

18.330 Spring 2012
Introduction to numerical analysis
Class notes

Laurent Demanet
Draft February 14, 2012

Preface

Good additional references are

- Burden and Faires, Introduction to numerical analysis
- Suli and Mayers, Introduction to numerical analysis
- Trefethen, Spectral methods in Matlab

as well as the pdf notes on numerical analysis by John Neu.

If the text points you to some external reference for further study, the latter is obviously not part of the material for the class.

We will very likely skip several sections in chapters 7 and 8 for 18.330.

Work in progress!

Chapter 1

Series and sequences

Throughout these notes we'll keep running into Taylor series and Fourier series. It's important to understand what is meant by convergence of series before getting to numerical analysis proper. These notes are self-contained, but two good extra references for this chapter are Tao, Analysis I; and Dahlquist and Björck, Numerical methods.

A sequence is a possibly infinite collection of numbers lined up in some order:

$$a_1, a_2, a_3, \dots$$

A series is a possibly infinite sum:

$$a_1 + a_2 + a_3 + \dots$$

We'll consider real numbers for the time being, but the generalization to complex numbers is a nice exercise which mostly consists in replacing each occurrence of an absolute value by a modulus.

The questions we address in this chapter are:

- What is the meaning of an infinite sum?
- Is this meaning ever ambiguous?
- How can we show convergence vs. divergence?
- When can we use the usual rules for finite sums in the infinite case?

1.1 Convergence vs. divergence

We view infinite sums as limits of partial sums. Since partial sums are sequences, let us first review convergence of sequences.

Definition 1. A sequence $(a_j)_{j=0}^{\infty}$ is said to be ϵ -close to a number b if there exists a number $N \geq 0$ (it can be very large), such that for all $n \geq N$,

$$|a_j - b| \leq \epsilon.$$

A sequence $(a_j)_{j=0}^{\infty}$ is said to converge to b if it is ϵ -close to b for all $\epsilon > 0$ (however small). We then write $a_j \rightarrow b$, or $\lim_{j \rightarrow \infty} a_j = b$.

If a sequence does not converge we say that it diverges. Unbounded sequences, i.e., sequences that contain arbitrarily large numbers, always diverge. (So we never say “converges to infinity”, although it’s fine to say “diverges to infinity”.)

Examples:

- $e^{-n} \rightarrow 0$ as $n \rightarrow \infty$, and convergence is very fast.
- $n/(n+2) \rightarrow 1$ as $n \rightarrow \infty$, and convergence is rather slow.
- $(-1)^n$ is bounded, but does not converge.
- $\log(n) \rightarrow \infty$ as $n \rightarrow \infty$, so the sequence diverges. For a proof that $\log(n)$ takes on arbitrarily large values, fix any large integer m . Does there exist an n such that $\log(n) \geq m$? Yes, it suffices to take $n \geq e^m$.

Definition 2. Consider a sequence $(a_j)_{j=0}^{\infty}$. We define the N -th partial sum S_N as

$$S_N = a_0 + a_1 + \dots + a_N = \sum_{j=0}^N a_j.$$

We say that the series $\sum_j a_j$ converges if the sequence of partial sums S_N converges to some number b as $N \rightarrow \infty$. We then write $\sum_{j=0}^{\infty} a_j = b$.

Again, if a series does not converge we say that it diverges.

Example 1. Consider $\sum_{j=0}^{\infty} 2^{-j}$, i.e.,

$$1 + \frac{1}{2} + \frac{1}{4} + \frac{1}{8} + \dots$$

This series converges to the limit 2. To prove this, consider the partial sum

$$S_N = \sum_{j=0}^N 2^{-j}.$$

Let us show by induction that $S_N = 2 - 2^{-N}$. The base case $N = 0$ is true since $2^{-0} = 2 - 2^{-0}$. For the induction case, assume $S_N = 2 - 2^{-N}$. We then write

$$S_{N+1} = S_N + 2^{-(N+1)} = (2 - 2^{-N}) + 2^{-(N+1)} = 2 - 2^{-(N+1)},$$

the desired conclusion.

Example 2. The previous example was the $x = 1/2$ special case of the so-called geometric series

$$1 + x + x^2 + x^3 + \dots$$

With a similar argument, we obtain the limit as

$$\sum_{j=0}^{\infty} x^j = \frac{1}{1-x},$$

provided the condition $|x| < 1$ holds. This expression can also be seen as the Taylor expansion of $1/(1-x)$, centered at zero, and with radius of convergence 1.

Example 3. Consider the so-called harmonic series

$$1 + \frac{1}{2} + \frac{1}{3} + \frac{1}{4} + \dots$$

This series diverges. To see this, let us show that the N partial sum is comparable to $\log(N)$. We use the integral test

$$S_N = \sum_{j=1}^N \frac{1}{j} \geq \int_1^{N+1} \frac{1}{x} dx.$$

(Insert picture here)

The latter integral is $\log(N+1)$, which diverges as a sequence. The partial sums, which are larger, must therefore also diverge.

Example 4. Consider

$$\sum_{j=1}^{\infty} \frac{1}{n^q},$$

for some $q > 0$. As a function of q , this is the Riemann zeta function $\zeta(q)$. (A fascinating object for number theorists.)

We've seen above that $q = 1$ leads to divergence. A similar integral test would show that the series converges when $q > 1$, while it diverges when $q \leq 1$.

We now switch to a finer understanding of convergence: certain series are absolutely convergent, while others are conditionally convergent. This will affect what type of algebraic manipulations can be done on them.

Definition 3. A series $\sum_{j=0}^{\infty} a_j$ is said to be absolutely convergent if $\sum_{j=0}^{\infty} |a_j|$ converges. If a series is not absolutely convergent, but nevertheless converges, we say that it is conditionally convergent.

The subtlety with conditional convergence is that alternating plus and minus signs may lead to convergence because of cancelations when summing consecutive terms.

Example 5. Consider

$$1 - \frac{1}{2} + \frac{1}{3} - \frac{1}{4} + \frac{1}{5} - \frac{1}{6} + \dots$$

This series is not absolutely convergent, because it reduces to the harmonic series when absolute values of the terms are taken. It is nevertheless convergent, hence conditionally convergent, as the following argument shows. Assume N is even, and let

$$S_N = \sum_{j=1}^N \frac{(-1)^j}{j}.$$

Terms can be grouped 2-by-2 to obtain, for $j \geq 1$,

$$\frac{1}{j} - \frac{1}{j+1} = \frac{1}{j(j+1)}.$$

A fortiori, $\frac{1}{j(j+1)} \leq \frac{1}{j^2}$, so

$$S_N \leq \sum_{j=1,3,5,\dots}^{N-1} \frac{1}{j^2},$$

which we know converges. If on the other hand N is odd, then $S_N = S_{N-1} + \frac{1}{N+1}$. Both terms S_N and $1/(N+1)$ are converging sequences, so their sum converges as well. This proves convergence.

Note that the series converges to

$$1 - \frac{1}{2} + \frac{1}{3} - \frac{1}{4} + \frac{1}{5} - \frac{1}{6} + \dots = \log(2).$$

This is the special case $x = 1$ in the Taylor expansion of $\log(1+x)$ about $x = 0$.

In passing, without proof, here is a useful test to check convergence of alternating series.

Theorem 1. (Alternating series test) Consider the series

$$\sum_{j=0}^{\infty} (-1)^j a_j,$$

where $a_j > 0$. If (a_j) converges to zero (as a sequence), then the series is convergent.

The main problem with conditionally convergent series is that if the terms are rearranged, then the series may converge to a different limit. The “safe zone” for handling infinite sums as if they were finite is when convergence is absolute.

Theorem 2. Let $f : \mathbb{Z}^+ \mapsto \mathbb{Z}^+$ be a bijection, i.e., f is a rearrangement of the nonnegative integers. Consider a series $\sum_{j=0}^{\infty} a_j$. If this series is absolutely convergent, then

$$\sum_{j=0}^{\infty} a_j = \sum_{j=0}^{\infty} a_{f(j)}.$$

Here is what usually happens when the assumption of absolute convergence is not satisfied.

Example 6. *Consider again*

$$\frac{1}{3} - \frac{1}{4} + \frac{1}{5} - \frac{1}{6} + \dots$$

which as we have seen equals $\log(2) - (1 - 1/2) = \log(2) - 1/2 = .193147\dots$. We can rearrange the terms of the series by assigning two negative terms for each positive term:

$$\frac{1}{3} - \frac{1}{4} - \frac{1}{6} + \frac{1}{5} - \frac{1}{8} - \frac{1}{10} + \frac{1}{7} + \dots$$

This series is also convergent, but happens to converge to $(\log(2) - 1)/2 = -.153426\dots$.

Other operations that can be safely performed on absolutely convergent series are passing absolute values inside the sum, and exchanging sums. Again, complications arise if the series is only conditionally convergent. (See Tao, Analysis I, for counter-examples.)

Theorem 3. *The following operations are legitimate for absolutely convergent series.*

- *Passing absolute values inside sums:*

$$\left| \sum_{j=0}^{\infty} a_j \right| \leq \sum_{j=0}^{\infty} |a_j|.$$

- *Swapping sums:*

$$\sum_{j=0}^{\infty} \sum_{k=0}^{\infty} a_{j,k} = \sum_{k=0}^{\infty} \sum_{j=0}^{\infty} a_{j,k}$$

Note in passing that the same is true for integrals of unbounded integrands or integrals over unbounded domains: they need to be absolutely convergent (integrability of the absolute value of the function) for the integral swap to be legitimate. This is the content of Fubini's theorem. Again, there are striking counter-examples when the integrals are not absolutely convergent and the swap doesn't work (See Tao, Analysis I).

1.2 The big-O notation

Here are a few useful pieces of notation for comparing growth or decay of sequences, used extensively by numerical analysts. They are called the big-O, little-o, and big-Theta notations. Big-O is used much more often than the other two. They occur when comparing decay rates of truncation errors, and runtimes of algorithms.

Definition 4. Consider two nonzero sequences f_n and g_n for $n = 0, 1, 2, \dots$

We write $f_n = O(g_n)$ when there exists $C > 0$ such that $|f_n| \leq C|g_n|$.

We write $f_n = o(g_n)$ when $f_n/g_n \rightarrow 0$ as $n \rightarrow \infty$.

We write $f_n = \Theta(g_n)$ when $f_n = O(g_n)$ and $g_n = O(f_n)$.

Examples:

- $f_n = n^2$ and $g_n = n^3$: we have $n^2 = O(n^3)$ and $n^2 = o(n^3)$ but $n^2 \neq \Theta(n^3)$.
- $f_n = \frac{n}{n+2}$ and $g_n = \frac{n}{n-3}$: we have $f_n = O(g_n)$ and $f_n = \Theta(g_n)$, but $f_n \neq o(g_n)$.
- Exponentials always dominate polynomials: $n^a = o(e^{bn})$ whatever $a > 0$ and $b > 0$.
- Conversely, $e^{-bn} = o(n^{-a})$.

Out loud, we can read the expression $f_n = O(g_n)$ as “ f_n in on the order of g_n ”.

The same notations can be used to compare sequences indexed by a parameter that goes to zero, such as (typically) the grid spacing h . The definition above is simply adapted by letting $h \rightarrow 0$ rather than $n \rightarrow \infty$.

Examples:

- $f(h) = h^2$ and $g(h) = h^3$: this time we have $g(h) = O(f(h))$ and $g(h) = o(f(h))$ when the limit of interest is $h \rightarrow 0$.
- Powers of h don't converge to zero nearly as fast as this exponential: $e^{a/h} = o(h^b)$ whatever $a > 0$ and $b > 0$.

Similarly, we may wish to compare functions f and g of a continuous variable x as either $x \rightarrow \infty$ or $x \rightarrow 0$; the definition is again modified in the obvious way. Whenever a $O(\cdot)$ or $o(\cdot)$ is written, some underlying limit is understood.

Chapter 2

Integrals as sums and derivatives as differences

We now switch to the simplest methods for integrating or differentiating a function from its function samples. A careful study of Taylor expansions reveals how accurate the constructions are.

2.1 Numerical integration

Consider a function f – we'll specify which assumptions we need to make about it in a minute. Let us reformulate the integral

$$\int_0^1 f(x) dx$$

by breaking up the interval $[a, b]$ into subintervals $[x_{j-1}, x_j]$, with $x_j = jh$, $h = 1/N$, and

$$0 = x_0 < x_1 < \dots < x_N = 1.$$

Together, the (x_j) are called a Cartesian grid. Still without approximation, we can write

$$\int_0^1 f(x) dx = \sum_{j=1}^N \int_{x_{j-1}}^{x_j} f(x) dx.$$

A *quadrature* is obtained when integrals are approximated by quantities that depend only on the samples $f(x_j)$ of f on the grid x_j . For instance, we can

approximate the integral over $[x_{j-1}, x_j]$ by the signed area of the rectangle of height $f(x_{j-1})$ and width h :

$$\int_{x_{j-1}}^{x_j} f(x) dx \simeq hf(x_{j-1}).$$

Putting the terms back together, we obtain the *rectangle method*:

$$\int_0^1 f(x) dx \simeq h \sum_{j=1}^N f(x_{j-1}).$$

(Insert picture here)

To understand the accuracy of this approximation, we need a detour through Taylor expansions. Given a very smooth function $f(x)$, it sometimes makes sense to write an expansion around the point $x = a$, as

$$f(x) = f(a) + f'(a)(x - a) + \frac{1}{2}f''(a)(x - a)^2 + \frac{1}{3!}f'''(a)(x - a)^3 + \dots$$

As a compact series, this would be written as

$$f(x) = \sum_{n=0}^{\infty} \frac{1}{n!} f^{(n)}(a)(x - a)^n,$$

where $n! = n \cdot (n - 1) \cdot \dots \cdot 2 \cdot 1$ is the factorial of n . The word “sometimes” is important: the formula above requires the function to be infinitely differentiable, and even still, we would need to make sure that the series converge (by means of the analysis of the previous chapter). The radius of convergence of the Taylor expansion is the largest R such that, for all x such that $|x - a| < R$, the series converges.

It is also possible to truncate a Taylor series after N terms, as

$$f(x) = \sum_{n=0}^{N-1} \frac{1}{n!} f^{(n)}(a)(x - a)^n + \frac{1}{N!} f^{(N)}(y)(x - a)^N, \quad (2.1)$$

where y is some number between x and a . The formula does not specify what the number y is; only that there exists one such that the remainder takes the form of the last term. In spite of the fact that y is unspecified, this is a much more useful form of Taylor expansion than the infinite series:

- We often only care about the size of the remainder, and write inequalities such as

$$\left| \frac{1}{N!} f^{(N)}(y)(x-a)^N \right| \leq \frac{1}{N!} \left(\max_{z \in [x, a]} |f^{(N)}(z)| \right) |x-a|^N,$$

or, for short,

$$\frac{1}{N!} f^{(N)}(y)(x-a)^N = O(|x-a|^N),$$

where all the “constants” are hidden in the O . So, the remainder is on the order of $|x-a|^N$. The limit implicit in the O notation is here $x-a \rightarrow 0$.

- The finite- N formula is valid whenever the function is N times differentiable, not infinitely differentiable! In fact, the Taylor series itself may diverge, but equation (2.1) is still valid.

In order to study the problem of approximating $\int_{x_{j-1}}^{x_j} f(x) dx$, let us expand f in a (short, truncated) Taylor series near $x = x_{j-1}$:

$$f(x) = f(x_{j-1}) + f'(y(x))h,$$

where $y(x) \in [x_{j-1}, x]$. We've written $y(x)$ to highlight that it depends on x . This works as long as f has one derivative. Integrating on both sides, and recognizing that $f(x_{j-1})$ is constant with respect to x , we get (recall $x_j - x_{j-1} = h$)

$$\int_{x_{j-1}}^{x_j} f(x) dx = hf(x_{j-1}) + h \int_{x_{j-1}}^{x_j} f'(y(x)) dx.$$

We don't know much about $y(x)$, but we can certainly write the inequality

$$h \left| \int_{x_{j-1}}^{x_j} f'(y(x)) dx \right| \leq h \max_{y \in [x_{j-1}, x_j]} |f'(y)| \int_{x_{j-1}}^{x_j} 1 dx = h^2 \max_{y \in [x_{j-1}, x_j]} |f'(y)|.$$

So, as long as f is differentiable, we have

$$\int_{x_{j-1}}^{x_j} f(x) dx = hf(x_{j-1}) + O(h^2),$$

where the derivative of f hides in the O sign. The integral over $[0, 1]$ has N such terms, so when they are added up the error compounds to $N \cdot O(h^2) = O(h)$ (because $h = 1/N$). Thus we have just proved that

$$\int_0^1 f(x) dx \simeq h \sum_{j=1}^N f(x_{j-1}) + O(h).$$

The error is on the order of the grid spacing h itself, so we say the method is first-order accurate (because the exponent of h is one.)

Choosing the (signed) height of each rectangle from the left endpoint x_{j-1} does not sound like the most accurate thing to do. Evaluating the function instead at $x_{j-1/2} = (j - 1/2)h$ is a better idea, called the midpoint method. It is possible to show that

$$\int_0^1 f(x) dx \simeq h \sum_{j=1}^N f(x_{j-1/2}) + O(h^2),$$

provided the function f is twice differentiable (because f'' hides in the O sign). The accuracy is improved, since h^2 gets much smaller than h as $h \rightarrow 0$. We say the midpoint method is second-order accurate.

Another solution to getting order-2 accuracy is to consider trapezoids instead of rectangles for the interval $[x_{j-1}, x_j]$. The area of the trapezoid spanned by the 4 points $(x_{j-1}, 0); (x_{j-1}, f(x_{j-1})); (x_j, f(x_j)); (x_j, 0)$ is $h(f(x_{j-1}) + f(x_j))/2$. This gives rise to the so-called trapezoidal method, or trapezoidal rule.

(Insert picture here)

Let us compare this quantity to the integral of f over the same interval. Consider the truncated Taylor expansions

$$f(x) = f(x_{j-1}) + f'(x_{j-1})(x - x_{j-1}) + O(h^2),$$

$$f(x_j) = f(x_{j-1}) + f'(x_{j-1})h + O(h^2),$$

where the second derivative of f appears in the O sign. Integrating the first relation gives

$$\int_{x_{j-1}}^{x_j} f(x) dx = hf(x_{j-1}) + \frac{h^2}{2}f'(x_{j-1}) + O(h^3).$$

The area of the trapezoid, on the other hand, is (using the second Taylor relation)

$$\frac{h}{2}(f(x_{j-1}) + f(x_j)) = hf(x_{j-1}) + \frac{h^2}{2}f'(x_{j-1})h + O(h^3).$$

Those two equations agree, except for the terms $O(h^3)$ which usually differ in the two expressions. Hence

$$\int_{x_{j-1}}^{x_j} f(x) dx = \frac{h}{2}(f(x_{j-1}) + f(x_j)) + O(h^3).$$

We can now sum over j (notice that all the terms appear twice, except for the endpoints) to obtain

$$\int_0^1 f(x) dx = \frac{h}{2}f(0) + h \sum_{j=1}^{N-1} f(x_j) + \frac{h}{2}f(1) + O(h^2).$$

The trapezoidal rule is second-order accurate. All it took is a modification of the end terms to obtain $O(h^2)$ accuracy in place of $O(h)$.

We'll return to the topic of numerical integration later, after we cover polynomial interpolation.

2.2 Numerical differentiation

The simplest idea for computing an approximation to the derivative $u'(x_j)$ of a function u from the samples $u_j = u(x_j)$ is to form finite difference ratios. On an equispaced grid $x_j = jh$, the natural candidates are:

- the one-sided, forward and backward differences

$$w_j^F = \frac{u_{j+1} - u_j}{h}, \quad w_j^B = \frac{u_j - u_{j-1}}{h},$$

- and the two-sided centered difference

$$w_j^C = \frac{u_{j+1} - u_{j-1}}{2h}.$$

Let us justify the accuracy of these difference quotients at approximating derivatives, as $h \rightarrow 0$. The analysis will depend on the smoothness of the underlying function u .

Assume without loss of generality that $x_{j-1} = -h$, $x_j = 0$, and $x_{j+1} = h$. For the forward difference, we can use a Taylor expansion about zero to get

$$u(h) = u(0) + hu'(0) + \frac{h^2}{2}u''(\xi), \quad \xi \in [0, h].$$

When substituted in the formula for the forward difference, we get

$$\frac{u(h) - u(0)}{h} - u'(0) = \frac{h}{2}u''(\xi).$$

The error is a $O(h)$ as soon as the function has two bounded derivatives. We say that the forward difference is a method of order one. The analysis for the backward difference is very similar.

For the centered difference, use two Taylor expansions about zero, for $u(h)$ and $u(-h)$, to get

$$\frac{u(h) - u(-h)}{2h} - u'(0) = \frac{h^2}{3}u'''(\xi).$$

The error is now $O(h^2)$, hence the method is of order 2.

The systematic way of deriving finite difference formulas is by differentiating a polynomial interpolant. We return to the topics of integration and differentiation in the next chapter.

Chapter 3

Interpolation

3.1 Polynomial interpolation

Given points $N + 1$ points $x_j \in \mathbb{R}$, $0 \leq j \leq N$, and sample values $y_j = f(x_j)$ of a function at these points, the polynomial interpolation problem consists in finding a polynomial $p_N(x)$ of degree N such that

$$y_j = p_N(x_j), \quad j = 0, \dots, N.$$

A central question is that of estimating the interpolation error

$$\|f - p_N\|_\infty,$$

where the L^∞ norm is taken over the interval $[x_0, x_N]$.

Call P_N the space of real-valued degree- N polynomials:

$$P_N = \left\{ \sum_{n=0}^N a_n x^n : a_n \in \mathbb{R} \right\}.$$

Lagrange's solution to the problem of polynomial interpolation is based on the following construction.

Lemma 1. (*Lagrange elementary polynomials*) Let $\{x_j, j = 0, \dots, N + 1\}$ be a collection of disjoint numbers. For each $k = 0, \dots, N$, there exists a unique polynomial $L_k(x)$ such that

$$L_k(x_j) = \delta_{ij} = \begin{cases} 1 & \text{if } j = k; \\ 0 & \text{if } j \neq k. \end{cases}$$

Proof. Fix k . L_k has roots at x_j for $j \neq k$, so L_k must be of the form¹

$$L_k(x) = C \prod_{j \neq k} (x - x_j).$$

Evaluating this expression at $x = x_k$, we get

$$1 = C \prod_{j \neq k} (x_k - x_j) \quad \Rightarrow \quad C = \frac{1}{\prod_{j \neq k} (x_k - x_j)}.$$

Hence the only possible expression for L_k is

$$L_k(x) = \frac{\prod_{j \neq k} (x - x_j)}{\prod_{j \neq k} (x_k - x_j)}.$$

□

These elementary polynomials will form a basis (in the sense of linear algebra) for expanding any polynomial interpolant p_N .

Theorem 4. (*Lagrange interpolation theorem*)

Let $\{x_j, j = 0, \dots, N + 1\}$ be a collection of disjoint real numbers. Let $\{y_j, j = 0, \dots, N\}$ be a collection of real numbers. Then there exists a unique $p_N \in P_N$ such that

$$p_N(x_j) = y_j, \quad j = 0, \dots, N.$$

Its expression is

$$p_N(x) = \sum_{k=0}^N y_k L_k(x), \tag{3.1}$$

where $L_k(x)$ are the Lagrange elementary polynomials.

¹That's because, if we fix j , we can divide $L_k(x)$ by $(x - x_j)$, $j \neq k$. We obtain $L_k(x) = (x - x_j)q(x) + r(x)$, where $r(x)$ is a remainder of lower order than $x - x_j$, i.e., a constant. Since $L_k(x_j) = 0$ we must have $r(x) = 0$. Hence $(x - x_j)$ must be a factor of $L_k(x)$. The same is true of any $(x - x_j)$ for $j \neq k$. There are N such factors, which exhausts the degree of N . The only remaining degree of freedom in L_k is the multiplicative constant.

Proof. The justification that (3.1) interpolates is obvious:

$$p_N(x_j) = \sum_{k=0}^N y_k L_k(x_j) = \sum_{k=0}^N y_k L_k(x_j) = y_j.$$

It remains to see that p_N is the unique interpolating polynomial. For this purpose, assume that both p_N and q_N take on the value y_j at x_j . Then $r_N = p_N - q_N$ is a polynomial of degree N that has a root at each of the $N + 1$ points x_0, \dots, x_N . The fundamental theorem of algebra, however, says that a nonzero polynomial of degree N can only have N (complex) roots. Therefore, the only way for r_N to have $N + 1$ roots is that it is the zero polynomial. So $p_N = q_N$. □

By definition,

$$p_N(x) = \sum_{k=0}^N f(x_k) L_k(x)$$

is called the *Lagrange interpolation polynomial* of f at x_j .

Example 7. (Example (6.1) in Suli-Mayers) Consider $f(x) = e^x$, and interpolate it by a parabola ($N = 2$) from three samples at $x_0 = -1, x_1 = 0, x_2 = 1$. We build

$$L_0(x) = \frac{(x - x_1)(x - x_2)}{(x_0 - x_1)(x_0 - x_2)} = \frac{1}{2}x(x - 1)$$

Similarly,

$$L_1(x) = 1 - x^2, \quad L_2(x) = \frac{1}{2}x(x + 1).$$

So the quadratic interpolant is

$$\begin{aligned} p_2(x) &= e^{-1}L_0(x) + e^0L_1(x) + e^1L_2(x), \\ &= 1 + \sinh(1)x + (\cosh(1) - 1)x^2, \\ &\simeq 1 + 1.1752x + 0.5431x^2. \end{aligned}$$

Another polynomial that approximates e^x reasonably well on $[-1, 1]$ is the Taylor expansion about $x = 0$:

$$t_2(x) = 1 + x + \frac{x^2}{2}.$$

Manifestly, p_2 is not very different from t_2 .

Let us now move on to the main result concerning the interpolation error of smooth functions.

Theorem 5. *Let $f \in C^{N+1}[a, b]$ for some $N > 0$, and let $\{x_j : j = 0, \dots, N\}$ be a collection of disjoint reals in $[a, b]$. Consider p_N the Lagrange interpolation polynomial of f at x_j . Then for every $x \in [a, b]$ there exists $\xi(x) \in [a, b]$ such that*

$$f(x) - p_N(x) = \frac{f^{(N+1)}(\xi(x))}{(N+1)!} \pi_{N+1}(x),$$

where

$$\pi_{N+1}(x) = \prod_{j=1}^{N+1} (x - x_j).$$

An estimate on the interpolation error follows directly from this theorem. Set

$$M_{N+1} = \max_{x \in [a, b]} |f^{(N+1)}(x)|$$

(which is well defined since $f^{(N+1)}$ is continuous by assumption, hence reaches its lower and upper bounds.) Then

$$|f(x) - p_N(x)| \leq \frac{M_{N+1}}{(N+1)!} |\pi_{N+1}(x)|$$

In particular, we see that the interpolation error is zero when $x = x_j$ for some j , as it should be.

Let us now prove the theorem.

Proof. (Can be found in Suli-Mayers, Chapter 6) □

In conclusion, the interpolation error:

- depends on the smoothness of f via the high-order derivative $f^{(N+1)}$;
- has a factor $1/(N+1)!$ that decays fast as the order $N \rightarrow \infty$;
- and is directly proportional to the value of $\pi_{N+1}(x)$, indicating that the interpolant may be better in some places than others.

The natural follow-up question is that of convergence: can we always expect convergence of the polynomial interpolant as $N \rightarrow \infty$? In other words, does the factor $1/(N+1)!$ always win over the other two factors?

Unfortunately, the answer is no in general. There are examples of very smooth (analytic) functions for which polynomial interpolation diverges, particularly so near the boundaries of the interpolation interval. This behavior is called the *Runge phenomenon*, and is usually illustrated by means of the following example.

Example 8. (*Runge phenomenon*) Let $f(x)$ for $x \in [-5, 5]$. Interpolate it at equispaced points $x_j = 10j/N$, where $j = -N/2, \dots, N/2$ and N is even. It is easy to check numerically that the interpolant diverges near the edges of $[-5, 5]$, as $N \rightarrow \infty$. See the Trefethen textbook on page 44 for an illustration of the Runge phenomenon.

(Figure here)

If we had done the same numerical experiment for $x \in [-1, 1]$, the interpolant would have converged. This shows that the size of the interval matters. Intuitively, there is divergence when the size of the interval is larger than the "features", or characteristic length scale, of the function (here the width of the bump near the origin.)

The analytical reason for the divergence in the example above is due in no small part to the very large values taken on by $\pi_{N+1}(x)$ far away from the origin — in contrast to the relatively small values it takes on near the origin. This is a problem intrinsic to equispaced grids. We will be more quantitative about this issue in the section on Chebyshev interpolants, where a remedy involving non-equispaced grid points will be explained.

As a conclusion, polynomial interpolants can be good for small N , and on small intervals, but may fail to converge (quite dramatically) when the interpolation interval is large.

3.2 Polynomial rules for integration

In this section, the goal is to approximate $\int_a^b u(x)dx$ by a weighted sum of samples $u(x_j)$, also called a quadrature. The master plan is the same as for differentiation: we will form interpolants of the data, integrate those interpolants, and deduce corresponding quadrature formulas.

3.2.1 Polynomial rules

The simplest idea for numerical integration is to represent a function by a piecewise polynomial, and integrate it. This gives rise to the following rules:

- The rectangles rule, where we approximate the function $u(x)$ by the constant $u(x_{j-1})$ in $[x_{j-1}, x_j]$. The result quadrature is

$$\sum_{j=1}^N h_j u(x_{j-1}).$$

- The trapezoidal rule, where we approximate the function $u(x)$ by a line joining $(x_{j-1}, u(x_{j-1}))$ and $(x_j, u(x_j))$ in $[x_{j-1}, x_j]$. The result is

$$\sum_{j=1}^N h_j \left(\frac{u(x_{j-1}) + u(x_j)}{2} \right).$$

On an equispaced grid, where $h_i = h$, the terms rearrange as

$$\frac{h}{2}u(x_0) + h \sum_{j=1}^{N-1} u(x_{j-1}) + \frac{h}{2}u(x_N).$$

- Simpson's rule, where we approximate the function by parabolas (see Suli-Mayers, Chap. 7), etc.

It is a good exercise to check, either by elementary Taylor expansions or by invoking the theory of polynomial interpolation, that the error incurred for integrating a smooth function in the subinterval $[x_{j-1}, x_j]$, is $O(h^2)$ for the rectangles rule, $O(h^3)$ for the trapezoidal rule, (and $O(h^5)$ for Simpson's rule,) where $h = \max_j h_j$. Since there are $O(1/h)$ such subintervals, the overall error is in general a $O(h)$ for the rectangles rule (a method of order 1) and a $O(h^2)$ for the trapezoidal rule (a method of order 2).

The higher-order variants of polynomial rules, called Newton-Cotes rules, are not very interesting because the Runge phenomenon kicks in again. Piecewise spline interpolation is not a good choice for numerical integration either, because of the two leftover degrees of freedom, whose arbitrary choice affects the accuracy of quadrature in an unacceptable manner.

3.3 Polynomial rules for differentiation

A systematic way of deriving finite difference formulas of higher order is to view them as derivatives of a polynomial interpolant passing through a small number of points neighboring x_j . For instance, it can be checked that the forward difference is obtained from the line joining $(x_j, f(x_j))$ and $(x_{j+1}, f(x_{j+1}))$. The centered difference is obtained from the parabola passing through the points $(x_{j-1}, f(x_{j-1}))$, $(x_j, f(x_j))$, and $(x_{j+1}, f(x_{j+1}))$. Differentiating one-sided interpolation polynomials is a good way to obtain one-sided difference formulas, which comes in handy at the boundaries of the computational domain. The following result establishes that the order of the finite difference formula matches the order of the polynomial being differentiated.

Theorem 6. *Let $f \in C^{N+1}[a, b]$, $\{x_j, j = 0, \dots, N\}$ some disjoint points, and p_N the corresponding interpolation polynomial. Then*

$$f'(x) - p'_N(x) = \frac{f^{(N+1)}(\xi)}{N!} \pi_N^*(x),$$

for some $\xi \in [a, b]$, and where $\pi_N^*(x) = (x - \eta_1) \dots (x - \eta_N)$ for some $\eta_j \in [x_{j-1}, x_j]$.

The proof of this result is an application of Rolle's theorem that we leave out. (It is not simply a matter of differentiating the error formula for polynomial interpolation, because we have no guarantee on $d\xi/dx$.)

A consequence of this theorem is that the error in computing the derivative is a $O(h^N)$, where $h = b - a$, as can be seen from majorizing the product $\pi_N^*(x)$.

It is interesting to notice, at least empirically, that the Runge's phenomenon is absent when the derivative is evaluated at the center of the interval over which the interpolant is built.

[Differentiation matrices: see the very good discussion in Trefethen's book.]

3.4 Piecewise polynomial interpolation

The idea of piecewise polynomial interpolation, also called spline interpolation, is to subdivide the interval $[a, b]$ into a large number of subintervals

$[x_{j-1}, x_j]$, and to use low-degree polynomials over each subintervals. This helps avoiding the Runge phenomenon. The price to pay is that the interpolant is no longer a C^∞ function – instead, we lose differentiability at the junctions between subintervals, where the polynomials are made to match.

If we use polynomials of order n , then the interpolant is piecewise C^∞ , and overall C^k , with $k \leq n - 1$. We could not expect to have $k = n$, because it would imply that the polynomials are identical over each subinterval. We'll see two examples: linear splines when $n = 1$, and cubic splines when $n = 3$. (We have seen in the homework why it is a bad idea to choose $n = 2$.)

3.4.1 Linear splines

We wish to interpolate a continuous function $f(x)$ of $x \in [a, b]$, from the knowledge of $f(x_j)$ at some points x_j , $j = 0, \dots, N$, not necessarily equispaced. Assume that $x_0 = a$ and $x_N = b$. The piecewise linear interpolant is build by tracing a straight line between the points $(x_{j-1}, f(x_{j-1}))$ and $(x_j, f(x_j))$; for $j = 1, \dots, N$ the formula is simply

$$s_L(x) = \frac{x_{j-1} - x}{x_j - x_{j-1}} f(x_{j-1}) + \frac{x - x_{j-1}}{x_j - x_{j-1}} f(x_j), \quad x \in [x_{j-1}, x_j].$$

In this case we see that $s_L(x)$ is a continuous function of x , but that its derivative is not continuous at the junction points, or nodes x_j .

If the function is at least twice differentiable, then piecewise linear interpolation has second-order accuracy, i.e., an error $O(h^2)$.

Theorem 7. *Let $f \in C^2[a, b]$, and let $h = \max_{j=1, \dots, N} (x_j - x_{j-1})$ be the grid diameter. Then*

$$\|f - s_L\|_{L^\infty[a, b]} \leq \frac{h^2}{8} \|f''\|_{L^\infty[a, b]}.$$

Proof. Let $x \in [x_{j-1}, x_j]$ for some $j = 1, \dots, N$. We can apply the basic result of accuracy of polynomial interpolation with $n = 1$: there exists $\xi \in [x_{j-1}, x_j]$ such that

$$f(x) - s_L(x) = \frac{1}{2} f''(\xi) (x - x_{j-1})(x - x_j), \quad x \in [x_{j-1}, x_j].$$

Let $h_j = x_j - x_{j-1}$. It is easy to check that the product $(x - x_{j-1})(x - x_j)$ takes its maximum value at the midpoint $\frac{x_{j-1} + x_j}{2}$, and that the value is $h_j^2/4$.

We then have

$$|f(x) - s_L(x)| \leq \frac{h_j^2}{8} \max_{\xi \in [x_{j-1}, x_j]} |f''(\xi)|, \quad x \in [x_{j-1}, x_j].$$

The conclusion follows by taking a maximum over j . \square

We can express any piecewise linear interpolant as a superposition of “tent” basis functions $\phi_k(x)$:

$$s_L(x) = \sum_k c_k \phi_k(x),$$

where $\phi_k(x)$ is the piecewise linear function equal to 1 at $x = x_k$, and equal to zero at all the other grid points x_j , $j \neq k$. Or in short, $\phi_k(x_j) = \delta_{jk}$. On an equispaced grid $x_j = jh$, an explicit formula is

$$\phi_k(x) = \frac{1}{h} S_{(1)}(x - x_k),$$

where

$$S_{(1)}(x) = x_+ - 2(x - h)_+ + (x - 2h)_+,$$

and where x_+ denotes the positive part of x (i.e., x if $x \geq 0$, and zero if $x < 0$.)

Observe that we can simply take $c_k = f(x_k)$ above. The situation will be more complicated when we pass to higher-order polynomials.

3.4.2 Cubic splines

Let us now consider the case $n = 3$ of an interpolant which is a third-order polynomial, i.e., a cubic, on each subinterval. The most we can ask is that the value of the interpolant, its derivative, and its second derivative be continuous at the junction points x_j .

Definition 5. (*Interpolating cubic spline*) Let $f \in C[a, b]$, and $\{x_j, j = 0, \dots, N\} \subset [a, b]$. An interpolating cubic spline is a function $s(x)$ such that

1. $s(x_j) = f(x_j)$;
2. $s(x)$ is a polynomial of degree 3 over each segment $[x_{j-1}, x_j]$;

3. $s(x)$ is globally C^2 , i.e., at each junction point x_j , we have the relations

$$s(x_j^-) = s(x_j^+), \quad s'(x_j^-) = s'(x_j^+), \quad s''(x_j^-) = s''(x_j^+),$$

where the notations $s(x_j^-)$ and $s(x_j^+)$ refer to the adequate limits on the left and on the right.

Let us count the degrees of freedom. A cubic polynomial has 4 coefficients. There are $N + 1$ points, hence N subintervals, for a total of $4N$ numbers to be specified.

The interpolating conditions $s(x_j) = f(x_j)$ specify two degrees of freedom per polynomial: one value at the left endpoint x_{j-1} , and one value at the right endpoint x_j . That's $2N$ conditions. Continuity of $s(x)$ follows automatically. The continuity conditions on s', s'' are imposed only at the interior grid points x_j for $j = 1, \dots, N - 1$, so this gives rise to $2(N - 1)$ additional conditions, for a total of $4N - 2$ equations.

There is a mismatch between the number of unknowns ($4N$) and the number of conditions ($4N - 2$). Two more degrees of freedom are required to completely specify a cubic spline interpolant, hence the precaution to write “an” interpolant in the definition above, and not “the” interpolant.

The most widespread choices for fixing these two degrees of freedom are:

- Natural splines:

$$s''(x_0) = s''(x_N) = 0.$$

If $s(x)$ measures the displacement of a beam, a condition of vanishing second derivative corresponds to a free end.

- Clamped spline:

$$s'(x_0) = p_0, \quad s'(x_N) = p_N,$$

where p_0 and p_N are specified values that depend on the particular application. If $s(x)$ measures the displacement of a beam, a condition of vanishing derivative corresponds to a horizontal clamped end.

- Periodic spline: assuming that $s(x_0) = s(x_N)$, then we also impose

$$s'(x_0) = s'(x_N), \quad s''(x_0) = s''(x_N).$$

Let us now explain the algorithm most often used for determining a cubic spline interpolant, i.e., the $4N$ coefficients of the N cubic polynomials, from the knowledge of $f(x_j)$. Let us consider the natural spline.

It is advantageous to write a system of equations for the second derivatives at the grid points, that we denote $\sigma_j = s''(x_j)$. Because $s(x)$ is piecewise cubic, we know that $s''(x)$ is piecewise linear. Let $h_j = x_j - x_{j-1}$. We can write

$$s''(x) = \frac{x_j - x}{h_j} \sigma_{j-1} + \frac{x - x_{j-1}}{h_j} \sigma_j, \quad x \in [x_{j-1}, x_j].$$

Hence

$$s(x) = \frac{(x_j - x)^3}{6h_j} \sigma_{j-1} + \frac{(x - x_{j-1})^3}{6h_j} \sigma_j + \alpha_j(x - x_{j-1}) + \beta_j(x_j - x).$$

We could have written $ax + b$ for the effect of the integration constants in the equation above, but writing it in terms of α_j and β_j makes the algebra that follows simpler. The two interpolation conditions for $[x_{j-1}, x_j]$ can be written

$$\begin{aligned} s(x_{j-1}) = f(x_{j-1}) &\quad \Rightarrow \quad f(x_{j-1}) = \frac{\sigma_{j-1} h_j^2}{6} + h_j \beta_j, \\ s(x_j) = f(x_j) &\quad \Rightarrow \quad f(x_j) = \frac{\sigma_j h_j^2}{6} + h_j \alpha_j. \end{aligned}$$

One then isolates α_j, β_j in this equation; substitutes those values in the equation for $s(x)$; and evaluates the latter at x_j to get $s(x_j) = f(x_j)$. Given that $\sigma_0 = \sigma_N = 0$, we end up with a system of $N - 1$ equations in the $N - 1$ unknowns $\sigma_1, \dots, \sigma_N$. Skipping the algebra, the end result is

$$h_j \sigma_{j-1} + 2(h_{j+1} + h_j) \sigma_j + h_{j+1} \sigma_{j+1} = 6 \left(\frac{f(x_{j+1}) - f(x_j)}{h_{j+1}} - \frac{f(x_j) - f(x_{j-1})}{h_j} \right).$$

We are in presence of a tridiagonal system for σ_j . It can be solved efficiently with Gaussian elimination, yielding a LU decomposition with bidiagonal factors. Unlike in the case of linear splines, there is no way around the fact that a linear system needs to be solved.

Notice that the tridiagonal matrix of the system above is diagonally dominant (each diagonal element is strictly greater than the sum of the other elements on the same row, in absolute value), hence it is always invertible.

One can check that the interpolation for cubic splines is $O(h^4)$ well away from the endpoints. This requires an analysis that is too involved for the present set of notes.

Finally, like in the linear case, let us consider the question of expanding a cubic spline interpolant as a superposition of basis functions

$$s(x) = \sum_k c_k \phi_k(x).$$

There are many ways of choosing the functions $\phi_k(x)$, so let us specify that they should have as small a support as possible, that they should have the same smoothness C^2 as $s(x)$ itself, and that they should be translates of each other when the grid x_j is equispaced with spacing h .

The only solution to this problem is the cubic B-spline. Around any interior grid point x_k , is supported in the interval $[x_{k-2}, x_{k+2}]$, and is given by the formula

$$\phi_k(x) = \frac{1}{4h^3} S_{(3)}(x - x_{k-2}),$$

where

$$S_{(3)}(x) = x_+^3 - 4(x-h)_+^3 + 6(x-2h)_+^3 - 4(x-3h)_+^3 + (x-4h)_+^3,$$

and where x_+ is the positive part of x . One can check that $\phi_k(x)$ takes the value 1 at x_k , $1/4$ at $x_{k\pm 1}$, zero outside of $[x_{k-2}, x_{k+2}]$, and is C^2 at each junction. It is a bell-shaped curve. It is an interesting exercise to check that it can be obtained as the convolution of two “tent” basis functions of linear interpolation:

$$S_{(3)}(x) = cS_{(1)}(x) \star S_{(1)}(x),$$

where c is some constant.

Now with cubic B-splines, we cannot put $c_k = f(x_k)$ anymore, since $\phi_k(x_j) \neq \delta_{jk}$. The particular values of c_k are the result of solving a linear system, as mentioned above. Again, there is no way around solving a linear system. In particular, if $f(x_k)$ changes at one point, or if we add a datum point (x_k, f_k) , then the update requires re-computing the whole interpolant. Changing one point has ripple effects throughout the whole interval $[a, b]$. This behavior is ultimately due to the more significant overlap between neighboring ϕ_k in the cubic case than in the linear case.

Chapter 4

Methods for ODE

Chapter 5

Ordinary differential equations

5.1 Initial-value problems

Initial-value problems (IVP) are those for which the solution is entirely known at some time, say $t = 0$, and the question is to solve the ODE

$$y'(t) = f(t, y(t)), \quad y(0) = y_0,$$

for other times, say $t > 0$. We will consider a scalar y , but considering systems of ODE is a straightforward extension for what we do in this chapter. We'll treat both theoretical questions of existence and uniqueness, as well as practical questions concerning numerical solvers.

We speak of t as being time, because that's usually the physical context in which IVP arise, but it could very well be a space variable.

5.1.1 Theory

Does a solution exist, is it unique, and does it tend to infinity (blow up) in finite time? These questions are not merely pedantic. As we now show with two examples, things can go wrong very quickly if we posit the wrong ODE.

Example 9. *Consider*

$$y' = \sqrt{y}, \quad y(0) = 0.$$

By separation of variables, we find

$$\int \frac{dy}{\sqrt{y}} = \int dt \quad \Rightarrow \quad y(t) = \frac{(t + C)^2}{4}.$$

Imposing the initial condition yields $C = 0$, hence $y(t) = t^2/4$. However, $y(t) = 0$ is clearly another solution, so we have non-uniqueness. In fact, there is an infinite number of solutions, corresponding to $y(t) = 0$ for $0 \leq t \leq t^*$ for some t^* , which then takes off along the parabola $y(t) = (t-t^*)^2/4$ for times $t \geq t^*$.

Example 10. Consider

$$y' = y^2, \quad y(0) = 1.$$

By separation of variables, we find

$$\int \frac{dy}{y^2} = \int dt \quad \Rightarrow \quad y(t) = \frac{-1}{t+C}.$$

The initial condition gives $C = -1$, hence $y(t) = 1/(1-t)$. It blows up at time $t = 1$, because $y(t)$ has a vertical asymptote. We say that the solution exists locally in any interval to the left of $t = 1$, but we don't have global existence.

Blowup and non-uniqueness are generally, although not always¹, unrealistic in applications. A theory of existence and uniqueness, including global existence (non blowup), is desired to guide us in formulating valid models for physical phenomena.

The basic result is the following.

Theorem 8. (Picard) For given T, C , and y_0 , consider the box B in (t, y) space, given by

$$B = [0, T] \times [y_0 - C, y_0 + C].$$

Assume that

- $f(t, y)$ is continuous over B ;
- $|f(t, y)| \leq K$ when $(t, y) \in B$; (boundedness)
- $|f(t, u) - f(t, v)| \leq L|u - v|$ when $(t, u), (t, v) \in B$. (Lipschitz continuity).

¹In nonlinear optics for instance, laser pulses may “collapse”. This situation is somewhat realistically modeled by an ODE that blows up, although not for times arbitrarily close to the singularity.

Assume furthermore that $C \geq \frac{K}{L}(e^{LT} - 1)$. Then there exists a unique $y \in C^1[0, T]$, such that

$$y'(t) = f(t, y(t)), \quad y(0) = y_0,$$

and such that $|y(t) - y_0| \leq C$. In short, the solution exists, is unique, and stays in the box for times $0 \leq t \leq T$.

Proof. The technique is called Picard's iteration. See p.311 in Suli-Mayers. \square

5.1.2 Numerics

Here is an overview of some of the most popular numerical methods for solving ODEs. Let $t_n = nh$, and denote by y_n the approximation of $y(t_n)$.

1. Explicit Euler.

$$y_{n+1} = y_n + hf(t_n, y_n).$$

This formula comes from approximating the derivative y' at $t = t_n$ by a forward difference. It allows to march in time from the knowledge of y_n , to get y_{n+1} .

2. Implicit Euler.

$$y_{n+1} = y_n + hf(t_{n+1}, y_{n+1}).$$

This time we use a backward difference for approximating the derivative at $t = t_{n+1}$. The unknown y_{n+1} appears implicitly in this equation, it still needs to be solved for as a function of y_n , using (for instance) Newton's method. The strategy is still to march in time, but at every step there is a nonlinear equation to solve.

3. Midpoint rule (implicit).

$$y_{n+1} = y_n + \frac{h}{2} [f(t_n, y_n) + f(t_{n+1}, y_{n+1})].$$

The derivative is approximated by a centered difference, at $t = \frac{t_n + t_{n+1}}{2}$ (the midpoint). This gives a more balanced estimate of the slope $\frac{y_{n+1} - y_n}{h}$. It is an implicit method: y_{n+1} needs to be solved for.

4. Improved Euler, Runge-Kutta 2 (explicit).

$$\begin{aligned}\tilde{y}_{n+1} &= y_n + hf(t_n, y_n), \\ y_{n+1} &= y_n + \frac{h}{2} [f(t_n, y_n) + f(t_{n+1}, \tilde{y}_{n+1})].\end{aligned}$$

This is the simplest of “predictor-corrector” methods. It is like the midpoint rule, except that we use a guess \tilde{y}_{n+1} for the unknown value of y_{n+1} in the right-hand side, and this guess comes from the explicit Euler method. Now y_{n+1} only appears in the left-hand side, so this is an explicit method.

5. Runge-Kutta 4 (explicit).

$$y_{n+1} = y_n + h[k_1 + 2k_2 + 2k_3 + k_4],$$

where the slopes k_1, \dots, k_4 are given in succession by

$$\begin{aligned}k_1 &= f(t_n, y_n), & k_2 &= f\left(t_n + \frac{h}{2}, y_n + \frac{h}{2}k_1\right), \\ k_3 &= f\left(t_n + \frac{h}{2}, y_n + \frac{h}{2}k_2\right), & k_4 &= f(t_n + h, y_n + hk_3).\end{aligned}$$

6. There are also methods that involve not just the past value y_n , but a larger chunk of history y_{n-1}, y_{n-2} , etc. These methods are called multi-step. They are in general less flexible than the one-step methods described so far, in that they require a constant step h as we march in time.

Two features of a numerical method are important when choosing a numerical method:

- Is it *convergent*, i.e., does the computed solution tend to the true solution as $h \rightarrow 0$, and at which rate?
- Is it *stable*, i.e., if we solve with different initial conditions y_0 and \tilde{y}_0 , are the computed solutions close in the sense that $|y_n - \tilde{y}_n| \leq C|y_0 - \tilde{y}_0|$, with $n = O(1/h)$, and C independent of h ?

To understand convergence better in the case of one-step methods, let us write the numerical scheme as

$$y_{n+1} = \Psi(t_n, y_n, h),$$

and introduce the *local error*

$$e_{n+1}(h) = \Psi(t_n, y(t_n), h) - y(t_{n+1}),$$

as well as the *global error*

$$E_n(h) = y_n - y(t_n).$$

Convergence is the study of the global error. The local error, however, is easier to understand. A numerical method is called consistent if the local error decays sufficiently fast as $h \rightarrow 0$ that there is hope that the global error would be small as well. The particular rate at which the local error decays is related to the notion of order of an ODE solver.

Definition 6. (*Consistency*) Ψ is consistent if, for any $n \geq 0$,

$$\lim_{h \rightarrow 0} \frac{e_n(h)}{h} = 0$$

Definition 7. (*Order*) Ψ is of order p if $e_n(h) = O(h^{p+1})$.

The basic convergence result for one-step solvers, that we will not prove, is that if the local error is $O(h^{p+1})$, then the global error is $O(h^p)$. This convergence result is only true as stated for one-step methods; for multi-step methods we would also need an assumption of stability (discussed below). Intuitively, the local errors compound over the $O(1/h)$ time steps necessary to reach a given fixed time t , hence the loss of one power of h . Of course the local errors don't exactly add up; they do up to a multiplicative constant. It is the behavior of the global error that dictates the notion of order of the numerical scheme.

It is a good exercise to show, using elementary Taylor expansions, that the explicit and implicit Euler methods are of order 1, and that the midpoint rule and improved Euler methods are of order 2. It turns out that Runge-Kutta 4 is of order 4, but it is not much fun to prove that.

Quantifying stability as stated above is a complex question. The important ideas already appear if we study the representative setting of linear stability.

Definition 8. (*A-stability*) Suppose $y' = \lambda y$ for some $\lambda \in \mathbb{C}$. Then Ψ is A-stable if $y_n \rightarrow 0$ as $n \rightarrow \infty$.

Of course A-stability depends on the value of λ . Stability for the original equation is guaranteed if $\operatorname{Re}(\lambda) < 0$, and the question is that of showing whether a numerical method Ψ is stable under the same condition or not. Let us consider the explicit Euler method.

Example 11. For explicit Euler applied to $y' = \lambda y$, we get

$$y_{n+1} = y_n + h\lambda y_n = (1 + h\lambda)y_n.$$

The iterates y_n tend to zero provided $|1 + h\lambda| < 1$, where the $|\cdot|$ denote the complex modulus. This is the equation of a disc in the complex $(h\lambda)$ -plane, with center at -1 , and radius 1. If $h\lambda$ sits inside this disk, the method is stable, and otherwise it isn't. We say that the explicit Euler method is conditionally stable: typically, we require both $\operatorname{Re}(\lambda) < 0$ and a small step size h in order to guarantee stability.

Similar reasonings show that the stability region for the implicit Euler method is the whole complex plane, take away a disk centered at $h\lambda = 1$, with radius 1. The stability region for the midpoint rule is the left-half of the complex plane, i.e. $\operatorname{Re}(\lambda) < 0$, the same condition as for the equation. The stability region for Runge-Kutta 4 is plotted on page 351 of Suli-Mayers.

5.2 Boundary-value problems

Chapter 6

Root-finding

Chapter 7

Fourier analysis

(Historical intro: the heat equation on a square plate or interval.)

Fourier's analysis was tremendously successful in the 19th century for formulating series expansions for solutions of some very simple ODE and PDE. This class shows that in the 20th century, Fourier analysis has established itself as a central tool for numerical computations as well, for vastly more general ODE and PDE when explicit formulas are not available.

7.1 The Fourier transform

We will take the Fourier transform of integrable functions of one variable $x \in \mathbb{R}$.

Definition 9. (*Integrability*) A function f is called *integrable*, or *absolutely integrable*, when

$$\int_{-\infty}^{\infty} |f(x)| dx < \infty,$$

in the sense of Lebesgue integration. One also writes $f \in L^1(\mathbb{R})$ for the space of integrable functions.

We denote the physical variable as x , but it is sometimes denoted by t in contexts in which its role is time, and one wants to emphasize that. The frequency, or wavenumber variable is denoted k . Popular alternatives choices for the frequency variable are ω (engineers) or ξ (mathematicians), or p (physicists).

Definition 10. The Fourier transform (FT) of an integrable function $f(x)$ is defined as

$$\hat{f}(k) = \int_{-\infty}^{\infty} e^{-ikx} f(x) dx. \quad (7.1)$$

When $\hat{f}(k)$ is also integrable, $f(x)$ can be recovered from $\hat{f}(k)$ by means of the inverse Fourier transform (IFT)

$$f(x) = \frac{1}{2\pi} \int_{-\infty}^{\infty} e^{ikx} \hat{f}(k) dk. \quad (7.2)$$

Intuitively, $\hat{f}(k)$ is the amplitude density of f at frequency k . The formula for recovering f is a decomposition of f into constituent waves.

The justification of the inverse FT formula belongs in a real analysis class (where it is linked to the notion of approximate identity.) We will justify the form of (7.2) heuristically when we see Fourier series in the next section.

The precaution of assuming integrability is so that the integrals can be understood in the usual Lebesgue sense. In that context, taking integrals over infinite intervals is perfectly fine. If (7.1) and (7.2) are understood as limits of integrals over finite intervals, it does not matter how the bounds are chosen to tend to $\pm\infty$.

One may in fact understand the formulas for the FT and IFT for much larger function classes than the integrable functions, namely distributions, but this is also beyond the scope of the class. We will generally not overly worry about these issues. It is good to know where to draw the line: the basic case is that of integrable functions, and anything beyond that requires care and adequate generalizations.

Do not be surprised to see alternative formulas for the Fourier transform in other classes or other contexts. Wikipedia lists them.

Here are some important properties of Fourier transforms:

- (Differentiation)

$$\widehat{f'}(k) = ik\hat{f}(k).$$

Justification: integration by parts in the integral for the FT.

- (Translation) If $g(x) = f(x + a)$, then

$$\hat{g}(k) = e^{ika} \hat{f}(k).$$

Justification: change of variables in the integral for the FT.

Let's see some examples of FT.

Example 12. *Let*

$$f(x) = \frac{1}{2a} \chi_{[-a,a]}(x) = \begin{cases} 1 & \text{if } x \in [-a, a]; \\ 0 & \text{otherwise.} \end{cases}$$

Then

$$\hat{f}(k) = \frac{1}{2a} \int_{-a}^a e^{-ikx} dx = \frac{\sin(ka)}{ka}.$$

This function is a scaled version of the sinc function,

$$\text{sinc}(k) = \frac{\sin k}{k}.$$

It is easy to check by L'Hospital's rule that

$$\text{sinc}(0) = 1.$$

At $k \rightarrow \infty$, $\text{sinc}(k)$ decays like $1/k$, but does so by alternating between positive and negative values. It is a good exercise to check that sinc is not absolutely integrable. It turns out that the Fourier transform can still be defined for it, so lack of integrability is not a major worry.

Example 13. *Consider the Gaussian function*

$$f(x) = e^{-x^2/2}.$$

By completing the square and adequately modifying the contour of integration in the complex plane, it can be shown that

$$\hat{f}(k) = \sqrt{\frac{\pi}{2}} e^{-k^2/2}.$$

Another basic property of Fourier transforms is the convolution theorem.

Theorem 9. *(The convolution theorem.) Denote convolution as $f \star g(x) = \int_{-\infty}^{\infty} f(y)g(x-y) dy$. Then*

$$\widehat{f \star g}(k) = \hat{f}(k) \hat{g}(k).$$

Proof. Let $h = f \star g$. We can use Fubini below provided every function is integrable.

$$\begin{aligned}\hat{h}(k) &= \int e^{-ikx} \int f(y)g(x-y) dy dx \\ &= \int \int e^{-iky} f(y) e^{-ik(x-y)} g(x-y) dy dx \\ &= \left(\int e^{-iky} f(y) dy \right) \left(\int e^{-ikx'} g(x') dx' \right) \\ &= \hat{f}(k) \hat{g}(k).\end{aligned}$$

□

The Fourier transform is an important tool in the study of linear differential equations because it turns differential problems into algebraic problems. For instance, consider a polynomial $P(x) = \sum a_n x^n$, and the ODE

$$P\left(\frac{d}{dx}\right)u(x) = f(x), \quad x \in \mathbb{R},$$

which means $\sum a_n \frac{d^n u}{dx^n} = f$. (Such ODE are not terribly relevant in real life because they are posed over the whole real line.) Upon Fourier transformation, the equation becomes

$$P(ik)\hat{u}(k) = \hat{f}(k),$$

which is simply solved as

$$\hat{u}(k) = \frac{\hat{f}(k)}{P(ik)},$$

and then

$$u(x) = \frac{1}{2\pi} \int_{-\infty}^{\infty} e^{ikx} \frac{\hat{f}(k)}{P(ik)} dk.$$

Beware the zeros of P when applying this formula! They always carry important physical interpretation. For instance, they could be resonances of a mechanical system.

The formula $\hat{u}(k) = \frac{\hat{f}(k)}{P(ik)}$ also lends itself to an application of the convolution theorem. Let $K(x)$ be the inverse Fourier transform of $1/P(ik)$. Then we have

$$u(x) = \int K(x-y)f(y) dy.$$

The function K is called *Green's function* for the original ODE.

7.2 Sampling and restriction

We aim to use Fourier transforms as a concept to help understand the accuracy of representing and manipulating functions on a grid, using a finite number of degrees of freedom. We also aim at using a properly discretized Fourier transform as a numerical tool itself.

For this purpose, $x \in \mathbb{R}$ and $k \in \mathbb{R}$ must be replaced by x and k on finite grids. Full discretization consists of *sampling* and *restriction*.

Let us start by sampling $x \in h\mathbb{Z}$, i.e., considering $x_j = jh$ for $j \in \mathbb{Z}$. The important consequence of sampling is that some complex exponential waves e^{ikx} for different k will appear to be the same on the grid x_j . We call *aliases* such functions that identify on the grid.

Definition 11. (*Aliases*) The functions e^{ik_1x} and e^{ik_2x} are aliases on the grid $x_j = jh$ if

$$e^{ik_1x_j} = e^{ik_2x_j}, \quad \forall j \in \mathbb{Z}.$$

Aliases happen as soon as

$$k_1jh = k_2jh + 2\pi \times \text{integer}(j).$$

Letting $j = 1$, and calling the integer n , we have

$$k_1 - k_2 = \frac{2\pi}{h}n,$$

for some $n \in \mathbb{Z}$. Two wave numbers k_1, k_2 are indistinguishable on the grid if they differ by an integer multiple of $2\pi/h$.

For this reason, we restrict without loss of generality the wavenumber to the interval

$$k \in [-\pi/h, \pi/h].$$

We also call this interval the fundamental cell in frequency (in reference to a similar concept in crystallography.)

Real-life examples of aliases are rotating wheels looking like they go backwards in a movie, Moiré patterns on jackets on TV, and stroboscopy.

The proper notion of Fourier transform on a grid is the following.

Definition 12. Let $x_j = hj$, $f_j = f(x_j)$. Semidiscrete Fourier transform (SFT):

$$\hat{f}(k) = h \sum_{j=-\infty}^{\infty} e^{-ikx_j} f_j, \quad k \in [-\pi/h, \pi/h]. \quad (7.3)$$

Inverse semidiscrete Fourier transform (ISFT):

$$f_j = \frac{1}{2\pi} \int_{-\pi/h}^{\pi/h} e^{ikx} \hat{f}(k) dk. \quad (7.4)$$

As we saw, sampling in x corresponds to a restriction in k . If one still wanted to peek outside $[-\pi/h, \pi/h]$ for the SFT, then the SFT would simply repeat by periodicity:

$$\hat{f}\left(k + \frac{2n\pi}{h}\right) = \hat{f}(k).$$

(why?). That's why we restrict it to the fundamental cell.

We can now define the proper notion of Fourier analysis for functions that are restricted to x in some interval, namely $[-\pi, \pi]$ for convention. Unsurprisingly, the frequency is sampled as a result. The following formulas are dual to those for the SFT.

Definition 13. *Fourier series (FS):*

$$\hat{f}_k = \int_{-\pi}^{\pi} e^{-ikx} f(x) dx. \quad (7.5)$$

Inverse Fourier series (IFS)

$$f(x) = \frac{1}{2\pi} \sum_{k=-\infty}^{\infty} e^{ikx} \hat{f}_k, \quad x \in [-\pi, \pi]. \quad (7.6)$$

If one uses the Fourier series inversion formula (7.6) for x outside of its intended interval $[-\pi, \pi]$, then the function simply repeats by periodicity:

$$f(x + 2n\pi) = f(x).$$

(again, why?)

The two formulas (7.5) and (7.6) can be justified quite intuitively. The expression $\int f(x)\overline{g(x)} dx$ is an inner product on functions. It is easy to see that the complex exponentials

$$\sqrt{\frac{h}{2\pi}} e^{-ikx_j} = v_j(k)$$

form an orthonormal set of functions on $[-\pi/h, \pi/h]$, for this inner product. Hence, up to normalization constants, (7.5) is simply calculation of the coefficients in an orthobasis (analysis), and (7.6) is the synthesis of the function

back from those coefficients. We'd have to understand more about the peculiarities of infinite-dimensional linear algebra to make this fully rigorous, but this is typically done in a real analysis class.

Here's an example of SFT.

Example 14.

$$f_j = \frac{1}{2a} \chi_{[-a,a]}(x_j)$$

Then

$$\begin{aligned} \hat{f}(k) &= \frac{h}{2a} \sum_{j=-a}^a e^{-ikjh} \\ &= \frac{h}{2a} e^{ikah} \sum_{j=0}^{2a} e^{-ikjh} \\ &= \frac{h}{2a} e^{ikah} \frac{(e^{-ikh})^{2a+1} - 1}{e^{-ikh} - 1} \quad (\text{geometric series}) \\ &= \frac{h}{2a} \frac{\sin(kh(a + 1/2))}{\sin(kh/2)}. \end{aligned}$$

This function is called the discrete sinc. It looks like a sinc, but it periodizes smoothly when $k = -\pi/h$ and $k = \pi/h$ are identified.

Our first slogan is therefore:

Sampling in x corresponds to restriction/periodization in k , and restriction/periodization in k corresponds to sampling in x .

7.3 The DFT and its algorithm, the FFT

The discrete Fourier transform is what is left of the Fourier transform when both space and frequency are sampled and restricted to some interval.

Consider

$$x_j = jh, \quad j = 1, \dots, N.$$

The point $j = 0$ is identified with $j = N$ by periodicity, so it is not part of the grid. If the endpoints are $x_0 = 0$ and $x_N = 2\pi$, then N and h relate as

$$h = \frac{2\pi}{N} \quad \Rightarrow \quad \frac{\pi}{h} = \frac{N}{2}.$$

For the dual grid in frequency, consider that N points should be equispaced between the bounds $[-\pi/h, \pi/h]$. The resulting grid is

$$k, \quad k = -\frac{N}{2} + 1, \dots, \frac{N}{2}.$$

We have the following definition.

Definition 14. *Discrete Fourier transform (DFT):*

$$\hat{f}_k = h \sum_{j=1}^N e^{-ikjh} f_j, \quad k = -\frac{N}{2}, \dots, \frac{N}{2}. \quad (7.7)$$

Inverse discrete Fourier transform (IDFT)

$$f_j = \frac{1}{2\pi} \sum_{k=-N/2+1}^{N/2} e^{ikjh} \hat{f}_k, \quad j = 1, \dots, N. \quad (7.8)$$

The DFT can be computed as is, by implementing the formula (7.7) directly on a computer. The complexity of this calculation is a $O(N^2)$, since there are N values of j , and there are N values of k over which the computation must be repeated.

There is, however, a smart algorithm that allows to group the computation of all the f_k in complexity $O(N \log N)$. It is called the fast Fourier transform (FFT). It is traditionally due to Tukey and Cooley (1965), but the algorithm had been discovered a few times before that by people who are not usually credited as much: Danielson and Lanczos in 1942, and well as Gauss in 1805.

The trick of the FFT is to split the samples of f into even samples (j even) and odd samples (j odd). Assume that N is a power of 2. Then

$$\begin{aligned} \hat{f}_k &= h \sum_{j=1}^N e^{-ikjh} f_j \\ &= h \sum_{j=1}^{N/2} e^{-ik(2j)h} f_{2j} + h \sum_{j=1}^{N/2} e^{-ik(2j+1)h} f_{2j+1} \\ &= h \sum_{j=1}^{N/2} e^{-ikj(2h)} f_{2j} + h e^{ikh} \sum_{j=1}^{N/2} e^{-ikj(2h)} f_{2j+1}. \end{aligned}$$

The first term on the last line is simply the DFT of length $N/2$, on a grid of spacing $2h$, of the even samples of f . The second term is, besides the multiplicative factor, the DFT of length $N/2$, on a grid of spacing $2h$, of the odd samples of f .

Note that those smaller DFTs would normally be only calculated for $k = -\frac{N}{4} + 1, \dots, \frac{N}{4}$, but we need them for $k = -\frac{N}{2} + 1, \dots, \frac{N}{2}$. This is not a big problem: we know that the DFT extends by periodicity outside the standard bounds for k , so all there is to do is copy \hat{f}_k by periodicity outside of $k = -\frac{N}{4} + 1, \dots, \frac{N}{4}$.

Already, we can see the advantage in this reduction: solving one problem of size N is essentially reduced to solving two problems of size $N/2$. Even better, the splitting into even and odd samples can be repeated recursively until the DFT are of size 1. When $N = 1$, the DFT is simply multiplication by a scalar.

At each stage, there are $O(N)$ operations to do to put together the summands in the equation of the last line above. So the complexities obey

$$\text{Compl}(N) \leq 2\text{Compl}(N/2) + O(N), \quad \text{Compl}(1) = O(1).$$

It is easy to see that the solution of this equation is

$$\text{Compl}(N) = O(N \log N),$$

which is much better than $O(N^2)$.

There are variants of the FFT when N is not a power of 2.

7.4 Smoothness and truncation

In this section, we study the accuracy of truncation of Fourier transforms to finite intervals. This is an important question not only because real-life numerical Fourier transforms are restricted in k , but also because, as we know, restriction in k serves as a proxy for sampling in x . It will be apparent in Chapter 2, section 2.1 that every claim that we make concerning truncation of Fourier transforms will have an implication in terms of accuracy of sampling a function on a grid, i.e., how much information is lost in the process of sampling a function $f(x)$ at points $x_j = jh$.

We will manipulate functions in the spaces L^1 , L^2 , and L^∞ . We have already encountered L^1 .

Definition 15. Let $1 \leq p < \infty$. A function f of $x \in \mathbb{R}$ is said to belong to the space $L^p(\mathbb{R})$ when

$$\int_{-\infty}^{\infty} |f(x)|^p dx < \infty.$$

Then the norm of f in $L^p(\mathbb{R})$ is $\|f\|_p = \left(\int_{-\infty}^{\infty} |f(x)|^p dx \right)^{1/p}$.

A function f of $x \in \mathbb{R}$ is said to belong to L^∞ when

$$\text{ess sup } |f(x)| < \infty.$$

Then the norm of f in $L^\infty(\mathbb{R})$ is $\text{ess sup } |f(x)|$.

In the definition above, “ess sup” refers to the essential supremum, i.e., the infimum over all dense sets $X \subset \mathbb{R}$ of the supremum of f over X . A set X is dense when $\mathbb{R} \setminus X$ has measure zero. The notions of supremum and infimum correspond to maximum and minimum respectively, when they are not necessarily attained. All these concepts are covered in a real analysis class. For us, it suffices to heuristically understand the L^∞ norm as the maximum value of the modulus of the function, except possibly for isolated points of discontinuity which don’t count in calculating the maximum.

It is an interesting exercise to relate the L^∞ norm to the sequence of L^p norms as $p \rightarrow \infty$.

We will need the very important Parseval and Plancherel identities. They express “conservation of energy” from the physical domain to the frequency domain.

Theorem 10. (Parseval’s identity). Let $f, g \in L^1(\mathbb{R}) \cap L^2(\mathbb{R})$. Then

$$\int_{-\infty}^{\infty} f(x) \overline{g(x)} dx = \frac{1}{2\pi} \int_{-\infty}^{\infty} \hat{f}(k) \overline{\hat{g}(k)} dk.$$

Proof. Let h be the convolution $f \star \tilde{g}$, where $\tilde{g}(x) = \overline{g(-x)}$. It is easy to see that the Fourier transform of \tilde{g} is simply $\overline{\hat{g}(k)}$. By the convolution theorem (Section 1.1), we have

$$\hat{h}(k) = \hat{f}(k) \overline{\hat{g}(k)}.$$

If we integrate this relation over $k \in \mathbb{R}$, and divide by 2π , we get the IFT at $x = 0$:

$$h(0) = \frac{1}{2\pi} \int_{-\infty}^{\infty} \hat{f}(k) \overline{\hat{g}(k)} dk.$$

On the other hand,

$$h(0) = \int_{-\infty}^{\infty} f(x) \overline{g(-(0-x))} dx = \int_{-\infty}^{\infty} f(x) \overline{g(x)} dx.$$

□

Theorem 11. (*Plancherel's identity*). Let $f \in L^1(\mathbb{R}) \cap L^2(\mathbb{R})$. Then

$$\int_{-\infty}^{\infty} |f(x)|^2 dx = \frac{1}{2\pi} \int_{-\infty}^{\infty} |\hat{f}(k)|^2 dk.$$

Proof. Apply Parseval's identity with $g = f$. □

(With the help of these formulas, it is in fact possible to extend their validity and the validity of the FT to $f, g \in L^2(\mathbb{R})$, and not simply $f, g \in L^1(\mathbb{R}) \cap L^2(\mathbb{R})$. This is a classical density argument covered in many good analysis texts.)

We need one more concept before we get to the study of truncation of Fourier transforms. It is the notion of total variation. We assume that the reader is familiar with the spaces $C^k(\mathbb{R})$ of bounded functions which are k times continuously differentiable.

Definition 16. (*Total variation*) Let $f \in C^1(\mathbb{R})$. The total variation of f is the quantity

$$\|f\|_{TV} = \int_{-\infty}^{\infty} |f'(x)| dx. \quad (7.9)$$

For functions that are not C^1 , the notion of total variation is given by either expression

$$\|f\|_{TV} = \lim_{h \rightarrow 0} \int_{-\infty}^{\infty} \frac{|f(x) - f(x-h)|}{|h|} dx = \sup_{\{x_p\} \text{ finite subset of } \mathbb{R}} \sum_p |f(x_{p+1}) - f(x_p)|, \quad (7.10)$$

These more general expressions reduce to $\int_{-\infty}^{\infty} |f'(x)| dx$ when $f \in C^1(\mathbb{R})$. When a function has finite total variation, we say it is in the space of functions of bounded variation, or $BV(\mathbb{R})$.

The total variation of a piecewise constant function is simply the sum of the absolute value of the jumps it undergoes. This property translates to a useful intuition about the total variation of more general functions if we view them as limits of piecewise constant functions.

The important meta-property of the Fourier transform is that

decay for large $|k|$ corresponds to smoothness in x .

There are various degrees to which a function can be smooth or rates at which it can decay, so therefore there are several ways that this assertion can be made precise. Let us go over a few of them. Each assertion either expresses a decay (in k) to smoothness (in x) implication, or the converse implication.

- Let $\hat{f} \in L^1(\mathbb{R})$ (decay), then $f \in L^\infty(\mathbb{R})$ and f is continuous (smoothness). That's because $|e^{ikx}| = 1$, so

$$|f(x)| \leq \frac{1}{2\pi} \int |e^{ikx} \hat{f}(k)| dk = \frac{1}{2\pi} \int |\hat{f}(k)| dk,$$

which proves boundedness. As for continuity, consider a sequence $y_n \rightarrow 0$ and the formula

$$f(x - y_n) = \frac{1}{2\pi} \int e^{ik(x-y_n)} \hat{f}(k) dk.$$

The integrand converges pointwise to $e^{ikx} \hat{f}(k)$, and is uniformly bounded in modulus by the integrable function $|\hat{f}(k)|$. Hence Lebesgue's dominated convergence theorem applies and yields $f(x - y_n) \rightarrow f(x)$, i.e., continuity in x .

- Let $\hat{f}(k)(1 + |k|^p) \in L^1(\mathbb{R})$ (decay). Then $f \in C^p$ (smoothness). We saw the case $p = 0$ above; the justification is analogous in the general case. We write

$$|f^{(n)}(x)| \leq \frac{1}{2\pi} \int |e^{ikx} (ik)^n \hat{f}(k)| dk \leq \int |k|^n |\hat{f}(k)| dk,$$

which needs to be bounded for all $0 \leq n \leq p$. This is obviously the case if $\hat{f}(k)(1 + |k|^p) \in L^1(\mathbb{R})$. Continuity of $f^{(p)}$ is proved like before.

- Let $f \in BV(\mathbb{R})$ (smoothness). Then $\hat{f}(k) \leq \|f\|_{TV} |k|^{-1}$ (decay). If $f \in C^1 \cap BV(\mathbb{R})$, then this is justified very simply from (7.9), and

$$ik\hat{f}(k) = \int e^{-ikx} f'(x) dx.$$

Take a modulus on both sides, and get the desired relation

$$|k|\hat{f}(k) \leq \int |f'(x)|dx = \|f\|_{TV} < \infty.$$

When $f \in BV(\mathbb{R})$, but $f \notin C^1$, either of the more general formulas (7.10) must be used instead. It is a great practice exercise to articulate a modified proof using the $\lim_{h \rightarrow 0}$ formula, and properly pass to the limit.

- Let f such that $f^{(k)} \in L^2(\mathbb{R})$ for $0 \leq k < p$, and assume that $f^{(p)} \in BV(\mathbb{R})$ (smoothness). Then there exists $C > 0$ such that $|\hat{f}(k)| \leq |k|^{-p-1}$ (decay). This claim is also formulated as point (a) in Theorem 1 on page 30 of Trefethen's "Spectral methods in Matlab". The justification is very simple when $f \in C^{p+1}$: we then get

$$(ik)^{p+1}\hat{f}(k) = \int e^{-ikx} f^{(p+1)}(x) dx,$$

so

$$|k|^{p+1}|\hat{f}(k)| \leq \int |f^{(p+1)}(x)| dx = \|f^{(p+1)}\|_{TV} < \infty.$$

Again, it is a good exercise to try and extend this result to functions not in C^{p+1} .

Example 15. *The canonical illustrative example for the two statements involving bounded variation is that of the B-splines. Consider*

$$s(x) = \frac{1}{2}\chi_{[-1,1]}(x),$$

the rectangle-shaped indicator of $[-1, 1]$ (times one half). It is a function in $BV(\mathbb{R})$, but it has no derivative in $L^2(\mathbb{R})$. Accordingly, its Fourier transform $\hat{s}(k)$ is predicted to decay at a rate $\sim |k|^{-1}$. This is precisely the case, for we know

$$\hat{s}(k) = \frac{\sin k}{k}.$$

Functions with higher regularity can be obtained by auto-convolution of s ; for instance $s_2 = s \star s$ is a triangle-shaped function which has one derivative in

L^2 , and such that $s'_2 \in BV(\mathbb{R})$. We anticipate that $\hat{s}_2(k)$ would decay like $|k|^{-2}$, and this the case since by the convolution theorem

$$\hat{s}_2(k) = (\hat{s}(k))^2 = \left(\frac{\sin k}{k}\right)^2.$$

Any number of autoconvolutions $s \star s \dots \star s$ can thus be considered: that's the family of B-splines.

The parallel between smoothness in x and decay in k goes much further. We have shown that p derivatives in x very roughly corresponds to an inverse-polynomial decay $|k|^{-p}$ in k . So C^∞ functions in x have so called super-algebraic decay in k : faster than $C_p|k|^{-p}$ for all $p \geq 0$.

The gradation of very smooth functions goes beyond C^∞ functions to include, in order:

- analytic functions that can be extended to a strip in the complex plane (like $f(x) = 1/(1+x^2)$), corresponding to a Fourier transform decaying exponentially (in this case $\hat{f}(k) = \pi e^{-|k|}$). That's Paley-Wiener theory, the specifics of which is not material for this course.
- analytic functions that can be extended to the whole complex plane with super-exponential growth (like $f(x) = e^{-x^2/2}$), whose Fourier transform decays faster than exponentially (in this case $\hat{f}(k) = \sqrt{\pi/2}e^{-k^2/2}$).
- analytic functions that can be extended to the whole complex plane with exponential growth at infinity (like $f(x) = \frac{\sin x}{x}$), whose Fourier transform is compactly supported (in this case $\hat{f}(k) = 2\pi\chi_{[-1,1]}(k)$). That's Paley-Wiener theory in reverse. Such functions are also called bandlimited.

More on this in Chapter 4 of Trefethen's book.

An important consequence of a Fourier transform having fast decay is that it can be truncated to some large interval $k \in [-N, N]$ at the expense of an error that decays fast as a function of N . The smoother f , the faster \hat{f} decays, and the more accurate the truncation to large N in frequency. On the other hand, there may be convergence issues if $f(x)$ is not smooth.

To make this quantitative, put

$$\hat{f}_N(k) = \chi_{[-N,N]}(k)\hat{f}(k).$$

Recall that $\frac{1}{2\pi} \int_{-N}^N e^{ikx} dx = \frac{\sin(Nx)}{\pi x}$. By the convolution theorem, we therefore have

$$f_N(x) = \frac{\sin Nx}{\pi x} \star f(x).$$

In the presence of $f \in L^2(\mathbb{R})$, letting $N \rightarrow \infty$ always gives rise to convergence $f_N \rightarrow f$ in $L^2(\mathbb{R})$. This is because by Plancherel's identity,

$$\|f - f_N\|_2^2 = \frac{1}{2\pi} \int_{-\infty}^{\infty} |\hat{f}_N(k) - \hat{f}(k)|^2 dk = \frac{1}{2\pi} \int_{|k|>N} |\hat{f}(k)|^2 dk.$$

This quantity tends to zero as $N \rightarrow \infty$ since the integral $\int_{\mathbb{R}} |\hat{f}(k)|^2$ over the whole line is bounded.

The story is quite different if we measure convergence in L^∞ instead of L^2 , when f has a discontinuity. Generically, L^∞ (called uniform) convergence fails in this setting. This phenomenon is called *Gibb's effect*, and manifests itself through ripples in reconstructions from truncated Fourier transforms. Let us illustrate this on the Heaviside step function

$$u(x) = \begin{cases} 1 & \text{if } x \geq 0; \\ 0 & \text{if } x < 0. \end{cases}$$

(It is not a function in $L^2(\mathbb{R})$ but this is of no consequence to the argument. It could be modified into a function in $L^2(\mathbb{R})$ by some adequate windowing.) Since $u(x)$ is discontinuous, we expect the Fourier transform to decay quite slowly. Consider the truncated FT

$$\hat{u}_N(k) = \chi_{[-N,N]}(k) \hat{u}(k).$$

Back in the x domain, this is

$$\begin{aligned} u_N(x) &= \frac{\sin Nx}{\pi x} \star u(x) \\ &= \int_0^\infty \frac{\sin N(x-y)}{\pi(x-y)} dy \\ &= \int_{-\infty}^{Nx} \frac{\sin y}{\pi y} dy \\ &\equiv s(Nx). \end{aligned}$$

(Draw picture)

The function $s(Nx)$ is a sigmoid function going from 0 at $-\infty$ to 1 at ∞ , going through $s(0) = 1/2$, and taking on min. value $\simeq -0.045$ at $x = -\pi/N$ and max. value $\simeq 1.045$ at $x = \pi/N$. The oscillations have near-period π/N . The parameter N is just a dilation index for the sigmoid, it does not change the min. and max. values of the function. Even as $N \rightarrow \infty$, there is no way that the sigmoid would converge to the Heaviside step in a uniform manner: we always have

$$\|u - u_N\|_\infty \gtrsim .045.$$

This example of Gibb's effect goes to show that truncating Fourier expansions could be a poor numerical approximation if the function is nonsmooth.

However, if the function is smooth, then everything goes well. Let us study the decay of the approximation error

$$\epsilon_N^2 = \|f - f_N\|_2^2$$

for truncation of the FT to $[-N, N]$. Again, we use Plancherel to get

$$\epsilon_N^2 = \int_{|k|>N} |\hat{f}(k)|^2 dk$$

Now assume that $\hat{f}(k) \leq C|k|^{-p}$, a scenario already considered earlier. Then

$$\epsilon_N^2 \leq C \int_{|k|>N} |k|^{-2p} \leq C' N^{-2p+1},$$

so, for some constant C'' (dependent on p but independent of N , hence called constant),

$$\epsilon_N \leq C'' N^{-p+1/2}.$$

The larger p , the faster the decay as $N \rightarrow \infty$.

When a function is C^∞ , the above decay rate of the approximation error is valid for all $p > 0$. When the function is analytic and the Fourier transform decays exponentially, the decay rate of the approximation error is even faster, itself exponential (a good exercise). In numerical analysis, either such behavior is called *spectral accuracy*.

7.5 Chebyshev expansions

In the previous sections, we have seen that smooth functions on the real line have fast decaying Fourier transforms.

On a finite interval, a very similar property hold for Fourier series: if a function is smooth in $[-\pi, \pi]$ and connects smoothly by periodicity at $x = -\pi$ and $x = \pi$, then its Fourier series (7.5) decays fast. Periodicity is essential, because the points $x = -\pi$ and $x = \pi$ play no particular role in (7.5). They might as well be replaced by $x = 0$ and $x = 2\pi$ for the integration bounds. So if a function is to qualify as smooth, it has to be equally smooth at $x = 0$ as it is at $x = -\pi$, identified with $x = \pi$ by periodicity.

For instance if a function f is smooth inside $[-\pi, \pi]$, but has $f(-\pi) \neq f(\pi)$, then for the purpose of convergence of Fourier series, f is considered discontinuous. We know what happens in this case: the Gibbs effect takes place, partial inverse Fourier series have unwelcome ripples, and convergence does not occur in L^∞ .

If f is smooth and periodic, then it is a good exercise to generalize the convergence results of the previous section, from the Fourier transform to Fourier series.

How shall we handle smooth functions in intervals $[a, b]$, which do not connect smoothly by periodicity? The answer is not unique, but the most standard tool for this in numerical analysis are the *Chebyshev polynomials*.

For simplicity consider $x \in [-1, 1]$, otherwise rescale the problem. Take a C^k nonperiodic $f(x)$ in $[-1, 1]$. The trick is to view it as $g(\theta) = f(\cos \theta)$. Since $\cos[0, \pi] = \cos[\pi, 2\pi] = [-1, 1]$, all the values of $x \in [-1, 1]$ are covered twice by $\theta \in [0, 2\pi]$. Obviously, at any point θ , g inherits the smoothness of f by the chain rule. Furthermore, g is periodic since $\cos \theta$ is periodic. So g is exactly the kind of function which we expect should have a fast converging Fourier series:

$$\hat{g}_k = \int_0^{2\pi} e^{-ik\theta} g(\theta) d\theta, \quad g(\theta) = \frac{1}{2\pi} \sum_{k=-\infty}^{\infty} e^{ik\theta} \hat{g}_k, \quad k \in \mathbb{Z}.$$

Since $g(\theta)$ is even in θ , we may drop the $\sin \theta$ terms in the expansion, as well as the negative k :

$$\hat{g}_k = \int_0^{2\pi} \cos(k\theta) g(\theta) d\theta, \quad g(\theta) = \frac{1}{2\pi} \hat{g}_0 + \frac{1}{\pi} \sum_{k=1}^{\infty} \cos(k\theta) \hat{g}_k, \quad k \in \mathbb{Z}^+.$$

Back to f , we find

$$\hat{g}_k = 2 \int_{-1}^1 \cos(k \arccos x) f(x) \frac{dx}{\sqrt{1-x^2}}, \quad f(x) = \frac{1}{2\pi} \hat{g}_0 + \frac{1}{\pi} \sum_{k=1}^{\infty} \cos(k \arccos x) \hat{g}_k, \quad k \in \mathbb{Z}^+$$

The function $\cos(k \arccos x)$ happens to be a polynomial in x , of order k , called the *Chebyshev polynomial* of order k . Switch to the letter n as is usually done:

$$T_n(x) = \cos(n \arccos x), \quad x \in [-1, 1].$$

The first few Chebyshev polynomials are

$$T_0(x) = 1, \quad T_1(x) = x, \quad T_2(x) = 2x^2 - 1.$$

It is a good exercise (involving trigonometric identities) to show that they obey the recurrence relation

$$T_{n+1}(x) = 2xT_n(x) - T_{n-1}(x).$$

The orthogonality properties of $T_n(x)$ over $[-1, 1]$ follows from those of $\cos(n\theta)$ over $[0, \pi]$, but watch the special integration weight:

$$\int_{-1}^1 T_m(x) T_n(x) \frac{dx}{\sqrt{1-x^2}} = c_n \delta_{mn},$$

with $c_n = \pi/2$ if $n = 0$ and $c_n = \pi$ otherwise. The Chebyshev expansion of f is therefore

$$\langle f, T_n \rangle = \int_{-1}^1 T_n(x) f(x) \frac{dx}{\sqrt{1-x^2}}, \quad f(x) = \frac{1}{\pi} \langle f, T_0 \rangle + \frac{2}{\pi} \sum_{n=1}^{\infty} T_n(x) \langle f, T_n \rangle, \quad k \in \mathbb{Z}^+.$$

Under the hood, this expansion is the FS of a periodic function, so we can apply results pertaining to Fourier series and Fourier transforms to obtain fast decay of Chebyshev expansions. This way, spectral accuracy is restored for nonperiodic functions defined on an interval.

Chapter 8

Spectral Interpolation, Differentiation, Quadrature

8.1 Interpolation

8.1.1 Bandlimited interpolation

While equispaced points generally cause problems for polynomial interpolation, as we just saw, they are the natural choice for discretizing the Fourier transform. For data on $x_j = jh$, $j \in \mathbb{Z}$, recall that the semidiscrete Fourier transform (SFT) and its inverse (ISFT) read

$$\hat{f}(k) = h \sum_{j \in \mathbb{Z}} e^{-ikx_j} f_j, \quad f_j = \frac{1}{2\pi} \int_{-\pi/h}^{\pi/h} e^{ikx_j} \hat{f}(k) dk.$$

The idea of spectral interpolation, or bandlimited interpolation, is to evaluate the ISFT formula above at some point x not equal to one of the x_j .

Definition 17. (*Spectral interpolation on \mathbb{R}*) Let $x_j = jh$, $j \in \mathbb{Z}$. Consider $f : \mathbb{R} \mapsto \mathbb{R}$, its restriction $f_j = f(x_j)$, and the SFT $\hat{f}(k)$ of the samples f_j . Then the spectral interpolant is

$$p(x) = \frac{1}{2\pi} \int_{-\pi/h}^{\pi/h} e^{ikx} \hat{f}(k) dk.$$

We can view the formula for $p(x)$ as the inverse Fourier transform of the compactly supported function equal to $\hat{f}(k)$ for $k \in [-\pi/h, \pi/h]$, and zero

otherwise. When the Fourier transform of a function is compactly supported, we say that that function is *bandlimited*, hence the name of the interpolation scheme.

Example 16. *Let*

$$f_j = \delta_{0j} = \begin{cases} 1 & \text{if } j = 0; \\ 0 & \text{if } j \neq 0. \end{cases}$$

Then the SFT is $\hat{f}(k) = h$ for $k \in [-\pi/h, \pi/h]$ as we saw previously. Extend it to $k \in \mathbb{R}$ by zero outside of $[-\pi/h, \pi/h]$. Then

$$p(x) = \frac{\sin(\pi x/h)}{\pi x/h} = \text{sinc}(\pi x/h).$$

This function is also called the Dirichlet kernel. It vanishes at $x_j = jh$ for $j \neq 0$, integer.

Example 17. *In full generality, consider now the sequence*

$$f_j = \sum_{k \in \mathbb{Z}} \delta_{jk} f_k.$$

By linearity of the integral,

$$p(x) = \sum_{k \in \mathbb{Z}} f_k \text{sinc}(\pi(x - x_k)/h).$$

The interpolant is a superposition of sinc functions, with the samples f_j as weights. Here sinc is the analogue of the Lagrange elementary polynomials of the previous section, and is called the interpolation kernel. For this reason, bandlimited interpolation sometimes goes by the name Fourier-sinc interpolation.

(Figure here for the interpolation of a discrete step.)

In the example above we interpolate a discontinuous function, and the result is visually not very good. It suffers from the same Gibbs effect that we encountered earlier. The smoother the underlying $f(x)$ which f_j are the samples of, however, the more accurate bandlimited interpolation.

In order to study the approximation error of bandlimited interpolation, we need to return to the link between SFT and FT. The relationship between $p(x)$ and f_j is *sampling*, whereas the relationship between the FT

$\hat{f}(k)\chi_{[-\pi/h, \pi/h]}(k)$ and the SFT $\hat{f}(k)$ is *periodization*. We have already alluded to this correspondence earlier, and it is time to formulate it more precisely.

(Figure here; sampling and periodization)

Theorem 12. (*Poisson summation formula, FT version*) Let $u : \mathbb{R} \mapsto \mathbb{R}$, sufficiently smooth and decaying sufficiently fast at infinity. (We are deliberately imprecise!) Let $v_j = u(x_j)$ for $x_j = jh$, $j \in \mathbb{Z}$, and

$$\hat{u}(k) = \int_{\mathbb{R}} e^{-ikx} u(x) dx, \quad (FT), \quad k \in \mathbb{R},$$

$$\hat{v}(k) = h \sum_{j \in \mathbb{Z}} e^{-ikx_j} u(x_j). \quad (SFT), \quad k \in [-\pi/h, \pi/h].$$

Then

$$\hat{v}(k) = \sum_{m \in \mathbb{Z}} \hat{u}(k + m \frac{2\pi}{h}), \quad k \in [-\pi/h, \pi/h] \quad (8.1)$$

In some texts the Poisson summation formula is written as the special case $k = 0$:

$$h \sum_{j \in \mathbb{Z}} u(jh) = \sum_{m \in \mathbb{Z}} \hat{u}(\frac{2\pi}{h}m).$$

By what we know concerning translations and Fourier transforms, this is equivalent to (8.1).

Proof. Consider the right-hand side in (8.1), and call it

$$\hat{\phi}(k) = \sum_{m \in \mathbb{Z}} \hat{u}(k + m \frac{2\pi}{h}), \quad k \in [-\pi/h, \pi/h].$$

It suffices to show that $\hat{\phi}(k) = \hat{v}(k)$, or equivalently in terms of their ISFT, that $\phi_j = v_j$ for $j \in \mathbb{Z}$. The ISFT is written

$$\phi_j = \frac{1}{2\pi} \int_{-\pi/h}^{\pi/h} \left[\sum_{m \in \mathbb{Z}} \hat{u}(k + m \frac{2\pi}{h}) \right] e^{ikjh} dk.$$

The function u is smooth, hence integrable, and the sum over m converges fast. So we can interchange sum and integral:

$$\phi_j = \frac{1}{2\pi} \sum_{m \in \mathbb{Z}} \int_{-\pi/h}^{\pi/h} \hat{u}(k + m \frac{2\pi}{h}) e^{ikjh} dk.$$

Now put $k' = k + m\frac{2\pi}{h}$, and change variable:

$$\phi_j = \frac{1}{2\pi} \sum_{m \in \mathbb{Z}} \int_{-\frac{\pi}{h} - m\frac{2\pi}{h}}^{\frac{\pi}{h} - m\frac{2\pi}{h}} \hat{u}(k') e^{ik'jh} e^{-i\frac{2\pi}{h}jh} dk'.$$

The extra exponential factor $e^{-i\frac{2\pi}{h}jh}$ is equal to 1 because $j \in \mathbb{Z}$. We are in presence of an integral over \mathbb{R} chopped up into pieces corresponding to sub-intervals of length $2\pi/h$. Piecing them back together, we get

$$\phi_j = \frac{1}{2\pi} \int_{\mathbb{R}} \hat{u}(k') e^{ik'jh} dk',$$

which is exactly the inverse FT of \hat{u} evaluated at $x_j = jh$, i.e., $\phi_j = u(x_j) = v_j$. \square

The Poisson summation formula shows that sampling a function at rate h corresponds to periodizing its spectrum (Fourier transform.) with a period $2\pi/h$. So the error made in sampling a function (and subsequently doing bandlimited interpolation) is linked to the possible overlap of the Fourier transform upon periodization.

- **Scenario 1.** Assume $\text{supp}(\hat{u}) \subset [-\pi/h, \pi/h]$. Then no error is made in sampling and interpolating u at rate h , because nothing happens upon $2\pi/h$ -periodization and windowing into $[-\pi/h, \pi/h]$:

$$\hat{p}(k) = \hat{u}(k) \quad \Rightarrow \quad p(x) = u(x).$$

(Draw picture)

- **Scenario 2.** Now assume that $\hat{u}(k)$ is not included in $[-\pi/h, \pi/h]$. In general the periodization of $\hat{u}(k)$ will result in some overlap inside $[-\pi/h, \pi/h]$. We call this *aliasing*. In that case, some information is lost and interpolation will not be exact.

Scenario 1 is known as the *Shannon sampling theorem*: a function bandlimited in $[-\pi/h, \pi/h]$ in k space is perfectly interpolated by bandlimited interpolation, on a grid of spacing h or greater. In signal processing h is also called the sampling rate, because x has the interpretation of time. When h is the largest possible rate such that no aliasing occurs, it can be referred to as the Nyquist rate.

More generally, the accuracy of interpolation is linked to the smoothness of $u(x)$. If the tails $\hat{u}(k)$ are small and h is large, we can expect that the error due to periodization and overlap won't be too big. The following result is an immediate consequence of the Poisson summation formula.

Theorem 13. (*Error of bandlimited interpolation*) Let $u \in L^2(\mathbb{R})$ such that u has $p - 1$ derivatives in L^2 , and $u^{(p)} \in BV(\mathbb{R})$. Denote by $\hat{u}(k)$ the FT of u . Let $v_j = u(x_j)$ at $x_j = jh$, $j \in \mathbb{Z}$. Denote by $\hat{v}(k)$ the SFT of v_j . Then

$$|\hat{v}(k) - \hat{u}(k)| = O(h^{p+1}) \quad \text{as } h \rightarrow 0.$$

Proof. By the Poisson summation formula (8.1),

$$\hat{v}(k) - \hat{u}(k) = \sum_{m \neq 0} \hat{u}(k + m \frac{2\pi}{h}), \quad k \in [-\pi/h, \pi/h].$$

As we saw earlier, the smoothness conditions on u imply that

$$|\hat{u}(k)| \leq C |k|^{-p-1}.$$

Since

$$k + m \frac{2\pi}{h} \in [-\frac{\pi}{h} + m \frac{2\pi}{h}, \frac{\pi}{h} + m \frac{2\pi}{h}],$$

we have

$$|\hat{u}(k + m \frac{2\pi}{h})| \leq C' |m \frac{\pi}{h}|^{-p-1},$$

for some different constant C' . Summing over $m \neq 0$,

$$|\hat{v}(k) - \hat{u}(k)| \leq C' \sum_{m \neq 0} |m|^{-p-1} (\frac{\pi}{h})^{-p-1} \leq C'' (\frac{\pi}{h})^{-p-1} \leq C''' h^{p+1}.$$

□

In the above theorem, the error of interpolation is expressed in the Fourier domain, but of course one can switch back to the x domain by means of Plancherel formula

$$\|u(jh) - v_j\|_{\ell_2(j)} = \frac{1}{\sqrt{2\pi}} \|\hat{u}(k) - \hat{v}(k)\|_{L^2[-\pi/h, \pi/h]}.$$

We have seen how to interpolate a function defined on \mathbb{R} , but as a closing remark let us notice that a similar notion exists for functions defined on

intervals, notably $x \in [-\pi, \pi]$ or $[0, 2\pi]$. In that case, wavenumbers are discrete, the FT is replaced by the FS, and the SFT is replaced by the DFT. Evaluating the DFT for x not on the grid would give an interpolant:

$$\frac{1}{2\pi} \sum_{k=-N/2+1}^{N/2} e^{ikx} \hat{f}_k.$$

Contrast with the formula (7.8) for the IDFT. This is almost what we want, but not quite, because the highest wavenumber $k = N/2$ is treated asymmetrically. It gives rise to an unnatural complex term, even if \hat{f}_k is real and even. To fix this, it is customary to set $\hat{f}_{-N/2} = \hat{f}_{N/2}$, to extend the sum from $-N/2$ to $N/2$, but to halve the terms corresponding to $k = -N/2$ and $N/2$. We denote this operation of halving the first and last term of a sum by a prime after the sum symbol:

$$\sum'$$

It is easy to check that this operation does not change the interpolating property. The definition of bandlimited interpolant becomes the following in the case of intervals.

Definition 18. (*Spectral interpolation on $[0, 2\pi]$*) Let $x_j = jh$, $j = 0, \dots, N-1$ with $h = 1/N$. Consider $f : [0, 2\pi] \mapsto \mathbb{R}$, its restriction $f_j = f(x_j)$, and the DFT \hat{f}_k of the samples f_j . Then the spectral interpolant is

$$p(x) = \frac{1}{2\pi} \sum'_{k=-N/2}^{N/2} e^{ikx} \hat{f}_k.$$

The atoms e^{ikx} for k integer are called *trigonometric polynomials*.

The theory for interpolation by trigonometric polynomials (inside $[0, 2\pi]$) is very similar to that for general bandlimited interpolants. The only important modification is that $[0, 2\pi]$ is a periodized interval, so a function qualifies as smooth only if it connects smoothly across $x = 0$ identified with $x = 2\pi$ by periodicity. The Poisson summation formula is still the central tool, and has a counterpart for Fourier series.

Theorem 14. (*Poisson summation formula, FS version*) Let $u : [0, 2\pi] \mapsto \mathbb{R}$, sufficiently smooth. Let $v_j = u(\theta_j)$ for $\theta_j = jh$, $j = 1, \dots, N$, $h = 2\pi/N$, and

$$\hat{u}_k = \int_0^{2\pi} e^{-ik\theta} u(\theta) d\theta, \quad (FS), \quad k \in \mathbb{Z},$$

$$\hat{v}_k = h \sum_{j=1}^N e^{-ik\theta_j} u(\theta_j). \quad (\text{DFT}), \quad k = -\frac{N}{2}, \dots, \frac{N}{2} - 1.$$

Then

$$\hat{v}_k = \sum_{m \in \mathbb{Z}} \hat{u}_{k+mN}, \quad k = -\frac{N}{2}, \dots, \frac{N}{2} - 1. \quad (8.2)$$

(Recall that $N = \frac{2\pi}{h}$ so this formula is completely analogous to (8.1).)

The unsurprising corollary is that if u has p derivatives, with a p -th derivative in BV , over the periodized $[0, 2\pi]$, then the bandlimited interpolation error is a $O(h^{p+1})$ in the ℓ_2 sense in k space.

Otherwise, if $u(\theta)$ has discontinuities like e.g. $u(2\pi^+) \neq u(0^-)$, interpolation will suffer from the Gibbs effect.

8.1.2 Chebyshev interpolation

Consider now smooth functions inside $[-1, 1]$ (for illustration), but not necessarily periodic. So the periodization $\sum_j f(x + 2j)$ may be discontinuous. Polynomial interpolation on equispaced points may fail because of the Runge phenomenon, and bandlimited interpolation will fail due to Gibbs's effect.

A good strategy for interpolation is the same trick as the one we used for truncation in the last chapter: pass to the variable θ such that

$$x = \cos \theta.$$

Then $x \in [-1, 1]$ corresponds to $\theta \in [0, \pi]$. We define $g(\theta) = f(\cos \theta)$ with g 2π -periodic and even.

We can now consider the bandlimited interpolant of $g(\theta)$ on an equispaced grid covering $[0, 2\pi]$, like for instance

$$\theta_j = \frac{\pi j}{N}, \quad \text{where } j = 1, \dots, 2N.$$

Using the definition we saw at the end of the last section (with the "sum prime"), we get

$$q(\theta) = \frac{1}{2\pi} \sum'_{k=-N}^N e^{ik\theta} \hat{g}_k, \quad \hat{g}_k = \frac{\pi}{N} \sum_{j=0}^{2N-1} e^{-ik\theta_j} g(\theta_j).$$

By even symmetry in θ and k (why?), we can write

$$q(\theta) = \sum_{k=0}^N \cos(k\theta)c_k,$$

for some coefficients c_k that are formed from the samples $g(\theta_j)$.

Back to x , we get the sample points $x_j = \cos(\theta_j)$. They are called *Chebyshev points*. They are not equispaced anymore, and because they are the projection on the x -axis of equispaced points on the unit circle, they cluster near the edges of $[-1, 1]$. There are $N + 1$ of them, from $j = 0$ to N , because of the symmetry in θ . It turns out that the x_j are the *extremal points* of the Chebyshev polynomial $T_N(x)$, i.e., the points in $[-1, 1]$ where T_N takes its maximum and minimum values.

In terms of the variable x , the interpolant can be expressed as

$$p(x) = q(\operatorname{acos} x) = \sum_{n=0}^N T_n(x)c_n, \quad T_n(x) = \cos(n \operatorname{acos} x),$$

with the same c_n as above. Since $T_n(x)$ are polynomials of degree $0 \leq n \leq N$, and $p(x)$ interpolates $f(x)$ at the $N + 1$ (non-equispaced) points x_j , we are in presence of the Lagrange interpolation polynomial for f at x_j !

Now, the interesting conclusion is not so much the new formula for the interpolation polynomial in terms of T_n , but that although this is a polynomial interpolant, the error analysis is inherited straight from Fourier analysis. If f is smooth, then g is smooth and periodic, and we saw that the bandlimited interpolant converges very fast. The Chebyshev interpolant of f is equal to the bandlimited interpolant of g so it converges at the exact same rate (in L^∞ in k or n space) — for instance $O(N^{-p-1})$ when f has p derivatives (in BV).

In particular, we completely bypassed the standard analysis of error of polynomial interpolation, and proved universal convergence for smooth f . The factor

$$\pi_{N+1}(x) = \prod_{j=0}^N (x - x_j)$$

that was posing problems in the error estimate then does not pose a problem anymore, because of the very special choice of Chebyshev points $\cos(\pi j/N)$ for the interpolation. Intuitively, clustering the grid points near the edges of the interval $[-1, 1]$ helps giving $\pi_{N+1}(x)$ more uniform values throughout $[-1, 1]$, hence reduces the gross errors near the edges.

Let us now explain the differences in the behavior of the monic polynomial $\prod_{j=0}^N (x - x_j)$ for equispaced vs. Chebyshev points, and argue that Chebyshev points are near-ideal for interpolation of smooth functions in intervals. The discussion below is mostly taken from Trefethen, p.43. (See also the last problem on homework 2 for an example of different analysis.)

Let $p(z) = \prod_{j=0}^N (z - x_j)$, where we have extended the definition of the monic polynomial to $z \in \mathbb{C}$. We compute

$$\log |p(z)| = \sum_{j=0}^N \log |z - x_j|,$$

Put

$$\phi_N(z) = (N + 1)^{-1} \sum_{j=0}^N \log |z - x_j|.$$

The function ϕ_N is like an electrostatic potential, due to charges at $z = x_j$, each with potential $(N + 1)^{-1} \log |z - x_j|$. Going back to $p(z)$ from ϕ_N is easy:

$$p(z) = e^{(N+1)\phi_N(z)}.$$

Already from this formula, we can see that small variations in ϕ_N will lead to exponentially larger variations in $p(z)$, particularly for large N .

Let us now take a limit $N \rightarrow \infty$, and try to understand what happens without being too rigorous. What matters most about the Chebyshev points is their *density*: the Chebyshev points are the projection onto the real-axis of a sequence of equispaced points on a circle. If the density of points on the circle is a constant $1/(2\pi)$, then the density of points on $[-1, 1]$ generated by vertical projection is

$$\rho_{Cheb}(x) = \frac{1}{\pi\sqrt{1-x^2}}. \quad (\text{normalized to integrate to 1 on } [-1, 1])$$

(This is a density in the sense that

$$N \int_a^b \rho_{Cheb}(x) dx$$

approximately gives the number of points in $[a, b]$.) Contrast with a uniform distribution of points, with density

$$\rho_{equi}(x) = \frac{1}{2}.$$

Then the potential corresponding to any given $\rho(x)$ is simply

$$\phi(z) = \int_{-1}^1 \rho(x) \log |z - x| dx.$$

The integral can be solved explicitly for both densities introduced above:

- For ρ_{equi} , we get

$$\phi_{equi}(z) = -1 + \frac{1}{2} \operatorname{Re}((z+1) \log(z+1) - (z-1) \log(z-1)).$$

It obeys $\phi_{equi}(0) = -1$, $\phi_{equi}(\pm 1) = -1 + \log 2$.

- For ρ_{Cheb} , we get

$$\phi_{Cheb}(z) = \log \frac{|z - \sqrt{z^2 - 1}|}{2}.$$

This function obeys (interesting exercise) $\phi_{Cheb}(x) = -\log 2$ for all $x \in [-1, 1]$ on the real axis.

The level curves of both ϕ_{equiv} and ϕ_{Cheb} in the complex plane are shown on page 47 of Trefethen.

Overlooking the fact that we have passed to a continuum limit for the potentials, we can give a precise estimate on the monic polynomial:

$$|p_{equi}(z)| \simeq e^{(N+1)\phi_{equi}(z)} = \begin{cases} (2/e)^N & \text{near } x = \pm 1; \\ (1/e)^N & \text{near } x = 0. \end{cases}$$

whereas

$$|p_{Cheb}(z)| \simeq e^{(N+1)\phi_{Cheb}(z)} = 2^{-N}, \quad z \in [-1, 1].$$

We see that p_{equi} can take on very different values well inside the interval vs. near the edges. On the other hand the values of p_{Cheb} are near-constant in $[-1, 1]$. The density $\rho_{Cheb}(x)$ is the only one that will give rise to this behavior, so there is something special about it. It is the difference in asymptotic behavior of $(2/e)^{-N}$ vs. 2^{-N} that makes the whole difference for interpolation, as $N \rightarrow \infty$.

One may argue that neither p_{equi} nor p_{Cheb} blow up as $N \rightarrow \infty$, but it is an interesting exercise to show that if we were interpolating in an interval

$[-a, a]$ instead of $[-1, 1]$, then the bounds would be multiplied by a^N , by homogeneity.

The Chebyshev points are not the only one that correspond to the density $\rho_{Cheb}(x)$ as $N \rightarrow \infty$. For instance, there is also the Chebyshev roots

$$\theta'_j = \frac{\pi}{2N} + \frac{\pi j}{N}, \quad j = 0, \dots, 2N - 1,$$

which are the *roots* of $T_N(x)$, instead of being the extremal points. They give rise to very good interpolation properties as well.

Finally, let us mention that the theory can be pushed further, and that the exponential rate of convergence of Chebyshev interpolation for analytic functions can be linked to the maximum value of $\phi(z)$ on the strip in which the extension $f(z)$ of $f(x)$ is analytic. We will not pursue this further.

8.2 Differentiation

8.2.1 Bandlimited differentiation

The idea that a numerical approximation of a derivative can be obtained from differentiating an interpolant can be pushed further. In this section we return to bandlimited interpolants. We've seen that they are extremely accurate when the function is smooth and periodic; so is the resulting differentiation scheme. It is called bandlimited differentiation, or spectral differentiation.

First consider the case of $x \in \mathbb{R}$ and $x_j = jh$, $j \in \mathbb{Z}$. As we've seen, the bandlimited/spectral interpolant of $u(x_j)$ is

$$p(x) = \frac{1}{2\pi} \int_{-\pi/h}^{\pi/h} e^{ikx} \hat{u}(k) dk,$$

where $\hat{v}(k)$ is the SFT of $v_j = u(x_j)$. Differentiating $p(x)$ reduces to a multiplication by ik in the Fourier domain. Evaluating $p'(x_j)$ is then just a matter of letting $x = x_j$ in the resulting formula. The sequence of steps for bandlimited differentiation ($x \in \mathbb{R}$) is the following:

- Obtain the SFT $\hat{v}(k)$ of $v_j = u(x_j)$;
- Multiply $\hat{v}(k)$ by ik ;
- Obtain the ISFT of $\hat{w}(k) = ik\hat{v}(k)$, call it w_j .

The numbers w_j obtained above are an approximation of $u'(x_j)$. The following result makes this precise.

Theorem 15. (*Accuracy of bandlimited differentiation, see also Theorem 4 in Trefethen's book*) Let $u, \dots, u^{(p-1)} \in L^2(\mathbb{R})$, and $u^{(p)} \in BV(\mathbb{R})$. Let $v_j = u(x_j)$, and w_j the result of bandlimited differentiation. Then

$$\sup_j |w_j - u'(x_j)| = O(h^{p-1}).$$

Proof. We did the proof in the homework. It hinges on the fact that

$$|\hat{v}(k) - \hat{u}(k)| = O(h^{p+1}).$$

One power of h is lost when differentiating u , and one power of h is lost in going back to the physical domain (j instead of k .) \square

The point of the above theorem is that the order of bandlimited differentiation is directly linked to the smoothness of the function, and can be arbitrarily large. This is called spectral accuracy. One can even push the analysis further and show that, when f is real-analytic, then the rate of convergence of w_j towards $u'(x_j)$ is in fact exponential/geometric.

Of course in practice we never deal with a function $u(x)$ defined on the real line. In order to formulate an algorithm, and not simply a sequence of abstract steps, we need to limit the interval over which u is considered. Spectral differentiation in the periodic interval $[0, 2\pi]$ works like before, except DFT are substituted for SFT. For $\theta \in [0, 2\pi]$, we've seen that the spectral interpolant is defined as

$$p(\theta) = \frac{1}{2\pi} \sum'_{k=-N/2}^{N/2} e^{ik\theta} \hat{f}_k.$$

(The sum prime is important here.) Again, a derivative can be implemented by multiplying by ik in the Fourier domain. The sequence of steps is very similar to what it was before, except that we can now label them as “compute”, and not just “obtain”:

- Compute the DFT \hat{v}_k of $v_j = u(x_j)$;
- Multiply \hat{v}_k by ik ;
- Compute the IDFT of $\hat{w}_k = ik\hat{v}_k$, call it w_j .

The result of accuracy are the same as before, with the provision that u needs to be not only smooth, but also smooth when extended by periodicity.

The FFT can be used to yield a fast $O(N \log N)$ algorithm for spectral differentiation.

Note that higher derivatives are obtained in the obvious manner, by multiplying in Fourier by the adequate power of ik .

8.2.2 Chebyshev differentiation

In view of what has been covered so far, the idea of Chebyshev differentiation is natural: it is simply differentiation of the Chebyshev interpolant at the Chebyshev nodes. It proves very useful for those smooth functions on an interval, which do not necessarily extend smoothly by periodicity.

Let us recall that the Chebyshev interpolant is $q(x) = p(\arccos x)$, where $p(\theta)$ is the bandlimited interpolant of $u(\cos \theta_j)$ at the Chebyshev points $x_j = \cos \theta_j$. As such, a differentiation on $q(x)$ is not the same thing as a differentiation on $p(\theta)$. Instead, by the chain rule,

$$q'(x) = \frac{-1}{\sqrt{1-x^2}} p'(\arccos x).$$

The algorithm is as follows. Start from the knowledge of $u(x_j)$ at $x_j = \cos \theta_j$, $\theta_j = jh$, $h = \pi/N$, and $j = 1, \dots, N$.

- Perform an even extension of $u(x_j)$ to obtain $u(\cos \theta_j)$ for $\theta_j = jh$, $h = \pi/N$, and $j = -N + 1, \dots, N$. Now we have all the equispaced sample of the periodic function $u(\cos \theta)$ for θ_j covering $[0, 2\pi]$, not just $[0, \pi]$.
- Take the DFT of those samples, call it \hat{v}_k ,
- Multiply by ik ,
- Take the IDFT of the result $ik\hat{v}_k$,
- Multiply by $-1/\sqrt{1-x_j^2}$ to honor the chain rule. At the endpoints $x_j = -1$ or 1 , take a special limit to obtain the proper values of $p'(-1)$ and $p(1)$. See Trefethen's book for the correct values.

This is a fast algorithm since we can use the FFT for the DFT and IDFT.

Since we are only a change of variables away from Fourier analysis in a periodic domain, the accuracy of Chebyshev differentiation is directly inherited from that of bandlimited differentiation. We also have spectral accuracy.

Note that higher derivatives can be treated similarly, by applying the chain rule repeatedly.

In practice, Chebyshev methods are particularly useful for boundary-value problems (we'll come back to this), when all samples of a function are to be determined at once, and when we have the freedom of choosing the sample points.

8.3 Integration

8.3.1 Spectral integration

To go beyond methods of finite order, the idea of spectral integration is to integrate a bandlimited interpolant. This strategy yields very high accuracy when the function is smooth and periodic.

Consider a function $u(\theta)$, $\theta \in [0, 2\pi]$. Its samples are $v_j = u(\theta_j)$ with $\theta_j = jh$, $h = 2\pi/N$, and $j = 1, \dots, N$. Form the DFT:

$$\hat{v}_k = h \sum_{j=1}^N e^{-ik\theta_j} v_j,$$

and the bandlimited interpolant,

$$p(\theta) = \frac{1}{2\pi} \sum_{k=-N/2}^{N/2} e^{ik\theta} \hat{v}_k.$$

Integrating $p(\theta)$ gives the remarkably simple following result.

$$\int_0^{2\pi} p(\theta) d\theta = \hat{v}_0 = h \sum_{j=1}^N u(\theta_j).$$

We are back to the trapezoidal rule! (The endpoints are identified, $\theta_0 = \theta_N$.) While we have already encountered this quadrature rule earlier, it is now derived from Fourier analysis. So the plot thickens concerning its accuracy properties.

Specifically, we have to compare \hat{v}_0 to $\hat{u}_0 = \int_0^{2\pi} u(\theta) d\theta$, where \hat{u}_k are the Fourier series coefficients of $u(\theta)$. The relationship between \hat{v}_0 and \hat{u}_0 is the Poisson summation formula (8.2):

$$\hat{v}_0 = \sum_{m \in \mathbb{Z}} \hat{u}_{mN}, \quad N = \frac{2\pi}{h}.$$

The most important term in this sum is \hat{u}_0 for $m = 0$, and our task is again to control the other ones, for $m \neq 0$. The resulting accuracy estimate is the following.

Theorem 16. *Assume $u^{(k)} \in L^2[0, 2\pi]$ for $0 \leq k < p$, and $u^{(p)} \in BV[0, 2\pi]$, where $[0, 2\pi]$ is considered a periodic interval. (So it matters that the function connects smoothly by periodicity.) Then*

$$\int_0^{2\pi} u(\theta) d\theta - h \sum_{j=1}^N u(\theta_j) = O(h^{p+1}).$$

Proof. By the Poisson summation formula, in the notations of the preceding few paragraphs,

$$\hat{u}_0 - \hat{v}_0 = \sum_{m \neq 0} \hat{u}_{mN}.$$

We have already seen that the smoothness properties of u are such that

$$|\hat{u}_k| \leq C|k|^{-p-1}.$$

So we have

$$|\hat{u}_0 - \hat{v}_0| \leq C \sum_{m \neq 0} (mN)^{-p-1} \leq C' N^{-p-1},$$

which is a $O(h^{p+1})$. □

As a conclusion, the trapezoidal rule is spectrally accurate (error $O(h^{p+1})$) for all $p \geq 0$ when $u \in C^\infty$, *provided* the function to be integrated is smooth *and* periodic. If the function is C^∞ in an interval but is for instance discontinuous upon periodization, then we revert to the usual $O(h^2)$ rate. So the true reason for the trapezoidal rule generally being $O(h^2)$ and not $O(h^\infty)$ is the presence of the boundaries.

An important example of such periodic smooth function, is a regular C^∞ function multiplied by a C^∞ window that goes smoothly to zero at the endpoints of $[a, b]$, like for instance a member of a partition of unity. (This plays an important role in some electromagnetism solvers, for computing the radar cross-section of scatterers.)

8.3.2 Chebyshev integration

Like for Chebyshev differentiation, we can warp θ into x by the formula $\theta = \arccos x$, and treat smooth functions that are not necessarily periodic. Integrating a Chebyshev interpolant gives rise to Chebyshev integration, also called Clenshaw-Curtis quadrature.

Assume that $f(x)$ is given for $x \in [-1, 1]$, otherwise rescale the x variable. The Chebyshev interpolant is built from the knowledge of f at the Chebyshev nodes $x_j = \cos \theta_j$, and takes the form

$$p(x) = \sum_{n \geq 0} a_n T_n(x).$$

We have seen that a_n are obtained by even extension of $f(x_j)$, followed by an FFT where the sine terms are dropped. As a result,

$$\int_{-1}^1 p(x) dx = \sum_{n \geq 0} a_n \int_{-1}^1 T_n(x) dx.$$

We compute

$$\int_{-1}^1 T_n(x) dx = \int_0^\pi \cos(n\theta) \sin \theta d\theta.$$

This integral can be evaluated easily by relating \cos and \sin to complex exponentials (or see Trefethen's book for an elegant shortcut via complex functions), and the result is

$$\int_{-1}^1 T_n(x) dx = \begin{cases} 0 & \text{if } n \text{ is odd;} \\ \frac{2}{1-n^2} & \text{if } n \text{ is even.} \end{cases}$$

The algorithm for Chebyshev differentiation is simply: (1) find the coefficients a_n for $0 \leq n \leq N$, and (2) form the weighted sum

$$\sum_{n \text{ even}, n \leq N} a_n \frac{2}{1-n^2}.$$

The accuracy estimate of Chebyshev integration is the same as that of bandlimited integration, except we do not need periodicity. So the method is spectrally accurate.

An important method related to Chebyshev quadrature is Gaussian quadrature, which is similar but somewhat different. It is also spectrally accurate for

smooth functions. Instead of using extrema or zeros of Chebyshev polynomials, it uses zeros of Legendre polynomials. It is usually derived directly from properties of orthogonal polynomials and their zeros, rather than Fourier analysis. This topic is fascinating but would take us a little too far.