions in $k + \ell + 2$ unknowns:

$$\begin{aligned} d_0 b_0 &= a_0, \\ d_1 b_0 + d_0 b_1 &= a_1, \\ &\vdots \\ d_k b_0 + \cdots + d_{k-\ell} b_\ell &= a_k, \\ d_{k+1} b_0 + \cdots + d_{k-\ell+1} b_\ell &= 0, \\ &\vdots \\ d_{k+\ell} b_0 + \cdots + d_k b_\ell &= 0. \end{aligned}$$

Moreover, all solutions $p_{k,\ell}(x)$ and $q_{k,\ell}(x)$ of (7) are equivalent in the sense that they have the same irreducible form. Every solution of (6) with $n = k + \ell$ therefore also satisfies (7), but not vice versa. From $p_{k,\ell}(x)$ and $q_{k,\ell}(x)$ satisfying (7) we find, for the unique irreducible form $p^*_{k,\ell}(x)/q^*_{k,\ell}(x)$, that

$$(f - r_{k,\ell})^{(i)}(x_0) = 0, \quad i = 0, \ldots, k' + \ell' + r,$$
$$k' = \partial p^*_{k,\ell},\ \ell' = \partial q^*_{k,\ell},\ r \geqslant 0,$$

where ∂p denotes the degree of the polynomial p. In some textbooks, the $[k/l]$ Padé approximation problem is said to have no solution if $k' + \ell' + r < k + \ell$; in others, the Padé approximant $r_{k,\ell}$ is identified with $r_{k',\ell'} = p^*_{k,\ell}/q^*_{k,\ell}$ if that is the case (this is the convention we adopt here). Let us illustrate the situation with a simple example. Take $x_0 = 0$ with $d_0 = 1$, $d_1 = 0$, $d_2 = 1$, and $k = 1 = \ell$. The linearized conditions (7) are then

$$b_0 = a_0, \quad b_1 = a_1, \quad b_0 = 0.$$

A solution is given by $p_{1,1}(x) = x$ and $q_{1,1}(x) = x$. We therefore find $r_{1,1}(x) = 1$, $k' = 0$, $\ell' = 0$, and

$$(f - r_{1,1})^{(2)}(x_0) = 2 \neq 0.$$

Since $r = 1$ we have $k' + \ell' + r = 1 < k + \ell = 2$.

This kind of complication does not occur when $\ell = 0$. The Padé approximant $r_{k,0}(x)$ is then merely the Taylor series partial sum of degree k. But when asymptotic behavior needs to be reproduced, a polynomial function is not very useful. In figure 5 one can compare the Taylor series partial sum of degree 9 with the $[5/4]$ Padé approximant for the function $f(x) = \arctan(x)$.

Padé approximants can be organized in a table, where the numerator degree indicates the row and the denominator degree the column. To illustrate this we give part of the Padé table for $f(x) = \sin(x)$ in table 1. A sequence of Padé approximants in the Padé table can converge uniformly or in measure only to a function

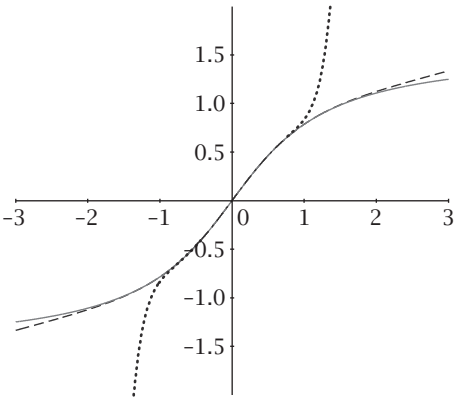

Figure 5 Padé approximants $r_{9,0}(x)$ (dotted line) and $r_{5,4}(x)$ (dashed line) for $\arctan(x)$ (solid line).

Table 1 The Padé table for $\sin(x)$.

	0	1	2	
1	x	x	$\dfrac{x}{1 + \frac{1}{6}x^2}$	$\cdots$
2	x	x	$\dfrac{x}{1 + \frac{1}{6}x^2}$	$\cdots$
3	$x - \frac{1}{6}x^3$	$x - \frac{1}{6}x^3$	$\dfrac{(-\frac{7}{60}x^3 + x)}{(1 + \frac{1}{20}x^2)}$	
	$\vdots$	$\vdots$		$\ddots$

$f(x)$ that is meromorphic in a substantial part of its domain.

2.5 Rational Interpolation

The rational equivalent of polynomial interpolation at mutually distinct interpolation points x_i consists of finding an irreducible rational function $r_{k,\ell}(x)$, of numerator degree at most k and denominator degree at most ℓ, that satisfies

$$r_{k,\ell}(x_i) = f_i, \quad i = 0, \ldots, k + \ell, \tag{8}$$

where $f_i = f(x_i)$. Instead of solving (8) one considers the linearized equations

$$(f q_{k,\ell} - p_{k,\ell})(x_i) = 0, \quad i = 0, \ldots, k + \ell, \tag{9}$$

where $p_{k,\ell}(x)$ and $q_{k,\ell}(x)$ are polynomials of respective degree k and ℓ. Condition (9) is a homogeneous linear system of $k + \ell + 1$ equations in $k + \ell + 2$ unknowns and it therefore always has a nontrivial solution. Moreover, as in the Padé approximation case, all solutions

of (9) are equivalent in the sense that they deliver the same unique irreducible rational function.

By computing the irreducible form $r_{k,\ell}(x)$ of $p_{k,\ell}(x)/q_{k,\ell}(x)$, common factors in numerator and denominator are canceled and it may well be that $r_{k,\ell}$ does not satisfy the interpolation conditions (8) anymore, despite $p_{k,\ell}$ and $q_{k,\ell}$ being solutions of (9), because one or more of the canceled factors may be of the form $x - x_i$ with x_i an interpolation point. A simple example illustrates this. Let $x_0 = 0$, $x_1 = 1$, $x_2 = 2$ with $f_0 = 0$, $f_1 = 3$, $f_2 = 3$, and take $k = 1 = \ell$. The homogeneous linear system of interpolation conditions is then

$$\begin{aligned} a_0 &= 0, \\ 3(b_0 + b_1) - (a_0 + a_1) &= 0, \\ 3(b_0 + 2b_1) - (a_0 + 2a_1) &= 0. \end{aligned}$$

A solution is given by $p_{1,1}(x) = 3x$ and $q_{1,1}(x) = x$. Hence, $r_{1,1}(x) = 3$ and clearly $r_{1,1}(x_0) \neq f_0$. The interpolation point x_0 is then called unattainable. This problem can only be fixed by increasing the degrees k and/or ℓ until the interpolation point is attainable. Note that unattainable interpolation points do not occur in polynomial interpolation ($\ell = 0$).

A well-known problem with rational interpolation and Padé approximation is the occurrence of undesirable poles in the interpolant $r_{k,\ell}(x)$. One way to avoid this is to work with preassigned poles, either explicitly or implicitly, by determining the denominator polynomial $q_{k,\ell}(x)$ a priori. All $k + \ell + 1$ interpolation conditions are then imposed on the coefficients of the numerator polynomial, and consequently the degree of the numerator is raised to $k + \ell$.

Let the interpolation points x_i be ordered such that $x_0 < x_1 < \cdots < x_n$ with $k + \ell = n$. A popular choice for the denominator polynomial that guarantees a pole-free real axis, unless the location of the poles needs to be controlled by other considerations, is

$$q_{n,n}(x) = \sum_{j=0}^{n} (-1)^j \prod_{\substack{k=0 \\ k \neq j}}^{n} (x - x_k).$$

With this choice, the rational interpolant can be written in a barycentric form similar to that in (2):

$$r_{n,n}(x) = \frac{\sum_{j=0}^{n} f_j (-1)^j/(x - x_j)}{\sum_{j=0}^{n} (-1)^j/(x - x_j)}.$$

Again, this form is very stable for interpolation. Its numerical sensitivity is measured by

$$M_n = \max_{a \leqslant x \leqslant b} \sum_{i=0}^{n} \frac{|q_{n,n}(x_i)\beta_i(x)|}{|q_{n,n}(x)|}.$$

And there is more good news now: in the case of equidistant interpolation points, M_n grows as slowly with n as the Lebesgue constant L_n in (3) for polynomial interpolation in the Chebyshev zeros. The latter makes the technique very useful in practice.

More practical choices for the denominator polynomial $q_{n,n}(x)$ are possible, guaranteeing other features, such as rapid convergence, comonotonicity, or coconvexity (coconcavity).

2.6 Sparse Interpolation

When interpolating

$$f(x) = \alpha_1 + \alpha_2 x^{100}$$

by a polynomial, the previous techniques require 101 samples of $f(x)$ to determine that $f(x)$ is itself a polynomial, while only four values need to be computed from the data points, namely the two exponents 0 and 100 and the two coefficients α_1 and α_2. So it would be nice if we could solve this polynomial reconstruction problem from only four samples.

The above is a special case of the more general *sparse interpolation*, which was studied as long ago as 1795, by Gaspard de Prony, in which the complex values ϕ_j and α_j in the interpolant

$$\phi(x) = \sum_{j=1}^{n} \alpha_j e^{\phi_j x}, \quad \alpha_j, \phi_j \in \mathbb{C}, \tag{10}$$

are to be determined from only $2n$ samples of $\phi(x)$.

While the nonlinear interpolation problems of Padé approximation and rational interpolation are solved by linearizing the conditions as in (7) and (9), the nonlinear problem of sparse interpolation is solved by separating the computation of the ϕ_j and the α_j into two linear algebra subproblems. Let $\phi(x)$ be sampled at the equidistant points $x_i = i\Delta$, $i = 0, \ldots, 2n - 1$, and let us denote $\phi(x_i)$ by f_i. We introduce the $n \times n$ Hankel matrices

$$H_n^{(r)} := \begin{pmatrix} f_r & \cdots & f_{r+n-1} \\ \vdots & \iddots & \vdots \\ f_{r+n-1} & \cdots & f_{r+2n-2} \end{pmatrix}$$

and $\lambda_j = e^{\phi_j \Delta}$, $j = 1, \ldots, n$.

The λ_j are then retrieved as the generalized eigenvalues of the problem

$$H_n^{(1)} v_j = \lambda_j H_n^{(0)} v_j, \quad j = 1, \ldots, n,$$

where the v_j are the generalized right eigenvectors. From the values λ_j, the complex numbers ϕ_j can be retrieved uniquely subject to the restriction that

$|\mathrm{Im}(\phi_j\Delta)| < \pi$. In order to satisfy this restriction, the sampling interval Δ is usually adapted to the range of the values $\mathrm{Im}(\phi_j)$.

The α_j are computed from the interpolation conditions

$$\sum_{j=1}^{n} \alpha_j e^{\phi_j x_i} = f_i, \quad i = 0, \ldots, 2n-1, \tag{11}$$

either by solving the system in the least-squares sense or by solving a subset of n consecutive interpolation conditions. Note that

$$e^{\phi_j x_i} = \lambda_j^i$$

and that the coefficient matrix of (11) is therefore a Vandermonde matrix.

With $f_i = \phi(x_i)$ we now define

$$f(x) = \sum_{j=0}^{\infty} f_j x^j,$$

where $x_i = i\Delta$, $i \geqslant 0$. Since

$$f_i = \sum_{j=1}^{n} \alpha_j e^{\phi_j x_i} = \sum_{j=1}^{n} \alpha_j \lambda_j^i,$$

we can rewrite $f(x)$ as follows:

$$f(x) = \sum_{j=1}^{n} \frac{\alpha_j}{1 - x\lambda_j}. \tag{12}$$

So we see that $f(x)$ is itself a rational function of degree $n-1$ in the numerator and n in the denominator, with poles $1/\lambda_j$. Hence, from Padé approximation theory we know (as is to be expected) that $r_{n-1,n}(x)$ reconstructs $f(x)$; in other words,

$$r_{n-1,n}(x) = f(x)$$

with

$$q(x) = \prod_{j=1}^{n} (1 - x\lambda_j).$$

The partial fraction decomposition (12) is the Laplace transform of the exponential model (10), which explains why this approach is known as the *Padé–Laplace method*.

The above connection between approximation theory and harmonic analysis is clearly not accidental. More constructions from harmonic analysis, including wavelets and Fourier series, also provide important insights into central problems in approximation theory. Other mathematical models in which the major features of a data set are represented using only a few terms are considered in the theory of COMPRESSED SENSING [**??**].

3 Least-Squares Approximation

When the quality of the data does not justify the imposition of an exact match on the approximating function, or when the quantity of the data is simply overwhelming and depicts a trend rather than very precise measurements, interpolation techniques are of no use. It is better to find a linear combination of suitable basis functions that approximates the data in some best sense. We first discuss the existence and uniqueness of a best linear approximant and the discrete linear least-squares problem. How the bestness or nearness of the approximation is measured is then explained. Different measures lead to different approximants and are to be used in different contexts. We discuss the importance of orthogonal basis functions and describe the continuous linear least-squares problem and the minimax approximation. A discussion of a connection with Fourier series and the interpolation and approximation of periodic data concludes the section.

3.1 Discrete Least-Squares Approximation

First and foremost we discuss the existence and uniqueness of a best approximant p^* from a finite-dimensional subspace P to an element f from a normed linear space V. More specifically, we ask for which of the ℓ_1-, ℓ_2-, or ℓ_∞-norms can we guarantee that either at least one or exactly one solution exists to the approximation problem of finding $p^* \in P$ such that

$$\|f - p^*\| \leqslant \|f - p\|, \quad p \in P.$$

The answer to the existence problem is affirmative for all three mentioned norms. To guarantee uniqueness of p^*, either the norm or the subspace P under consideration must satisfy additional conditions. And we must distinguish between discrete and continuous approximation and norms.

When V is *strictly convex*, in other words, when a sphere in V does not contain line segments, so that

$$\|x_1 - c\| = r = \|x_2 - c\| \quad \Rightarrow \quad \|\lambda x_1 + (1-\lambda)x_2 - c\| < r, \quad 0 < \lambda < 1,$$

then the best approximant p^* to f is unique. This applies, for instance, to the ℓ_2- or Euclidean norm, in both the discrete and continuous cases.

In the discussion of the role of P with respect to the uniqueness of p^*, we deal with the continuous case first. When a basis $\{b_0(x), \ldots, b_n(x)\}$ for P satisfies the *Haar condition*, meaning that every linear combination

$$q_n(x) = \lambda_0 b_0(x) + \cdots + \lambda_n b_n(x)$$

has at most n zeros, then the continuous best ℓ_1 and best ℓ_∞ approximation problems also have a unique solution.

Let us now look at the discrete best approximation problem in somewhat more detail. We consider a large data set of values f_i that we want to approximate by a linear combination of some linearly independent basis functions $b_j(x)$:

$$\lambda_0 b_0(x_i) + \cdots + \lambda_n b_n(x_i) = f_i, \quad i = 0, \ldots, m > n. \tag{13}$$

This $(m+1) \times (n+1)$ linear system can be written compactly as

$$\left.\begin{aligned} A\lambda = f, \quad \lambda = \begin{bmatrix} \lambda_0 \\ \vdots \\ \lambda_n \end{bmatrix}, \; f = \begin{bmatrix} f_0 \\ \vdots \\ f_m \end{bmatrix}, \\ A = (b_{j-1}(x_{i-1})) \in \mathbb{R}^{(m+1)\times(n+1)}. \end{aligned}\right\} \tag{14}$$

Unless the right-hand side f lies in the column space of A, the system cannot be solved exactly. The residual vector is given by

$$r = f - A\lambda \in \mathbb{R}^{m+1},$$

and the solution λ we are looking for is the one that solves the system best, in other words, the system that makes the magnitude (or norm) of the residual vector minimal. The least-squares problem corresponds to using the Euclidean norm or the ℓ_2-norm $\|r\|_2 = (r_1^2 + \cdots + r_m^2)^{1/2}$ to measure the residual vector, and the optimization problem translates to

$$(A^{\mathrm{T}}A)\lambda = A^{\mathrm{T}}f,$$

which is a square linear system of equations called the NORMAL EQUATIONS [??]. If the matrix A of the overdetermined linear system has maximal column rank, then the matrix $A^{\mathrm{T}}A$ is nonsingular and the solution is unique.

When every $(n+1)\times(n+1)$ submatrix of the matrix A in (14) is nonsingular, then the discrete best ℓ_∞ approximation problem has a unique solution as well. An example showing the lack of uniqueness of the best ℓ_1 approximation under the same condition is easy to find. Take

$$A = \begin{pmatrix} 1 \\ -1 \end{pmatrix}, \qquad f = \begin{pmatrix} 1 \\ 2 \end{pmatrix}$$

in (14). Then the minimum of $\|A\lambda - f\|_1$ with $n = 0$ and $b_0(x) = 1$ is the same for all $-2 \leqslant \lambda_0 \leqslant 1$.

In practice, instead of solving the normal equations, more numerically stable techniques based on ORTHOGONAL TRANSFORMATIONS [??] are applied directly to the overdetermined system (14). These transformations do not alter the Euclidean norm of the residual vector r and hence have no impact on the optimization criterion.

Let us now see whether the Euclidean norm is the correct norm to use.

3.2 Choice of Norm

If the optimal solution to the overdetermined linear system is the one that makes the norm $\|r\|$ of the residual minimal, then we must decide which norm to use to measure r. Although norms are in a sense equivalent, because they differ only by a scalar multiple depending only on the dimension, it makes quite a difference whether we minimize $\|r\|_1$, $\|r\|_2$, or $\|r\|_\infty$. Let us perform the following experiment.

Using a Gaussian random number generator with mean μ and standard deviation σ, we generate $m+1$ numbers f_i. The approximation problem we consider is the computation of an estimate for μ from the data points f_i, where σ expresses a tight or loose spread around μ. Compare this with a real-life situation where the data f_i are collected by performing some measurements of a magnitude μ, and σ represents the accuracy of the measuring tool used to obtain the f_i.

In the notation of (13), we want to fit the f_i by a multiple of the basis function $b_0(x) = 1$ because we are looking for the constant μ. The overdetermined linear system takes the form

$$\lambda_0 \cdot 1 = f_i, \quad i = 0, \ldots, m.$$

It is clear that this linear system does not have an exact solution. The residual vector is definitely nonzero. We shall see that different criteria or norms can be used to express the closeness of the estimate λ_0 for μ to the data points f_i or, in other words, the magnitude of the residual vector r with components $f_i - \lambda_0$, and that the standard deviation σ will also play a role.

If the Euclidean norm is used, then the optimal estimate $\lambda_0^{(2)}$ is the mean of the m measurements f_i:

$$\lambda_0^{(2)} = \frac{1}{m+1} \sum_{i=0}^{m} f_i.$$

If we choose the ℓ_1-norm $\|r\|_1 = \sum_{i=1}^{m} |r_i|$ as a way to measure distances, then the value $\lambda_0^{(1)}$ that renders the ℓ_1-norm of the residual vector minimal is the median of the values f_i. Any change that makes the larger values extremely large or the smaller values extremely small therefore has no impact on $\lambda_0^{(1)}$, which is rather insensitive to outliers.

When choosing as distance function the ℓ_∞-norm $\|r\|_\infty = \max_{i=1,\ldots,m} |r_i|$, the optimal solution $\lambda_0^{(\infty)}$ to the problem is given by

$$\lambda_0^{(\infty)} = \tfrac{1}{2}\Big(\min_{i=0,\ldots,m} f_i + \max_{i=0,\ldots,m} f_i \Big).$$

This can also be understood intuitively. The value for λ_1 that makes $\|r\|_\infty$ minimal is the one that makes the largest deviation minimal, so it should be right in the middle between the extremes.

So the ℓ_∞-norm criterion performs particularly well in the context of rather accurate data (in this experiment meaning small standard deviation σ) that suffer relatively small input errors (such as roundoff errors). When outliers or additional errors (such as from manual data input) are suspected, use of the ℓ_1-norm is recommended. If the measurement errors are believed to be normally distributed with mean zero, then the ℓ_2-norm is the usual choice. Approximation problems of this type are therefore called least-squares problems.

3.3 Orthogonal Basis Functions

In the same way that we prefer to draw a graph using an orthogonal set of axes (the smaller the angle between the axes, the more difficult it becomes to make a clear drawing), it is preferred to use a so-called orthogonal set of basis functions $b_j(x)$ in (13). Orthogonal basis functions $b_j(x)$ can tremendously improve the conditioning or sensitivity of the problem (14). They are also useful in continuous least-squares problems.

The notion of orthogonality in a function space parallels that of orthogonality in the vector space $\mathbb{R}^k$: for a positive weight function $w(x)$ defined on the interval $[a,b]$, we say that the functions f and g are w-orthogonal if

$$\langle f, g \rangle_w = \int_a^b f(x) g(x) w(x)\, \mathrm{d}x = 0.$$

The function $w(x)$ can assign a larger weight to certain parts of the interval $[a,b]$. For instance, the function $w(x) = 1/\sqrt{1-x^2}$ on $[-1,1]$ assigns more weight toward the endpoints of the interval.

For $w(x) = 1$ and $[a,b] = [-1,1]$, a sequence of orthogonal polynomials $L_i(x)$ satisfying

$$\int_{-1}^{1} L_j(x) L_k(x)\, \mathrm{d}x = 0, \quad j \neq k,$$

is given by

$$L_0(x) = 1, \qquad L_1(x) = x,$$
$$L_{i+1}(x) = \frac{2i+1}{i+1} x L_i(x) - \frac{i}{i+1} L_{i-1}(x), \quad i \geqslant 1.$$

The polynomials $L_i(x)$ are called the *Legendre polynomials*. For $w(x) = 1/\sqrt{1-x^2}$ and $[a,b] = [-1,1]$, a sequence of orthogonal polynomials $T_i(x)$ satisfying

$$\int_{-1}^{1} T_j(x) T_k(x) \frac{1}{\sqrt{1-x^2}}\, \mathrm{d}x = 0, \quad j \neq k,$$

is given by

$$T_0(x) = 1, \qquad T_1(x) = x,$$
$$T_{i+1}(x) = 2x T_i(x) - T_{i-1}(x), \quad i \geqslant 1.$$

The polynomials $T_i(x)$ are called the *Chebyshev polynomials* (of the first kind). They are also very useful in (continuous as well as discrete) least-squares problems, as discussed below.

When the polynomials are to be used on an interval $[a,b]$ different from $[-1,1]$, the simple change of variable

$$x \to \frac{2}{b-a}\left(x - \frac{a+b}{2} \right)$$

transforms the interval $[a,b]$ to the interval $[-1,1]$, on which the orthogonal polynomials are defined.

Orthogonal polynomials also satisfy the Haar condition, so every linear combination

$$q_n(x) = a_0 p_0(x) + \cdots + a_n p_n(x)$$

of the orthogonal polynomials $p_i(x)$ of degree $i = 0,\ldots,n$ has at most n zeros. Therefore, orthogonal polynomials are also a suitable basis in which to express an interpolating function: the system of interpolation conditions

$$\sum_{j=0}^{n} a_j p_j(x_i) = f_i, \quad i = 0,\ldots,n,$$

has a coefficient matrix that is guaranteed to be nonsingular for mutually distinct points x_i.

The importance of orthogonal basis functions in interpolation and approximation cannot be overstated. Problems become numerically better conditioned and formulas simplify. For instance, the Chebyshev polynomials $T_i(x)$ also satisfy the discrete orthogonality

$$\sum_{i=0}^{n} T_j(x_i) T_k(x_i) = (1+\delta_{k0}) \frac{n+1}{2} \delta_{jk}, \qquad j,k = 0,1,\ldots,$$

where δ_{ij} is the KRONECKER DELTA [**??**] and

$$x_i = \cos\left(\frac{(2i+1)\pi}{2(n+1)} \right)$$

are the zeros of the Chebyshev polynomial T_{n+1}. When expressing the polynomial interpolant $p_n(x)$ in (1) of

degree n in the Chebyshev basis,

$$p_n(x) = \sum_{j=0}^{n} a_j T_j(x),$$

an easy explicit formula for $p_n(x)$ interpolating the values f_i at the points x_i can be given:

$$p_n(x) = \frac{\sum_{i=0}^{n} f_i}{n+1} + \sum_{j=1}^{n} \Big(\frac{2\sum_{i=0}^{n} f_i T_j(x_i)}{n+1} \Big) T_j(x).$$

Another elegant explicit formula, based on the continuous orthogonality property of the Chebyshev polynomials, is given in (16).

The practical utility of Chebyshev polynomials is illustrated by the open source Chebfun software system (www.chebfun.org) for numerical computation with functions, which is built on piecewise-polynomial interpolation at the extrema of Chebyshev polynomials, or what is equivalent (via the FAST FOURIER TRANSFORM [**??**]), expansions in Chebyshev polynomials.

3.4 Chebyshev Series

Let us now choose the basis functions $b_j(x) = T_j(x)$ and look for the coefficients λ_j that make the ℓ_2-norm of

$$f(x) - \sum_{j=0}^{n} \lambda_j T_j(x), \quad -1 \leqslant x \leqslant 1,$$

minimal for $f(x)$ defined on $[-1,1]$, for simplicity. We are looking for the polynomial $p_n(x)$ of degree n that is closest to $f(x)$, where we measure the distance between the functions using the inner product

$$\begin{aligned} \|f - p_n\|_2^2 &= \langle f - p_n, f - p_n \rangle \\ &= \int_{-1}^{1} \frac{(f-p_n)^2(x)}{\sqrt{1-x^2}}\, \mathrm{d}x. \end{aligned} \tag{15}$$

This is a continuous least-squares problem because the norm of a function is minimized instead of the norm of a finite-dimensional vector. Since

$$\begin{aligned} \Big\| f - \sum_{j=0}^{n} \lambda_j T_j \Big\|_2^2 &= \Big\langle f - \sum_{j=0}^{n} \lambda_j T_j, f - \sum_{j=0}^{n} \lambda_j T_j \Big\rangle \\ &= \|f\|_2^2 - \sum_{j=0}^{n} \langle f, T_j/\|T_j\|_2 \rangle^2 \\ &\quad + \sum_{j=0}^{n} (\langle f, T_j/\|T_j\|_2 \rangle - \lambda_j \|T_j\|_2)^2, \end{aligned}$$

in which only the last sum of squares depends on λ_j, the minimum is attained for the so-called Chebyshev coefficients

$$\lambda_j = \langle f, T_j \rangle / \langle T_j, T_j \rangle. \tag{16}$$

The partial sum of degree n of the Chebyshev series development of a function,

$$f(x) = \sum_{j=0}^{\infty} \frac{\langle f, T_j \rangle}{\langle T_j, T_j \rangle} T_j(x),$$

is therefore the best polynomial approximation of degree n to $f(x)$ in the ℓ_2 sense. Since

$$\Big| f(x) - \sum_{j=0}^{n} \frac{\langle f, T_j \rangle}{\langle T_j, T_j \rangle} T_j(x) \Big| \leqslant \sum_{j=n+1}^{\infty} \Big| \frac{\langle f, T_j \rangle}{\langle T_j, T_j \rangle} \Big|,$$

this error can be made arbitrarily small when the series of Chebyshev coefficients converges absolutely, a condition which is automatically satisfied for functions that are continuously differentiable in $[-1,1]$.

The above technique can be generalized to any weight function and its associated family of orthogonal polynomials: when switching the weight function, the norm criterion (15) changes and the orthogonal basis is changed.

The Chebyshev series partial sums are good overall approximations to a function $f(x)$ defined on the interval $[-1,1]$ (or $[a,b]$ after a suitable change of variable). To illustrate this, in figure 6 we compare, for the function $f(x) = \arctan(x)$, the error plots of the Chebyshev series partial sum of degree 9 with the Taylor series partial sum of the same degree. Its Chebyshev series and Taylor series developments are, respectively, given by

$$\begin{aligned} \arctan(x) &= \sum_{i=0}^{\infty} (-1)^i \frac{2(\sqrt{2}-1)^{2i+1}}{2i+1} T_{2i+1}(x), \\ \arctan(x) &= \sum_{i=0}^{\infty} (-1)^i \frac{1}{2i+1} x^{2i+1}. \end{aligned}$$

Although explicit formulas for the Taylor series expansion of most elementary and special functions are known, the same is not true for Chebyshev series expansions. For most functions, the coefficients (16) have to be computed numerically because no analytic expression for (16) can be given.

3.5 The Minimax Approximation

Instead of minimizing the ℓ_2-distance (15) between a function $f(x) \in C([a,b])$ and a polynomial model for $f(x)$, we can consider the problem of minimizing the ℓ_∞-distance. Every continuous function $f(x)$ defined on a closed interval $[a,b]$ has a unique so-called *minimax polynomial approximant* of degree n. This means that there exists a unique polynomial $p_n = p_n^*$ of degree $\partial p_n \leqslant n$ that minimizes

$$\|f - p_n\|_\infty = \max_{x \in [a,b]} \Big| f(x) - \sum_{j=0}^{n} \lambda_j x^j \Big|. \tag{17}$$

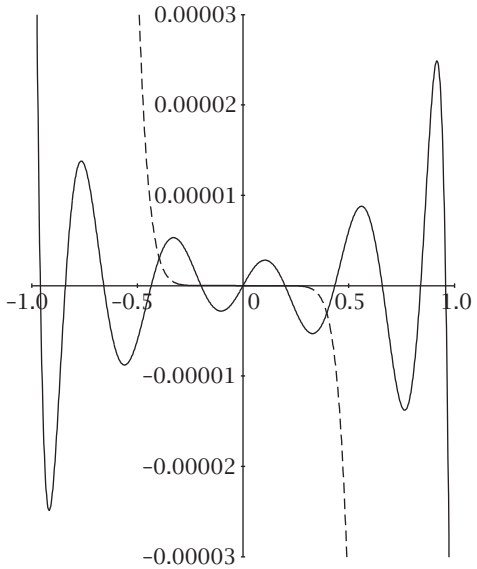

Figure 6 Error plots of Chebyshev (solid line) and Taylor (dashed line) partial sums of degree 9 for $\arctan(x)$.

More generally, if the set of basis functions $\{b_0(x), \ldots, b_n(x)\}$ satisfies the Haar condition, then there exists a unique approximant

$$q_n^*(x) = \lambda_0^* b_0(x) + \cdots + \lambda_n^* b_n(x)$$

that minimizes

$$\|f - q_n\|_\infty = \max_{x\in[a,b]} \left| f(x) - \sum_{j=0}^{n} \lambda_j b_j(x) \right|.$$

The minimum is attained and is not an infimum. When $b_j(x) = x^j$, the polynomial $p_n^*(x)$ is computed using the Remes algorithm, which is based on its characterization given by the alternation property of the function $(f - p_n^*)(x)$: p_n^* is the best polynomial approximant of degree n if the error $\|f - p_n^*\|_\infty$ is attained by the function $f - p_n^*$ in at least $n + 2$ points $y_0, \ldots, y_{n+1}$ in the interval $[a, b]$ and this with alternating sign, meaning that

$$\exists y_0 > y_1 > \cdots > y_{n+1} \in [a,b]:$$
$$(f - p_n^*)(y_i) = s(-1)^i \|f - p_n^*\|_\infty,$$
$$s = \pm 1,\ i = 0, \ldots, n+1.$$

The Remes algorithm is an iterative procedure and the polynomial $p_n^*(x)$ is obtained as the limit. The above characterization is also called the *equioscillation property*. We illustrate it in figure 7, where we plot the error $e^x - p_3^*(x)$ on $[-1, 1]$. Compare this figure with figures 8 and 9, in which the error oscillates but does not equioscillate.

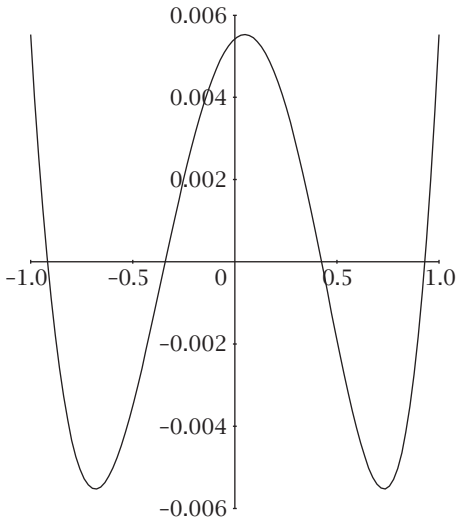

Figure 7 Error plot of $e^x - p_3^*(x)$.

How much better the (nonlinear) minimax approximation is, compared with a linear approximation procedure of degree n such as polynomial interpolation, Chebyshev approximation, and the like, is expressed by the norm $\|P_n\|_\infty = \sup_{\|f\|_\infty \leqslant 1} \|P_n(f)\|_\infty$ of the linear operator P_n that associates with a function its particular linear approximant. Since $P_n(p_n^*) = p_n^*$, we have

$$\begin{aligned} \|f - P_n(f)\|_\infty &= \|f - p_n^* + p_n^* - P_n(f)\|_\infty \\ &= \|f - p_n^* + P_n(p_n^* - f)\|_\infty \\ &\leqslant (1 + \|P_n\|_\infty)\|f - p_n^*\|_\infty. \end{aligned}$$

The value $\|P_n\|_\infty$ is called the *Lebesgue constant*. When dealing with polynomial interpolation, $\|P_n\|_\infty = L_n$, with L_n given by (3).

The quality of the continuous best ℓ_2 approximant (such as that in figure 8) is expressed by

$$\|P_n\|_\infty = \frac{1}{\pi} \int_0^\pi \left| \frac{\sin((n + \frac{1}{2})\theta)}{\sin(\frac{1}{2}\theta)} \right|,$$

which again grows only logarithmically with n. Continuous ℓ_2 polynomial approximation and Lagrange interpolation in the Chebyshev zeros (such as in figure 9) can therefore be considered near-best polynomial approximants.

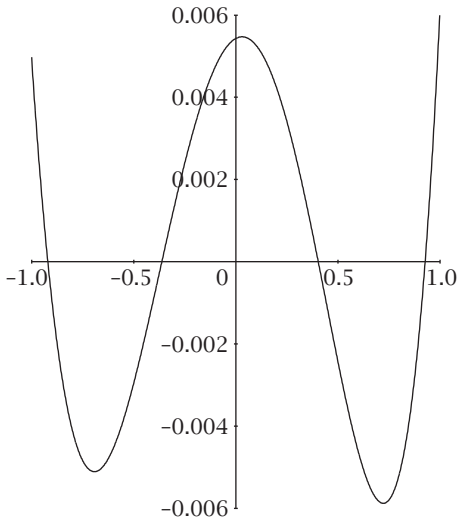

Figure 8 Error plot of the Chebyshev partial sum of degree 3 for e^x.

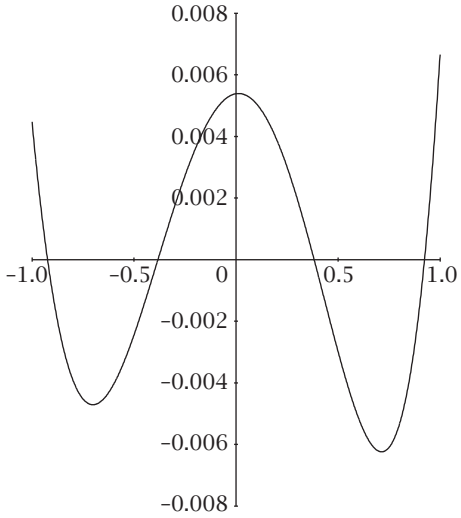

Figure 9 Error plot of the polynomial interpolant of degree 3 in the Chebyshev zeros for e^x.

3.6 Fourier Series

Let us return to a discrete approximation problem. Our interest is now in data exhibiting some periodic behavior, such as the description of rotation-invariant geometric figures or the sampling of a sound waveform. A suitable set of orthogonal basis functions is the set

$$\left.\begin{array}{c} 1,\ \cos(x),\ \cos(2x),\ldots,\ \cos(nx), \\ \sin(x),\ \sin(2x),\ \ldots,\ \sin(nx),\end{array}\right\} \tag{18}$$

as long as the distinct datapoints x_i with $i = 0,\ldots,m$ are evenly spaced on an interval of length 2π, because then, for any two basis functions $b_j(x)$ and $b_k(x)$ from (18), we have

$$\langle b_j, b_k\rangle = \sum_{i=0}^{m} b_j(x_i)b_k(x_i) = \frac{m+1}{2}\delta_{jk}, \quad j \neq 0,$$
$$\langle b_0, b_k\rangle = \sum_{i=0}^{m} b_0(x_i)b_k(x_i) = (m+1)\delta_{0k},$$

where δ_{ij} is the Kronecker delta. For simplicity we assume that the real data $f_0,\ldots,f_m$ are given on $[0, 2\pi)$ at

$$x_0 = 0, \quad x_1 = \frac{2\pi}{m}, \quad x_2 = \frac{4\pi}{m}, \quad \ldots, \quad x_m = \frac{2m\pi}{m+1}.$$

Because of the periodicity, the value at $x_{m+1} = 2\pi$ equals the value at x_0 and therefore it is not repeated. Let $m \geqslant 2n$ and consider the approximation

$$\frac{\lambda_0}{2} + \sum_{j=1}^{n} \lambda_{2j}\cos(jx) + \sum_{j=1}^{n} \lambda_{2j-1}\sin(jx).$$

The values

$$\lambda_{2j} = \frac{2}{m+1}\sum_{i=0}^{m} f_i \cos(jx_i), \quad j = 0,\ldots,n,$$
$$\lambda_{2j-1} = \frac{2}{m+1}\sum_{i=0}^{m} f_i \sin(jx_i), \quad j = 1,\ldots,n,$$

minimize the ℓ_2-norm

$$\sum_{i=1}^{m}\left(\frac{\lambda_0}{2} + \sum_{j=1}^{n}\lambda_{2j}\cos(jx_i) + \sum_{j=1}^{n}\lambda_{2j-1}\sin(jx_i) - f_i\right)^2.$$

When $m = 2n$, the minimum is zero because the least-squares approximant becomes a trigonometric interpolant. Note that we have replaced λ_0 by $\frac{1}{2}\lambda_0$ because $\langle b_j, b_j\rangle$ is smaller by a factor of 2 when $j = 0$.

If we form a single complex quantity $\Lambda_j = \lambda_{2j} - \mathrm{i}\lambda_{2j-1}$ for $j = 1,\ldots,n$, where $\mathrm{i} = \sqrt{-1}$, these summations can be computed using a discrete Fourier transform that maps the data f_i at the points x_i to the Λ_j:

$$\Lambda_j = \frac{1}{m+1}\sum_{i=0}^{m} f_i e^{-\mathrm{i}2\pi ij/(m+1)}, \quad j = 0,\ldots,n.$$

The functions in (18) also satisfy a continuous orthogonality property:

$$\int_0^{2\pi} \cos(jx)\cos(kx)\,\mathrm{d}x = \pi\delta_{jk}, \quad j \neq 0,$$
$$\int_0^{2\pi} \cos(jx)\cos(kx)\,\mathrm{d}x = \frac{\pi}{2}\delta_{jk}, \quad j = 0,$$
$$\int_0^{2\pi} \cos(jx)\sin(kx)\,\mathrm{d}x = 0,$$
$$\int_0^{2\pi} \sin(jx)\sin(kx)\,\mathrm{d}x = \pi\delta_{jk}.$$

They can therefore be used for a Fourier series representation of a function $f(x)$:

$$f(x) = \frac{\lambda_0}{2} + \sum_{j=1}^{\infty} \lambda_{2j}\cos(jx) + \sum_{j=1}^{\infty} \lambda_{2j-1}\sin(jx),$$
$$\lambda_{2j} = \frac{\langle f(x), \cos(jx)\rangle}{\pi}, \quad j = 0, 1, \ldots,$$
$$\lambda_{2j-1} = \frac{\langle f(x), \sin(jx)\rangle}{\pi}, \quad j = 1, 2, \ldots.$$

The partial sum of trigonometric degree n of this series minimizes the ℓ_2-norm

$$\int_0^{2\pi} \Big(f(x) - \frac{\lambda_0}{2} - \sum_{j=1}^{n} \lambda_{2j}\cos(jx) - \sum_{j=1}^{n} \lambda_{2j-1}\sin(jx) \Big)^2 \mathrm{d}x.$$

4 Multivariate Interpolation and Approximation

The approximation of multivariate functions—continuous ones as well as discontinuous ones—is an active field of research due to its large variety of applications in the computational sciences and engineering. A wide range of multivariate generalizations of the above interpolation and approximation problems to functions of several variables $x, y, z, \ldots$ have therefore been developed: polynomial and rational ones, discrete and continuous ones.

A fundamental issue in multivariate interpolation and approximation is the so-called *curse of dimensionality*, meaning that when the dimensionality increases, the number of different combinations of variables grows exponentially in the dimensionality. A polynomial of degree 3 in eight variables already has 165 terms! Another problem is that the polynomial basis of the multinomials does not satisfy the Haar condition and there is no easy generalization of this property to the multivariate case.

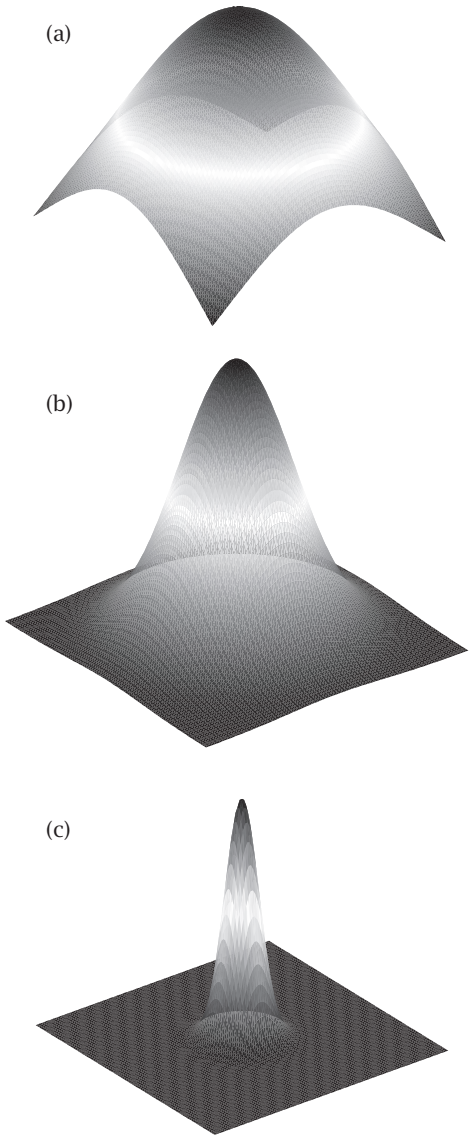

Figure 10 Gaussian basis function with (a) $s = 0.4$, (b) $s = 1$, and (c) $s = 3$.

In order to counter both problems, the theory of radial basis functions has been developed.

Let us consider data f_i given at corresponding multidimensional vectors x_i, $i = 0, \ldots, n$. The data vectors x_i do not have to form a grid but can be scattered. A *radial basis function* is a function whose value depends only on the distance from the origin or from another point, so its variable is $r = \|x\|_2$ or $r = \|x - c\|_2$. When centering a radial basis function $B(r)$ at each data

point, there is a basis function $B(\|x-x_j\|_2)$ for each $j = 0,\ldots,n$. The coefficients a_j in a radial basis function interpolant $\sum_{j=0}^{n} a_j B(\|x-x_j\|_2)$ are then computed from the linear system

$$\sum_{j=0}^{n} a_j B(\|x_i - x_j\|_2) = f_i, \quad i = 0,\ldots,n.$$

Several commonly used types of radial basis functions $B(r)$ guarantee nonsingular systems of interpolation conditions, in other words, nonsingular matrices

$$\left(B(\|x_i - x_j\|_2)\right)_{i,j=0,\ldots,n}.$$

We mention as examples the Gaussian, multiquadric, inverse multiquadric, and a member of the Matern family, respectively, given by

$$B(r) = \mathrm{e}^{-(sr)^2},$$
$$B(r) = \sqrt{1+(sr)^2},$$
$$B(r) = 1/\sqrt{1+(sr)^2},$$
$$B(r) = (1+sr)\mathrm{e}^{-sr}.$$

The real parameter s is called a shape parameter. As can be seen in figure 10, different choices for s greatly influence the shape of $B(r)$. Smaller shape parameters correspond to a flatter or wider basis function. The choice of s has a significant impact on the accuracy of the approximation, and finding an optimal shape parameter is not an easy problem. Another concern is the numerical conditioning of the radial basis interpolation problem, especially when the shape parameter is small. The user often has to find the right trade-off between accuracy and conditioning.

The concept of radial basis function also allows one to work mesh free. In a mesh or grid of data points each point has a fixed number of neighbors and this connectivity between neighbors is used to define mathematical operators such as the derivative and the divided difference. Multivariate mesh-free methods allow one to generalize these concepts and are especially useful when the mesh is difficult to maintain (e.g., in high-dimensional problems, when there is nonlinear behavior, discontinuities, singularities, etc.).

5 Future Research

Especially in multivariate approximation theory, many research questions remain unsolved: theory for the multivariate case is not nearly as well understood as it is for the univariate case. But researchers continue to push the boundaries in the one-variable case as well: what is the largest function class or the most general domain for which a result holds, for example? Many papers can be found on Jackson-type inequalities (approximation error bounds in terms of the function's smoothness), Bernstein-type inequalities (bounds on derivatives of polynomials), and convergence properties of particular approximations (polynomial, spline, rational, trigonometric), to name just a few fundamental topics. The development of orthogonal basis functions, on disconnected regions or in more variables, also deserves (and is getting) a lot of attention.

Further Reading

Baker Jr, G., and P. Graves-Morris. 1996. *Padé Approximants*, 2nd edn. Encyclopedia of Mathematics and Its Applications, volume 59. Cambridge: Cambridge University Press.

Braess, D. 1986. *Nonlinear Approximation Theory*. Berlin: Springer.

Buhmann, M. D. 2003. *Radial Basis Functions: Theory and Implementations*. Cambridge: Cambridge University Press.

Cheney, E. 1982. *Introduction to Approximation Theory*. New York: McGraw-Hill.

Chihara, T. 2011. *An Introduction to Orthogonal Polynomials*. New York: Dover.

de Boor, C. 2001. *A Practical Guide to Splines*, revised edn. New York: Springer.

DeVore, R. A. 1998. Nonlinear approximation. *Acta Numerica* 7:51–150.

Powell, M. 1981. *Approximation Theory and Methods*. Cambridge: Cambridge University Press.

Szabados, J., and P. Vértesi. 1990. *Interpolation of Functions*. Teaneck, NJ: World Scientific.

Trefethen, L. N. 2013. *Approximation Theory and Approximation Practice*. Philadephia, PA: SIAM.