Mathematical Foundations of Data Sciences

Gabriel Peyré
CNRS & DMA, École Normale Supérieure

CHAPTER 2

Fourier and Convolution

September 9, 2026

Convolution combines shifted copies of a signal, but evaluating it directly can be costly and obscure the structure of the operation. Fourier coordinates turn convolution into multiplication, providing both efficient algorithms and a way to solve differential equations. We develop this connection on continuous and finite domains, study sampling and the fast Fourier transform, and extend the viewpoint to groups, surfaces, and graphs.

The main reference for this chapter is [1].

2.1 Hilbert Spaces and Fourier Transforms

2.1.1 Orthonormal Bases

Many methods in data science begin by expanding the input signal in a basis. An orthonormal basis gives a simple reconstruction formula and preserves energy, which simplifies the analysis. We first describe these properties in a Hilbert space, such as $\mathcal{H} = L^2([0,1]^d)$ for signals on a continuous domain or $\mathcal{H} = \mathbb{R}^N$ for discrete signals.

A complex Hilbert space $(\mathcal{H}, \langle\cdot,\cdot\rangle)$ is complete for the norm induced by its Hermitian inner product. We take this inner product to be linear in the first argument and conjugate-linear in the second. A separable Hilbert space admits an orthonormal basis $(\varphi_k)_k$ (finite in finite dimension), so every $f \in \mathcal{H}$ has the expansion

$$f = \sum_k \langle f, \varphi_k\rangle \varphi_k$$

where the norm is determined by $\|f\|^2 \overset{\text{def.}}{=} \langle f, f\rangle$. Convergence means that $\|f - \sum_{k=0}^N \langle f, \varphi_k\rangle\varphi_k\| \to 0$ as $N \to +\infty$. Parseval's identity expresses conservation of energy:

$$\|f\|^2 = \sum_k |\langle f, \varphi_k\rangle|^2.$$

The Gram–Schmidt procedure constructs such a basis from a linearly independent family $(\bar\varphi_k)_k$ with dense span. Set $\varphi_0 = \bar\varphi_0/\|\bar\varphi_0\|$ and $\varphi_k = \tilde\varphi_k/\|\tilde\varphi_k\|$, where $\tilde\varphi_k$ is the orthogonal residual $\tilde\varphi_k = \bar\varphi_k - \sum_{i<k} a_i\varphi_i$. Orthogonality requires $a_i = \langle\bar\varphi_k, \varphi_i\rangle$.

On $L^2([-1,1])$ with the usual inner product, orthogonalizing the monomials gives the Legendre polynomials. Their conventional normalization differs from unit L^2 norm:

$$\varphi_0(x) = 1, \quad \varphi_1(x) = x, \quad \varphi_2(x) = \frac{1}{2}(3x^2 - 1), \quad \text{etc.}$$

On $L^2(\mathbb{R}, e^{-x^2/2}\,\mathrm{d}x)$, this construction gives the probabilists' Hermite polynomials $P_0(x) = 1$, $P_1(x) = x$, $P_2(x) = x^2 - 1$, up to normalization. Multiplying normalized polynomials by $e^{-x^2/4}$ instead gives orthonormal Hermite functions for Lebesgue measure. A degree-k orthogonal polynomial has exactly k distinct real zeros in the interior of the support of the measure.

The Shannon interpolation theorem gives another example: $(\mathrm{sinc}(x-k))_k$ is an orthonormal basis of the subspace $\left\{f \in L^2(\mathbb{R})\ ;\ \mathrm{supp}(\hat f) \subset [-\pi,\pi]\right\}$. Using the continuous representative of f, the reconstruction

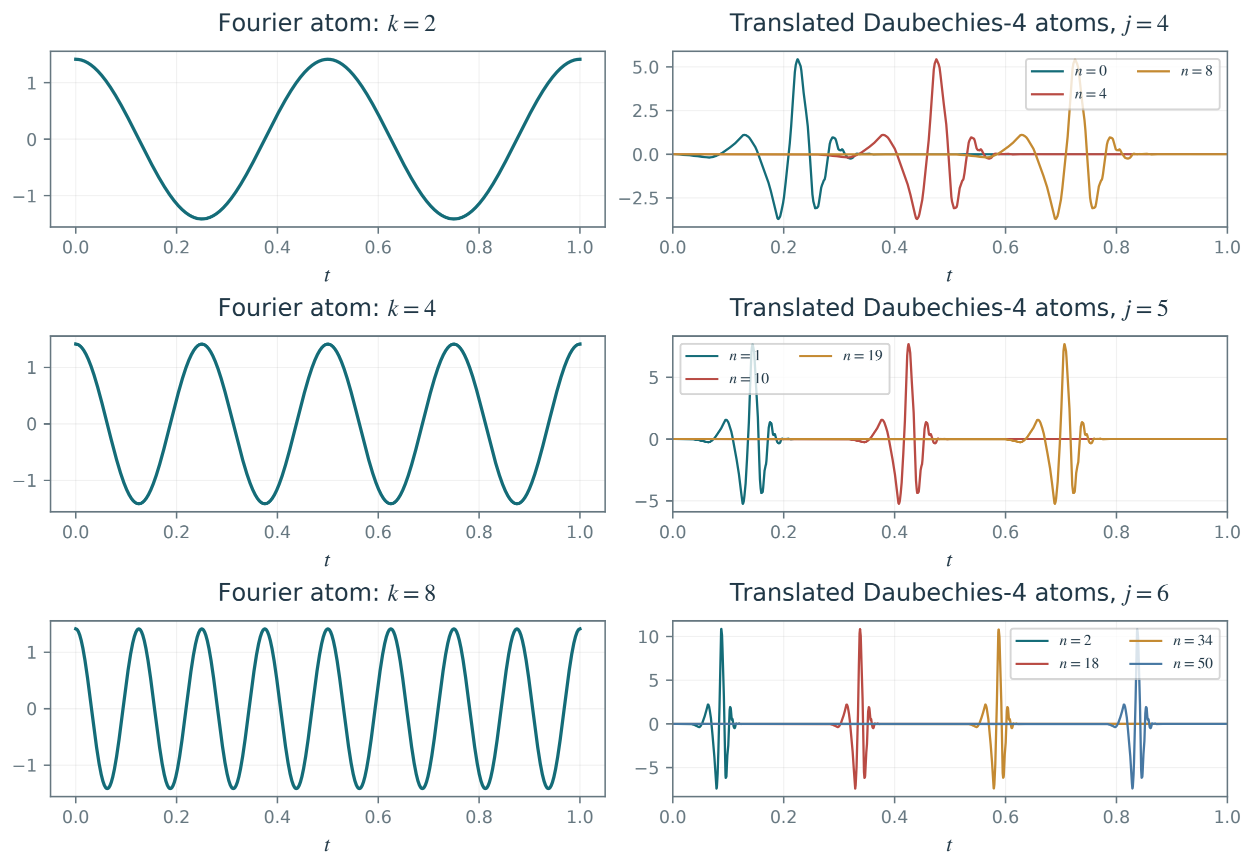

Figure 2.1. Fourier modes (left) and several translated Daubechies wavelets at each scale (right).

formula $f = \sum_k f(k)\,\mathrm{sinc}(x-k)$ identifies the coefficients as $\langle f, \mathrm{sinc}(x-k)\rangle = f(k)$. The series converges in L^2 and pointwise; the previous chapter proves the pointwise formula under additional decay assumptions.

Orthogonal bases can also arise as complete families of eigenvectors of self-adjoint operators. Below, we explain how translation-invariant operators (convolutions) characterize Fourier bases.

2.1.2 Fourier Basis on $\mathbb{R}/2\pi\mathbb{Z}$

Fourier analysis takes different forms on different domains. On the circle, it gives a countable orthonormal basis. On the real line, frequencies form a continuum, and the Fourier transform is not an expansion in a countable basis of exponential functions.

On $L^2(\mathbb{T})$ where $\mathbb{T} = \mathbb{R}/2\pi\mathbb{Z}$, equipped with $\langle f, g\rangle \stackrel{\text{def.}}{=} \frac{1}{2\pi}\int_{\mathbb{T}} f(x)\bar{g}(x)\mathrm{d}x$, one can use the Fourier basis

$$\varphi_k(x) \stackrel{\text{def.}}{=} e^{ikx} \quad \text{for} \quad k \in \mathbb{Z}. \tag{2.1}$$

One thus has

$$f = \sum_{k\in\mathbb{Z}} \hat{f}_k e^{ik\cdot} \quad \text{where} \quad \hat{f}_k \stackrel{\text{def.}}{=} \frac{1}{2\pi}\int_0^{2\pi} f(x)e^{-ikx}\mathrm{d}x, \tag{2.2}$$

with convergence in $L^2(\mathbb{T})$. Pointwise convergence is delicate; see Section 1.2.

Figure 2.1, left, shows examples of the real parts of Fourier atoms.

We recall that for $f \in L^1(\mathbb{R})$, its Fourier transform is defined as

$$\forall\, \omega \in \mathbb{R}, \quad \hat{f}(\omega) \stackrel{\text{def.}}{=} \int_{\mathbb{R}} f(x)e^{-ix\omega}\mathrm{d}x.$$

The transform extends to $L^2(\mathbb{R})$ by density.

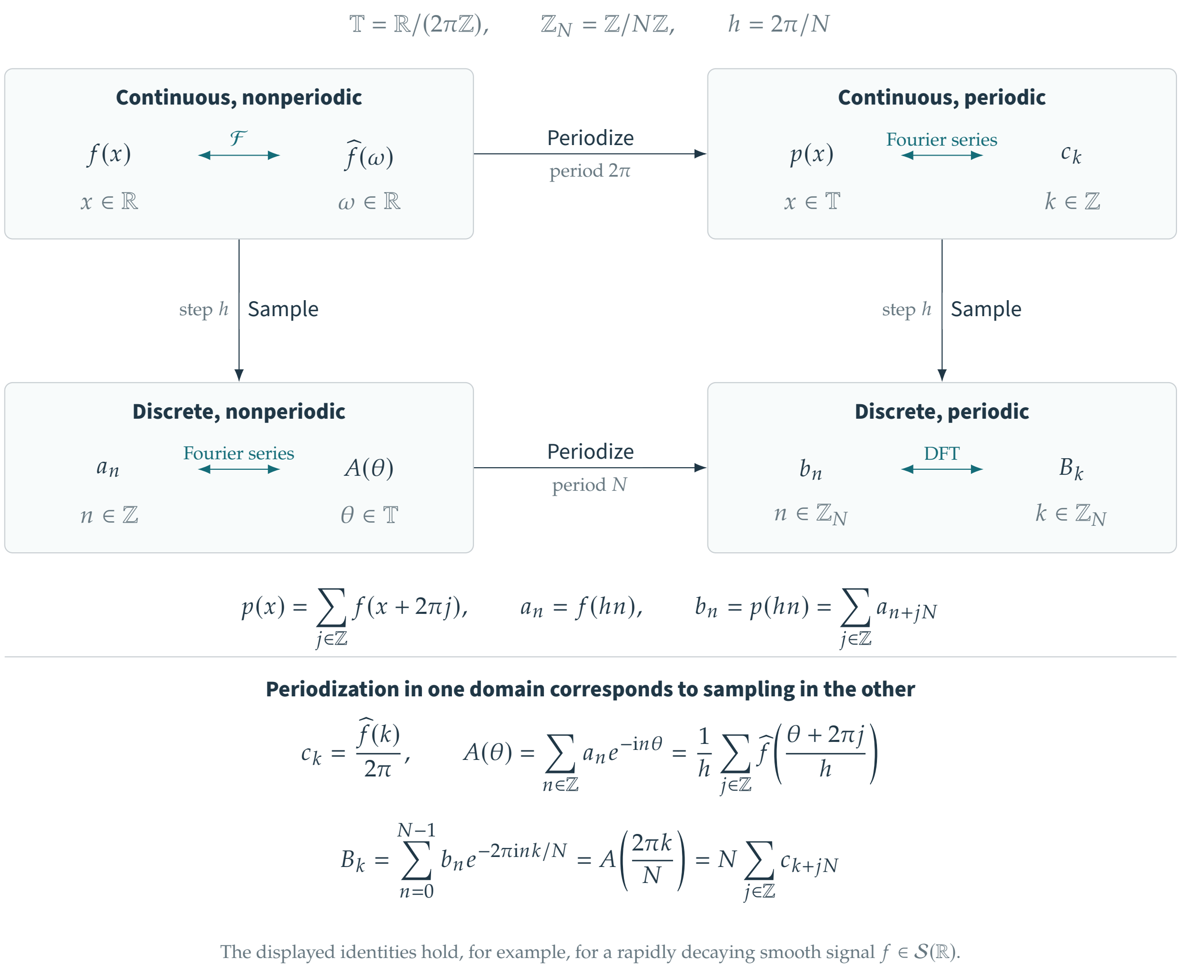

Figure 2.2. Four settings for Fourier analysis, linked by sampling and periodization.

The following diagram relates the Fourier transform on $\mathbb{R}$ to Fourier coefficients on $\mathbb{T}$:

$$\text{sampling}\quad \begin{array}{ccc} f(x) & \xrightarrow{\mathcal{F}} & \hat{f}(\omega) \\ \downarrow & & \downarrow \\ (f(n))_n & \xrightarrow{\text{Fourier series}} & \sum_n f(n)e^{-\mathrm{i}\omega n} \end{array} \quad \text{periodization} .$$

The two routes through the diagram give the identity

$$\sum_n f(n)e^{-\mathrm{i}\omega n} = \sum_n \hat{f}(\omega - 2\pi n) \tag{2.3}$$

which is Poisson's summation formula (Proposition 1.1).

2.2 Convolution on $\mathbb{R}$ and $\mathbb{T}$

2.2.1 Convolution

On $\mathbb{X} = \mathbb{R}$ or $\mathbb{T}$, one defines

$$f \star g(x) = \int_{\mathbb{X}} f(t)g(x-t)\mathrm{d}t. \tag{2.4}$$

Here and throughout this section, integrals use Lebesgue measure on $\mathbb{R}$ and normalized Haar measure $\mathrm{d}t/(2\pi)$ on $\mathbb{T}$. With this convention, Young's inequality states that, for $1 \leqslant p, q, r \leqslant \infty$ and $(f, g) \in L^p(\mathbb{X}) \times L^q(\mathbb{X})$,

$$\frac{1}{p} + \frac{1}{q} = 1 + \frac{1}{r} \quad \Longrightarrow \quad f \star g \in L^r(\mathbb{X}) \quad \text{and} \quad \|f \star g\|_{L^r(\mathbb{X})} \leqslant \|f\|_{L^p(\mathbb{X})}\|g\|_{L^q(\mathbb{X})}. \tag{2.5}$$

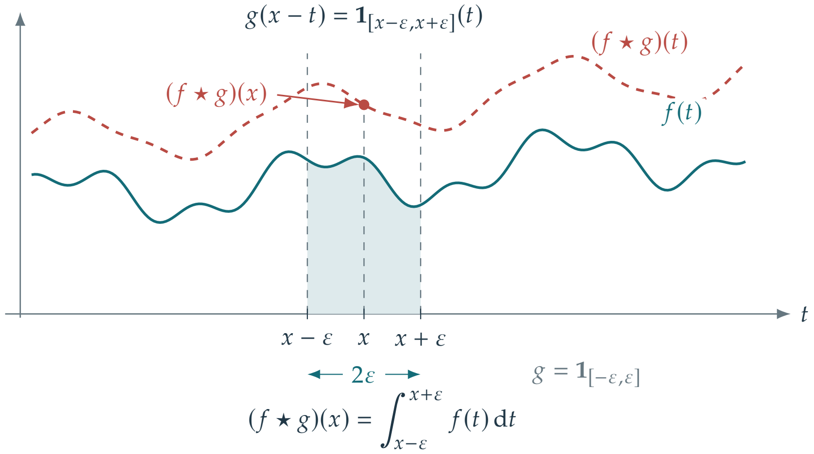

Figure 2.3. Convolution on $\mathbb{R}$. The dashed curve is $f \star g$; the marked value $(f \star g)(x)$ equals the shaded integral.

In particular, convolution by $f \in L^1(\mathbb{X})$ defines a bounded linear map from $L^p(\mathbb{X})$ to itself. When $r = \infty$ and $1 < p, q < \infty$, the convolution $f \star g$ has a bounded continuous representative, illustrating its regularizing effect. On the circle, $p < q \implies L^q(\mathbb{T}) \subset L^p(\mathbb{T})$, so $L^\infty(\mathbb{X})$ is contained in every space in this scale.

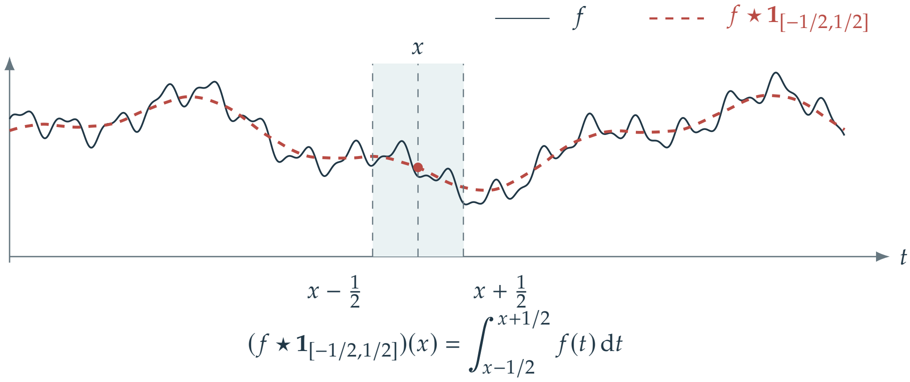

Figure 2.4. Signal filtering with a box filter (running average).

Convolution can regularize functions. For instance, if $f \in L^1(\mathbb{X})$ and $g \in C^1(\mathbb{X})$ is bounded with bounded derivative, then $f \star g$ is differentiable and $(f \star g)' = f \star g'$. To construct an approximate identity, choose a smooth compactly supported $\rho \geqslant 0$ with $\int_{\mathbb{R}} \rho = 1$ and set $\rho_\varepsilon(x) = \varepsilon^{-1}\rho(x/\varepsilon)$. Then $f \star \rho_\varepsilon \to f$ in $L^p(\mathbb{R})$ for $1 \leqslant p < \infty$, and uniformly for bounded uniformly continuous f. On the circle, use $2\pi \sum_{k\in\mathbb{Z}} \rho_\varepsilon(\cdot - 2\pi k)$, whose integral against normalized Haar measure is one. This smoothing effect is also useful for denoising signals and images.

Proposition 2.1. *If $f, g \in L^1(\mathbb{X})$, then*

$$\mathcal{F}(f \star g) = \hat{f} \odot \hat{g}. \tag{2.6}$$

For $\mathbb{X} = \mathbb{R}$, $\odot$ denotes pointwise multiplication of the continuous Fourier transforms. For $\mathbb{X} = \mathbb{T}$, $\odot$ acts on bounded Fourier-coefficient sequences (pointwise multiplication of sequences).

This means that $\mathcal{F}$ is an algebra homomorphism. For instance, if $\mathbb{X} = \mathbb{R}$, its range lies in the algebra of continuous functions that vanish at $\pm\infty$.

As shown in Figure 2.7, successive convolutions of a box function produce cardinal splines: piecewise polynomials of increasing smoothness.

Convolution also describes the sum of independent random variables. Writing f_X for the probability density of a random vector X with respect to Lebesgue measure, independence of X and Y gives $f_{X+Y} = f_X \star f_Y$. This identity is fundamental in probability and statistics.

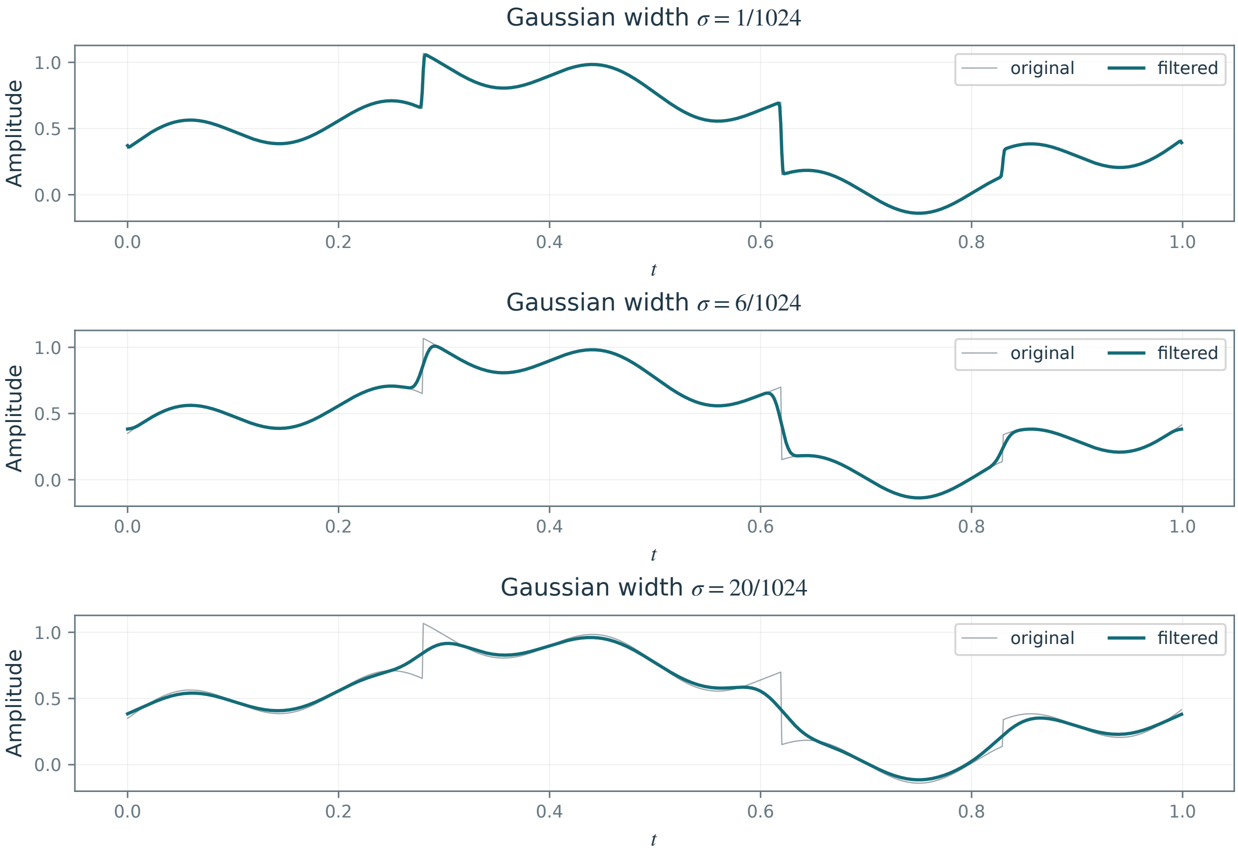

Figure 2.5. Filtering an irregular signal with Gaussian filters of increasing width σ.

$$\begin{tikzcd}
(f, g) \arrow[r, "\star"] \arrow[d, "\mathcal{F} \times \mathcal{F}"'] & f \star g \arrow[d, "\mathcal{F}"] \\
(\widehat{f}, \widehat{g}) \arrow[r, "\cdot"'] & \widehat{f}\,\widehat{g}
\end{tikzcd}$$

Figure 2.6. How the Fourier transform converts convolution into multiplication.

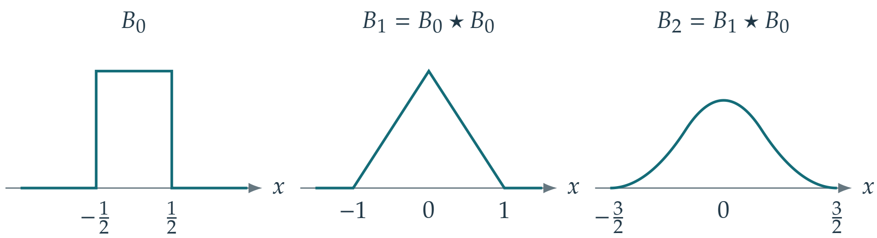

$$B_0 = \mathbf{1}_{[-1/2,1/2]}, \qquad B_m = B_0^{\star(m+1)}$$

Figure 2.7. Cardinal splines are defined by successive convolutions.

$$
\begin{tikzcd}
f \arrow[r, "H"] \arrow[d, "T_\tau"'] & Hf \arrow[d, "T_\tau"] \\
T_\tau f \arrow[r, "H"'] & T_\tau(Hf)
\end{tikzcd}
$$

$$T_\tau f(x) = f(x-\tau)$$

Figure 2.8. Commutative diagram for translation invariance.

```
Associated code:  fourier/test_denoising.m and fourier/test_fft2.m
```

2.2.2 Translation Invariant Operators

Translation-invariant operators commute with translations. In signal and image processing, they apply the same rule at every location.

The following propositions characterize translation-invariant operators as convolutions against kernels that may be distributions. The term "translation equivariant" is more precise: translating the input translates the output by the same amount. The kernel's regularity depends on the topologies of the input and output spaces, so the proof differs between the two settings below. We begin with operators whose outputs are continuous functions.

Proposition 2.2. *We define $T_\tau f = f(\cdot - \tau)$. A bounded linear operator $H : (L^2(\mathbb{X}), \|\cdot\|_2) \to (C_b(\mathbb{X}), \|\cdot\|_\infty)$ commutes with every translation, meaning that for all τ, $H \circ T_\tau = T_\tau \circ H$, if and only if*

$$\forall f \in L^2(\mathbb{X}), \quad H(f) = f \star g$$

with $g \in L^2(\mathbb{X})$.

Formally, the identity $f = f \star \delta$ suggests that $Hf = f \star H\delta$ for a translation-invariant operator. Since $H\delta$ need not be defined under the present hypotheses, the following argument uses bounded evaluation functionals instead.

Proof. By (2.5) and the continuity result for $r = \infty$, the convolution operator $H : f \mapsto f \star g$ with $g \in L^2(\mathbb{X})$ is bounded from $L^2(\mathbb{X})$ to $C_b(\mathbb{X})$. Moreover,

$$(H \circ T_\tau)(f) = \int_{\mathbb{X}} f((\cdot - \tau) - y) g(y) \mathrm{d}y = (f \star g)(\cdot - \tau) = T_\tau(Hf),$$

so that such an H is translation-invariant.

Conversely, we define $\ell : f \mapsto H(f)(0) \in \mathbb{C}$, which is well defined since $H(f) \in C^0(\mathbb{X})$. Continuity of H gives a constant C such that $\|Hf\|_\infty \leqslant C\|f\|_2$, hence $|\ell(f)| \leqslant C\|f\|_2$. Thus ℓ is a continuous linear functional on the Hilbert space $L^2(\mathbb{X})$. By the Fréchet–Riesz theorem, there exists $h \in L^2(\mathbb{X})$ such that $\ell(f) = \langle f,\, h \rangle$. Hence, $\forall\, x \in \mathbb{X}$,

$$H(f)(x) = T_{-x}(Hf)(0) = H(T_{-x} f)(0) = \ell(T_{-x} f) = \langle T_{-x} f,\, h \rangle = \int_{\mathbb{X}} f(y + x) \overline{h(y)}\, \mathrm{d}y = (f \star g)(x),$$

where $g(x) \overset{\text{def.}}{=} \overline{h(-x)} \in L^2(\mathbb{X})$. □

On $\mathbb{T}$, an operator need not produce continuous outputs. Its convolution kernel may then be a distribution, and Fourier coefficients provide a convenient way to define the operator.

Proposition 2.3. *A bounded linear operator $H : L^2(\mathbb{T}) \to L^2(\mathbb{T})$ commutes with every translation, meaning that for all τ, $H \circ T_\tau = T_\tau \circ H$, if and only if*

$$\forall f \in L^2(\mathbb{T}), \quad \mathcal{F}(H(f)) = \hat{f} \odot c$$

where $c \in \ell^\infty(\mathbb{Z})$ (a bounded sequence).

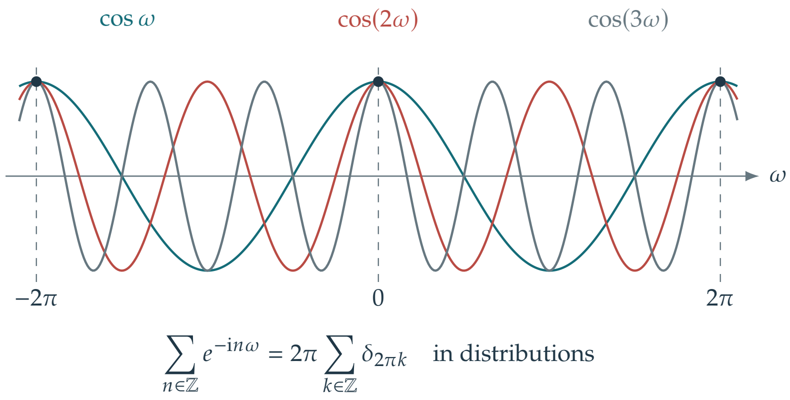

Figure 2.9. Sine waves summed in the Poisson formula.

Proof. We denote $\varphi_n \stackrel{\text{def.}}{=} e^{\mathrm{i}n\cdot}$. One has

$$H(\varphi_n) = H(T_\tau T_{-\tau}\varphi_n) = H(T_\tau e^{\mathrm{i}n\tau}\varphi_n) = e^{\mathrm{i}n\tau}H(T_\tau(\varphi_n)) = e^{\mathrm{i}n\tau}T_\tau(H(\varphi_n)).$$

Thus, for all n,

$$\langle H(\varphi_n),\, \varphi_m\rangle = e^{\mathrm{i}n\tau}\langle T_\tau \circ H(\varphi_n),\, \varphi_m\rangle = e^{\mathrm{i}n\tau}\langle H(\varphi_n),\, T_{-\tau}(\varphi_m)\rangle = e^{\mathrm{i}(n-m)\tau}\langle H(\varphi_n),\, \varphi_m\rangle$$

So for $n \neq m$, $\langle H(\varphi_n),\, \varphi_m\rangle = 0$, and we define $c_n \stackrel{\text{def.}}{=} \langle H(\varphi_n),\, \varphi_n\rangle$. We thus have $H(\varphi_n) = c_n\varphi_n$. Since H is continuous, $\|Hf\|_{L^2(\mathbb{T})} \leqslant C\|f\|_{L^2(\mathbb{T})}$ for some constant C, and thus by Cauchy–Schwarz

$$|c_n| = |\langle H(\varphi_n),\, \varphi_n\rangle| \leqslant \|H(\varphi_n)\|\|\varphi_n\| \leqslant C$$

because $\|\varphi_n\| = 1$, so that $c \in \ell^\infty(\mathbb{Z})$. By continuity of H, recalling that by definition $\hat{f}_n \stackrel{\text{def.}}{=} \langle f,\, \varphi_n\rangle$,

$$H(f) = \lim_N H\Big(\sum_{n=-N}^{N} \hat{f}_n\varphi_n\Big) = \lim_N \sum_{n=-N}^{N} \hat{f}_n H(\varphi_n) = \lim_N \sum_{n=-N}^{N} c_n\hat{f}_n\varphi_n = \sum_{n\in\mathbb{Z}} c_n\hat{f}_n\varphi_n$$

where we used $H(\varphi_n) = c_n\varphi_n$. This proves the desired representation. □

Conversely, by Parseval's identity, a bounded sequence c defines a bounded Fourier multiplier, which commutes with translations. Translation-invariant operators are therefore diagonal in the Fourier basis.

2.2.3 Poisson's Formula and Distributions

Informally, the Fourier series

$$\sum_n f(n)e^{-\mathrm{i}\omega n}$$

can be viewed as the Fourier transform $\mathcal{F}(\Pi_1 \odot f)$ of the discrete distribution

$$\Pi_1 \odot f = \sum_n f(n)\delta_n \quad \text{where} \quad \Pi_s = \sum_n \delta_{sn}$$

for $s > 0$, where δ_a is the Dirac mass at location $a \in \mathbb{R}$, i.e. the distribution such that $\int f\,\mathrm{d}\delta_a = f(a)$ for any continuous f. A measure can be multiplied by a continuous function; a general distribution can be multiplied by a smooth function. For a tempered distribution μ, its Fourier transform is defined by duality on Schwartz test functions:

$$\forall\, g \in \mathcal{S}(\mathbb{R}), \quad \int_{\mathbb{R}} g(x)\mathrm{d}\mathcal{F}(\mu) = \int_{\mathbb{R}} \mathcal{F}(g)\mathrm{d}\mu, \quad \text{where} \quad \mathcal{F}(g) \stackrel{\text{def.}}{=} \int_{\mathbb{R}} g(x)e^{-\mathrm{i}x\cdot}\,\mathrm{d}x,$$

Here $\mathcal{S}(\mathbb{R})$ is the Schwartz space; the integral notation denotes the distributional pairing.

The Poisson formula (2.3) can thus be interpreted as

$$\mathcal{F}(\Pi_1 \odot f) = \sum_n \hat{f}(\cdot - 2\pi n) = \int_{\mathbb{R}} \hat{f}(\cdot - \omega) \mathrm{d}\Pi_{2\pi}(\omega) = \hat{f} \star \Pi_{2\pi}$$

Since $\mathcal{F}^{-1} = \frac{1}{2\pi}\mathcal{S} \circ \mathcal{F}$ where $\mathcal{S}(f) = f(-\cdot)$, applying this operator to both sides gives

$$\Pi_1 \odot f = \frac{1}{2\pi}\mathcal{S} \circ \mathcal{F}(\hat{f} \star \Pi_{2\pi}) = \mathcal{S}(\frac{1}{2\pi}\mathcal{F}(\hat{f}) \odot \hat{\Pi}_{2\pi}) = \mathcal{S}(\frac{1}{2\pi}\mathcal{F}(\hat{f})) \odot \mathcal{S}(\hat{\Pi}_{2\pi}) = \hat{\Pi}_{2\pi} \odot f.$$

This can be interpreted as the relation

$$\hat{\Pi}_{2\pi} = \Pi_1 \quad \Longrightarrow \quad \hat{\Pi}_1 = 2\pi\Pi_{2\pi}.$$

For intuition, consider a finite Fourier series

$$\sum_{n=-N}^{N} e^{-\mathrm{i}n\omega} = \frac{\sin((N+1/2)\omega)}{\sin(\omega/2)}$$

which is the Dirichlet kernel, with removable singularities at multiples of 2π. Its value there is $2N + 1$; in the sense of distributions it converges to $2\pi\Pi_{2\pi}$.

2.3 Finite Fourier Transform and Convolution

2.3.1 Discrete Orthonormal Bases

A discrete signal is a finite-dimensional vector $f \in \mathbb{C}^N$, where N is the number of samples and f_n is the value at a specified location. For a 2-D image $f \in \mathbb{C}^N \simeq \mathbb{C}^{N_0 \times N_0}$, we have $N = N_0 \times N_0$, with N_0 pixels in each direction.

The discrete inner product is the analogue of the continuous L^2 inner product:

$$\langle f, g \rangle = \sum_{n=0}^{N-1} f_n \bar{g}_n.$$

The corresponding squared Euclidean distance is

$$\|f - g\|^2 = \sum_{n=0}^{N-1} |f_n - g_n|^2.$$

Exactly as in the continuous case, a discrete orthonormal basis $\{\psi_k\}_{0 \leqslant k < N}$ of $\mathbb{C}^N$ satisfies

$$\langle \psi_k, \psi_{k'} \rangle = \delta_{k-k'}. \tag{2.7}$$

A signal expands in this orthonormal basis as

$$f = \sum_{k=0}^{N-1} \langle f, \psi_k \rangle \psi_k.$$

It preserves energy

$$\|f\|^2 = \sum_{n=0}^{N-1} |f_n|^2 = \sum_{k=0}^{N-1} |\langle f, \psi_k \rangle|^2$$

Computing the set of all inner products $\{\langle f, \psi_k \rangle\}_{0 \leqslant k < N}$ directly requires $O(N^2)$ operations. This cost is prohibitive when N reaches millions of samples. Structured bases permit faster decompositions: Fourier and wavelet transforms require $O(N \log(N))$ and $O(N)$ operations, respectively.

2.3.2 Discrete Fourier transform

We denote $f = (f_n)_{n=0}^{N-1} \in \mathbb{R}^N$, and regard its entries as indexed by $n \in \mathbb{Z}/N\mathbb{Z}$, which is a finite abelian group under addition. This corresponds to using periodic boundary conditions.

The discrete Fourier transform is defined as

$$\forall\, k = 0, \dots, N-1, \quad \hat{f}_k \overset{\text{def.}}{=} \sum_{n=0}^{N-1} f_n e^{-\frac{2i\pi}{N} kn} = \langle f, \varphi_k \rangle \quad \text{where} \quad \varphi_k \overset{\text{def.}}{=} (e^{\frac{2i\pi}{N} kn})_{n=0}^{N-1} \in \mathbb{C}^N \tag{2.8}$$

where the canonical inner product on $\mathbb{C}^N$ is $\langle u, v \rangle = \sum_{n=0}^{N-1} u_n \bar{v}_n$ for $(u, v) \in (\mathbb{C}^N)^2$. This definition is motivated by sampling the Fourier basis $x \mapsto e^{ikx}$ on $\mathbb{R}/2\pi\mathbb{Z}$ at equally spaced points $(\frac{2\pi}{N} n)_{n=0}^{N-1}$. The next proposition establishes the orthogonality of the discrete Fourier atoms and the associated reconstruction formula.

Proposition 2.4. *One has*

$$\langle \varphi_k, \varphi_\ell \rangle = \begin{cases} N & \text{if} \quad k = \ell, \\ 0 & \text{otherwise}. \end{cases}$$

In particular, this implies

$$\forall\, n = 0, \dots, N-1, \quad f_n = \frac{1}{N} \sum_k \hat{f}_k e^{\frac{2i\pi}{N} kn}. \tag{2.9}$$

Proof. One has, if $k \neq \ell$

$$\langle \varphi_k, \varphi_\ell \rangle = \sum_n e^{\frac{2i\pi}{N}(k-\ell)n} = \frac{1 - e^{\frac{2i\pi}{N}(k-\ell)N}}{1 - e^{\frac{2i\pi}{N}(k-\ell)}} = 0$$

which is the sum of a geometric series (equivalently, the sum of equispaced points on a circle). The inversion formula is simply $f = \sum_k \langle f, \varphi_k \rangle \frac{\varphi_k}{\|\varphi_k\|^2}$. ◻

2.3.3 Fast Fourier transform

Assuming $N = 2N'$, one has

$$\hat{f}_{2k} = \sum_{n=0}^{N'-1} (f_n + f_{n+N/2}) e^{-\frac{2i\pi}{N'} kn}$$
$$\hat{f}_{2k+1} = \sum_{n=0}^{N'-1} e^{-\frac{2i\pi}{N} n} (f_n - f_{n+N/2}) e^{-\frac{2i\pi}{N'} kn}.$$

For the second line, we used the computation

$$e^{-\frac{2i\pi}{N}(2k+1)(n+N/2)} = e^{-\frac{2i\pi}{N}(2kn+kN+n+N/2)} = -e^{-\frac{2i\pi}{N} n} e^{-\frac{2i\pi}{N'} kn}.$$

Let $\mathcal{F}_N(f) = \hat{f}$ denote the discrete Fourier transform on $\mathbb{C}^N$. Introduce the vectors $f_e = (f_n + f_{n+N/2})_n \in \mathbb{C}^{N'}$ and $f_o = (f_n - f_{n+N/2})_n \in \mathbb{C}^{N'}$ and the twiddle factors $\alpha_N = (e^{-2i\pi n/N})_{n=0}^{N'-1} \in \mathbb{C}^{N'}$. The recursion becomes

$$\mathcal{F}_N(f) = \mathcal{I}_N(\mathcal{F}_{N/2}(f_e), \mathcal{F}_{N/2}(f_o \odot \alpha_N))$$

where $\mathcal{I}_N$ is the "interleaving" operator, defined by $\mathcal{I}_N(a, b) \overset{\text{def.}}{=} (a_0, b_0, a_1, b_1, \dots, a_{N'-1}, b_{N'-1})$. Recursing yields the radix-two fast Fourier transform (FFT) when N is a power of two. Other FFT factorizations handle general lengths. Zero-padding to a power of two is also useful, although it changes the transform length and the sampled frequencies.

This algorithm can also be interpreted as a factorization of the Fourier matrix into a product of $O(\log N)$ sparse matrices.

Let $C(N)$ denote the number of elementary operations required to compute $\hat{f}$. The recursion gives

$$C(N) = 2C(N/2) + NK \tag{2.10}$$

where KN accounts for N complex additions and $N/2$ multiplications. With the change of variable

$$\ell \overset{\text{def.}}{=} \log_2(N) \quad \text{and} \quad T(\ell) \overset{\text{def.}}{=} \frac{C(N)}{N}$$

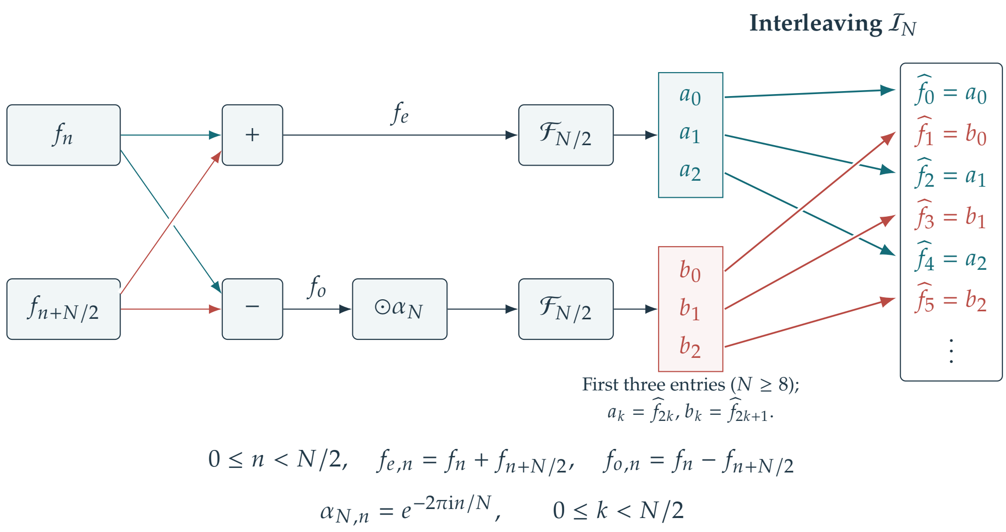

Figure 2.10. Diagram of one radix-two FFT step. Both halves of the input enter the sum and difference; the outputs of the two smaller transforms are interleaved into even and odd frequencies.

i.e. $C(N) = 2^\ell T(\ell)$, the relation (2.10) becomes

$$2^\ell T(\ell) = 2 \times 2^{\ell-1} T(\ell-1) + 2^\ell K \quad \Longrightarrow \quad T(\ell) = T(\ell-1) + K \quad \Longrightarrow \quad T(\ell) = T(0) + K\ell$$

and using the fact that $T(0) = C(1)/1 = 0$, one obtains

$$C(N) = KN \log_2(N).$$

Direct evaluation instead requires $O(N^2)$ operations to compute the N coefficients (2.8), each involving a sum of size N.

2.3.4 Finite convolution

For $(f, g) \in (\mathbb{R}^N)^2$, one defines $f \star g \in \mathbb{R}^N$ as

$$\forall n = 0, \ldots, N-1, \quad (f \star g)_n \stackrel{\text{def.}}{=} \sum_{k=0}^{N-1} f_k g_{n-k} = \sum_{k+\ell=n} f_k g_\ell \tag{2.11}$$

where $+$ and $-$ are interpreted modulo N (vectors are defined on the group $\mathbb{Z}/N\mathbb{Z}$, or equivalently, one uses periodic boundary conditions). This defines a commutative algebra structure $(\mathbb{R}^N, +, \star)$, with neutral element the "Dirac" $\delta_0 \stackrel{\text{def.}}{=} (1, 0, \ldots, 0)^\top \in \mathbb{R}^N$. The following proposition shows that $\mathcal{F} : f \mapsto \hat{f}$ is an algebra isomorphism and an isometry (up to a scaling by $\sqrt{N}$ of the norm) from $(\mathbb{C}^N, +, \star)$ onto $(\mathbb{C}^N, +, \odot)$ with neutral element $\mathbb{1}_N = (1, \ldots, 1) \in \mathbb{R}^N$.

Proposition 2.5. *One has* $\mathcal{F}(f \star g) = \hat{f} \odot \hat{g}$.

Proof. We denote $T : g \mapsto f \star g$. One has

$$(T\varphi_\ell)_n = (f \star \varphi_\ell)_n = \sum_p f_p e^{\frac{2i\pi}{N}\ell(n-p)} = e^{\frac{2i\pi}{N}\ell n} \hat{f}_\ell = \hat{f}_\ell (\varphi_\ell)_n.$$

Thus $(\varphi_\ell)_\ell$ form a basis of N eigenvectors of T, with eigenvalues $\hat{f}_\ell$. The operator T is diagonal in this basis. Let $F = (e^{-\frac{2i\pi}{N}kn})_{k,n}$ be the Fourier transform matrix. The inversion formula (2.9) gives $F^{-1} = \frac{1}{N}F^*$, where $F^* = \bar{F}^\top$ is the adjoint (conjugate transpose). The diagonalization of T is therefore

$$T = F^{-1} \operatorname{diag}(\hat{f}) F \quad \Longrightarrow \quad \mathcal{F}(Tg) = \operatorname{diag}(\hat{f}) F g \quad \Longrightarrow \quad \mathcal{F}(f \star g) = \operatorname{diag}(\hat{f}) \hat{g}.$$

□

This proposition gives an $O(N\log N)$ convolution algorithm:

$$f \star g = \mathcal{F}^{-1}(\hat{f} \odot \hat{g}).$$

This is more efficient than directly evaluating formula (2.11) for filters with large support. When $|\operatorname{Supp}(g)| = P$ is small, direct evaluation costs $O(PN)$ and may be faster. An example is $g = [1,1,0,\ldots,0,1]/3$, the moving average, where

$$(f \star g)_n = \frac{f_{n-1} + f_n + f_{n+1}}{3}$$

needs $3N$ operations.

Convolution as translation invariant operator. Define translation on $\mathbb{Z}/N\mathbb{Z}$ by $(T_\tau f)_n \stackrel{\text{def.}}{=} f_{n-\tau}$, with $n-\tau$ computed modulo N. The next proposition is the finite counterpart of the results on $\mathbb{R}$ and $\mathbb{R}/2\pi\mathbb{Z}$. On $\mathbb{Z}/N\mathbb{Z}$, no convergence issues arise, and the impulse response can be defined directly.

Proposition 2.6. *Let $H : \mathbb{R}^N \longrightarrow \mathbb{R}^N$ be linear. Then $H \circ T_\tau = T_\tau \circ H$ for all τ if and only if there exists $h \in \mathbb{R}^N$ (often called the impulse response) so that $Hf = f \star h$.*

Proof. For the convolution-to-invariance implication, we first note that $T_\tau f = f \star \delta_\tau$ is a convolution against a translated Dirac $\delta_\tau = T_\tau(\delta)$ where $\delta = (1,0,\ldots,0)$. If $Hf = f \star h$, commutativity gives

$$H \circ T_\tau(f) = h \star \delta_\tau \star f = \delta_\tau \star h \star f = T_\tau \circ H(f).$$

Conversely, we define $h \stackrel{\text{def.}}{=} H(\delta)$. Then

$$H(f) = H(\sum_i f_i(T_i\delta)) = \sum_i f_i H(T_i\delta) = \sum_i f_i T_i H(\delta) = \sum_i f_i h_{\cdot-i} = f \star h.$$

□

Polynomial multiplication. The FFT can multiply large polynomials. It can also multiply large integers by interpreting their digit expansions in a fixed base as polynomials. Indeed,

$$\Big(\sum_{i=0}^{A} a_i X^i\Big)\Big(\sum_{j=0}^{B} b_j X^j\Big) = \sum_{k=0}^{A+B}\Big(\sum_{i+j=k} a_i b_j\Big)X^k$$

Setting aside the carry operation, which can be performed in linear time, one can write $\sum_{i+j=k} a_i b_j = (\bar{a} \star \bar{b})_k$ when one defines $\bar{a}, \bar{b} \in \mathbb{R}^{A+B+1}$ by zero padding.

2.4 Discretization Issues

The FFT also approximates the continuous Fourier transform and its inverse. Reversing the roles of space and frequency leads to an efficient spectral interpolation method.

2.4.1 Fourier approximation via spatial zero padding.

The discrete Fourier transform (2.8) approximates samples of the continuous Fourier transform (1.2). For a sufficiently smooth f supported on $[0,1]$, consider the vector $f^Q \stackrel{\text{def.}}{=} (f(n/N))_{n=0}^{Q-1} \in \mathbb{R}^Q$. Taking $Q \geqslant N$ pads the original samples with zeros, since continuity and the support condition give $f(n/N) = 0$ for $n \geqslant N$. For signed frequency indices $-\lfloor Q/2\rfloor \leqslant k \leqslant \lceil Q/2\rceil - 1$ and $T = Q/N$,

$$\begin{aligned} \frac{1}{N}\hat{f}^Q_k &= \frac{1}{N}\sum_{n=0}^{N-1} f(n/N)e^{-2\pi i n k/Q} \\ &\approx \int_0^1 f(x)e^{-2\pi i k x/T}\,\mathrm{d}x = \hat{f}(2\pi k/T). \end{aligned}$$

For a fixed physical frequency, the Riemann-sum error is $O(1/N)$ for $f \in C^1$. The bound need not be uniform over frequencies growing with N. Increasing Q at fixed N samples the same spectrum more densely; it does not increase the Nyquist frequency.

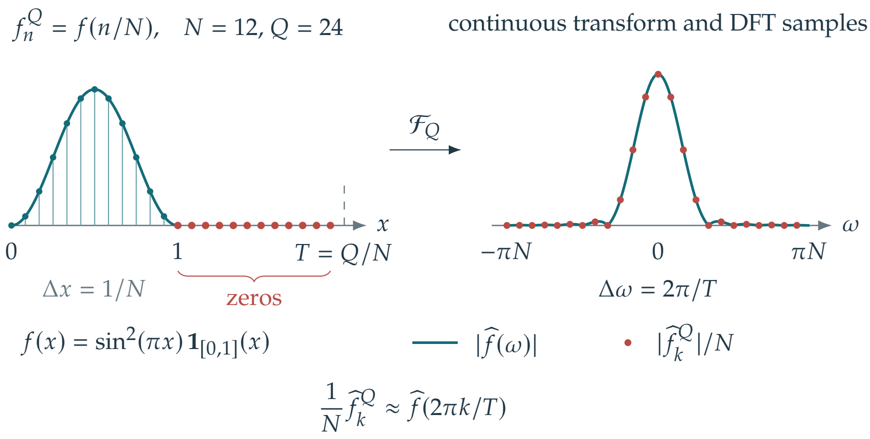

Figure 2.11. Fourier transform approximation by zero-padding in the spatial domain.

2.4.2 Fourier interpolation via spectral zero padding.

One can reverse the roles of space and frequency in the previous construction. Given N uniform discrete samples $f^N = (f_n^N)_{n=0}^{N-1}$, one can compute their discrete Fourier transform $\mathcal{F}(f^N) = \hat{f}^N$ (in $O(N\log(N))$ operations with the FFT),

$$\hat{f}_k^N \overset{\text{def.}}{=} \sum_{n=0}^{N-1} f_n^N e^{-\frac{2i\pi}{N}nk},$$

and then zero-pad the spectrum to obtain a vector of length $Q \geqslant N$. For simplicity, assume that $N = 2N'+1$ is odd. Even lengths admit a similar, slightly more involved construction. With frequency indices $-N' \leqslant k \leqslant N'$, define the padded spectrum by

$$\tilde{f}_k^Q = \hat{f}_k^N \quad (-N' \leqslant k \leqslant N'), \qquad \tilde{f}_k^Q = 0 \quad \text{otherwise}, \qquad \tilde{f}^Q \in \mathbb{C}^Q$$

The signed frequency indices are interpreted modulo Q when calling an FFT routine. An inverse transform of size Q, multiplied by Q/N, then gives in $O(Q\log Q)$ operations

$$\begin{aligned}
\frac{Q}{N}\mathcal{F}^{-1}(\tilde{f}^Q)_\ell &= \frac{Q}{N} \times \frac{1}{Q}\sum_{k=-N'}^{N'} \hat{f}_k^N e^{\frac{2i\pi}{Q}\ell k} = \frac{1}{N}\sum_{k=-N'}^{N'}\sum_{n=0}^{N-1} f_n^N e^{-\frac{2i\pi}{N}nk} e^{\frac{2i\pi}{Q}\ell k} \\
&= \sum_{n=0}^{N-1} f_n^N \frac{1}{N}\sum_{k=-N'}^{N'} e^{2i\pi\left(-\frac{n}{N}+\frac{\ell}{Q}\right)k} = \sum_{n=0}^{N-1} f_n^N \frac{\sin\left[\pi N\left(\frac{\ell}{Q}-\frac{n}{N}\right)\right]}{N\sin\left[\pi\left(\frac{\ell}{Q}-\frac{n}{N}\right)\right]} \\
&= \sum_{n=0}^{N-1} f_n^N \operatorname{sinc}_N\left(\frac{\ell}{T}-n\right) \quad \text{where} \quad T \overset{\text{def.}}{=} \frac{Q}{N} \quad \text{and} \quad \operatorname{sinc}_N(u) \overset{\text{def.}}{=} \frac{\sin(\pi u)}{N\sin(\pi u/N)}.
\end{aligned}$$

We use the geometric-series identity with $\rho = e^{i\omega}$, $a = -b$, $\omega = 2\pi\left(-\frac{n}{N}+\frac{\ell}{Q}\right)$,

$$\sum_{i=a}^{b}\rho^i = \frac{\rho^{a-\frac{1}{2}} - \rho^{b+\frac{1}{2}}}{\rho^{-\frac{1}{2}} - \rho^{\frac{1}{2}}} = \frac{\sin((b+\frac{1}{2})\omega)}{\sin(\omega/2)}.$$

This gives a discrete counterpart of the Shannon interpolation formula (1.9). It evaluates the trigonometric interpolant of the samples exactly on a grid of size Q, at cost $O(Q\log(Q))$. Furthermore, $\operatorname{sinc}_N(u) \to \operatorname{sinc}(u)$ for fixed u as $N \to \infty$, with removable singularities filled by continuity.

2.5 Fourier in Multiple Dimensions

Tensor products extend the one-dimensional Fourier transform to any finite dimension $d > 1$.

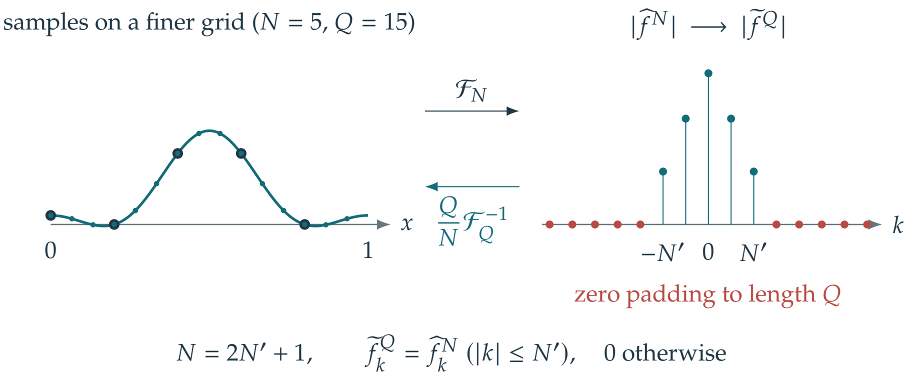

Figure 2.12. Interpolation by zero-padding in the frequency domain.

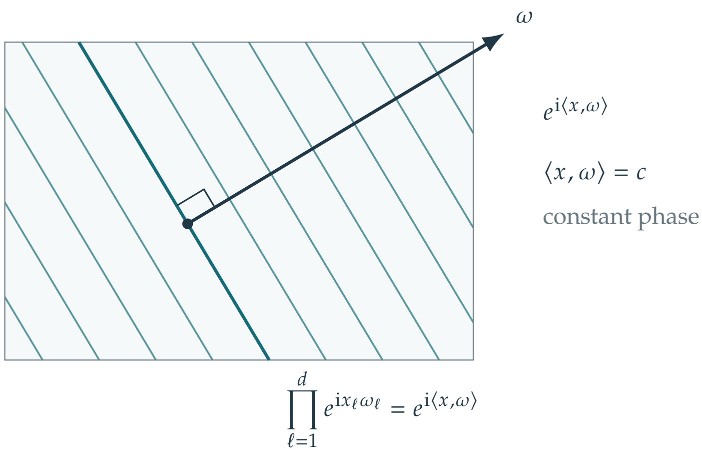

Figure 2.13. 2-D sine wave.

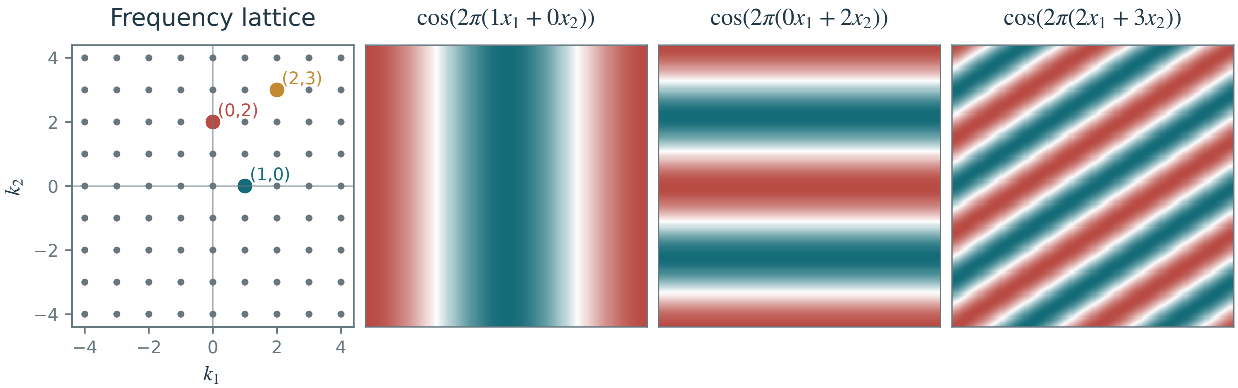

Figure 2.14. 2D Fourier orthogonal bases.

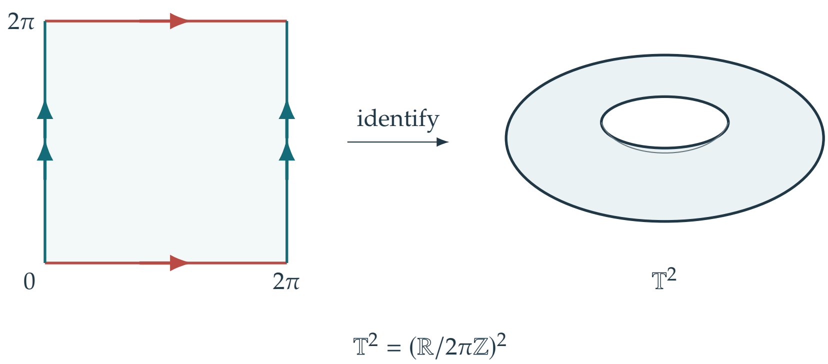

Figure 2.15. The two-dimensional torus $\mathbb{T}^2 = (\mathbb{R}/2\pi\mathbb{Z})^2$.

2.5.1 On Continuous Domains

On $\mathbb{R}^d$. Tensor products retain a useful feature of Fourier atoms: a product of elementary 1-D complex exponentials is a plane wave

$$\prod_{\ell=1}^{d} e^{\mathrm{i}x_\ell\omega_\ell} = e^{\mathrm{i}\langle x,\,\omega\rangle}$$

whose constant-phase hyperplanes are orthogonal to the wave vector $\omega = (\omega_\ell)_{\ell=1}^d \in \mathbb{R}^d$ when this vector is nonzero. Here $\langle x,\, \omega\rangle = \sum_\ell x_\ell\omega_\ell$ is the canonical inner product on $\mathbb{R}^d$.

The Fourier transform and its inverse are defined by

$$\forall\,\omega \in \mathbb{R}^d, \quad \hat{f}(\omega) \overset{\text{def.}}{=} \int_{\mathbb{R}^d} f(x)e^{-\mathrm{i}\langle x,\,\omega\rangle}\mathrm{d}x,$$
$$\forall\,x \in \mathbb{R}^d, \quad f(x) = \frac{1}{(2\pi)^d}\int_{\mathbb{R}^d} \hat{f}(\omega)e^{\mathrm{i}\langle x,\,\omega\rangle}\,\mathrm{d}\omega,$$

under the same integrability assumptions as in one dimension.

On $(\mathbb{R}/2\pi\mathbb{Z})^d$. Given an orthonormal basis $(\varphi_{n_1})_{n_1\in\mathbb{N}}$ of $L^2(\mathbb{X})$, one constructs an orthonormal basis of $L^2(\mathbb{X}^d)$ by tensorization

$$\forall\,k = (k_1,\ldots,k_d) \in \mathbb{N}^d, \quad \forall\,x \in \mathbb{X}^d, \quad \varphi_k(x) = \varphi_{k_1}(x_1)\ldots\varphi_{k_d}(x_d). \tag{2.12}$$

Fubini's theorem gives orthogonality. Completeness gives convergence of the rectangular partial sums $\sum_{\|k\|_\infty\leqslant N}\langle f,\, \varphi_k\rangle\varphi_k \to f$ in $L^2(\mathbb{X}^d)$.

For the multidimensional torus $(\mathbb{R}/2\pi\mathbb{Z})^d$, using the Fourier basis (2.1), this gives the basis

$$\forall\,k \in \mathbb{Z}^d, \quad \varphi_k(x) = e^{\mathrm{i}\langle x,\,k\rangle}$$

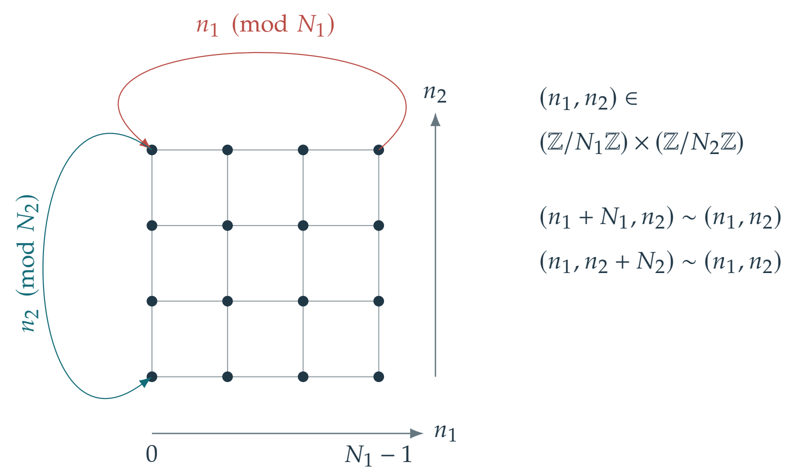

Figure 2.16. Discrete 2-D torus.

which is indeed a Hilbertian orthonormal basis for the inner product $\langle f,\, g\rangle \stackrel{\text{def.}}{=} \frac{1}{(2\pi)^d}\int_{\mathbb{T}^d} f(x)\bar{g}(x)\mathrm{d}x$. This defines the Fourier transform and the reconstruction formula on $L^2(\mathbb{T}^d)$

$$\hat{f}_k \stackrel{\text{def.}}{=} \frac{1}{(2\pi)^d}\int_{\mathbb{T}^d} f(x)e^{-\mathrm{i}\langle x,\,k\rangle}\mathrm{d}x \quad \text{and} \quad f = \sum_{k\in\mathbb{Z}^d} \hat{f}_k e^{\mathrm{i}\langle x,\,k\rangle}.$$

2.5.2 On Discrete Domains

Discrete Fourier Transform. On a d-dimensional discrete domain of the form

$$n = (n_1, \ldots, n_d) \in \mathbb{Y}_d \stackrel{\text{def.}}{=} [\![1, N_1]\!] \times \ldots \times [\![1, N_d]\!]$$

(we denote $[\![a, b]\!] \stackrel{\text{def.}}{=} \{i \in \mathbb{Z}\,;\; a \leqslant i \leqslant b\}$) of $N = N_1 \ldots N_d$ points, with periodic boundary conditions, one defines an orthogonal basis $(\varphi_k)_k$ by the same tensor product formula as (2.12) but using the 1-D discrete Fourier basis (2.8)

$$\forall\,(k, n) \in \mathbb{Y}_d^2, \quad \varphi_k(n) = \varphi_{k_1}(n_1)\ldots\varphi_{k_d}(n_d) = \prod_{\ell=1}^{d} e^{\frac{2\mathrm{i}\pi}{N_\ell}k_\ell n_\ell} = e^{2\mathrm{i}\pi\langle k,\,n\rangle_{\mathbb{Y}_d}} \tag{2.13}$$

where we used the (rescaled) inner product

$$\langle k,\; n\rangle_{\mathbb{Y}_d} \stackrel{\text{def.}}{=} \sum_{\ell=1}^{d} \frac{k_\ell n_\ell}{N_\ell}. \tag{2.14}$$

The pairing in (2.14) acts on frequency and spatial index vectors. The basis $(\varphi_k)_k$ is orthonormal for the signal inner product $N^{-1}\sum_n f_n \bar{g}_n$. As in one dimension, we define the Fourier coefficients without normalizing by $(N_\ell)_\ell$:

$$\forall\,k \in \mathbb{Y}_d, \quad \hat{f}_k \stackrel{\text{def.}}{=} \sum_{n\in\mathbb{Y}_d} f_n e^{-2\pi\mathrm{i}\langle k,\,n\rangle_{\mathbb{Y}_d}},$$

$$\forall\,n \in \mathbb{Y}_d, \quad f_n = \frac{1}{N}\sum_{k\in\mathbb{Y}_d} \hat{f}_k e^{2\pi\mathrm{i}\langle k,\,n\rangle_{\mathbb{Y}_d}}.$$

Fast Fourier Transform. We describe the algorithm for $d = 2$; the same construction works in any finite dimension. A fast algorithm for orthogonal decompositions in two 1-D bases $(\varphi^1_{k_1})_{k_1=1}^{N_1}$, $(\varphi^2_{k_2})_{k_2=1}^{N_2}$ also gives a fast decomposition in the tensor product basis $(\varphi^1_{k_1} \otimes \varphi^2_{k_2})_{k_1,k_2}$: apply the algorithm first to the rows and then to the columns of the matrix $(f_n)_{n=(n_1,n_2)} \in \mathbb{R}^{N_1\times N_2}$. Either order gives the same result. Indeed

$$\forall\,k = (k_1, k_2), \quad \langle f,\, \varphi^1_{k_1} \otimes \varphi^2_{k_2}\rangle = \sum_{n=(n_1,n_2)} f_n \overline{\varphi^1_{k_1}(n_1)}\,\overline{\varphi^2_{k_2}(n_2)} = \sum_{n_1}\left(\sum_{n_2} f_{n_1,n_2}\overline{\varphi^2_{k_2}(n_2)}\right)\overline{\varphi^1_{k_1}(n_1)}.$$

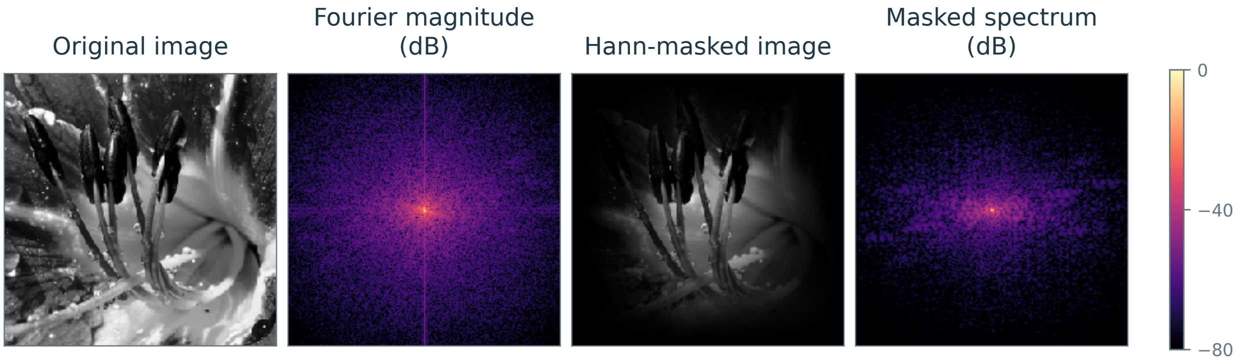

Figure 2.17. 2-D Fourier analysis of an image (left), and attenuation of the periodicity artifact using masking (right).

If $C(N_1)$ is the cost of the 1-D algorithm on $\mathbb{R}^{N_1}$, the resulting 2-D transform costs $N_2C(N_1) + N_1C(N_2)$. For the FFT, this is $O(N_1N_2 \log(N_1N_2)) = O(N \log(N))$, where $N = N_1N_2$.

Represent $f \in \mathbb{R}^{N_1\times N_2}$ as a matrix, and write $F_N = (e^{-2i\pi kn/N})_{k,n}$ for the Fourier matrix, whose rows are the φ_k^*. The 2-D transform can then be expressed as matrix products:

$$\hat{f} = F_{N_1} \times f \times F_{N_2}^\top \in \mathbb{C}^{N_1\times N_2}.$$

In practice, the FFT performs these operations without explicit matrix multiplication.

`Associated code: coding/test_fft2.m`

2.5.3 Shannon sampling theorem.

Theorem 1.2 extends to sampling on a uniform Cartesian grid in $\mathbb{R}^d$ by tensorization. Use the continuous representative of the bandlimited function f. In two dimensions, if $\mathrm{supp}(\hat{f}) \subset [-\pi/s_1, \pi/s_1] \times [-\pi/s_2, \pi/s_2]$ and f decays sufficiently fast,

$$\forall\, x \in \mathbb{R}^2, \quad f(x) = \sum_{n\in\mathbb{Z}^2} f(n_1s_1, n_2s_2)\,\mathrm{sinc}(x_1/s_1 - n_1)\,\mathrm{sinc}(x_2/s_2 - n_2) \quad \text{where} \quad \mathrm{sinc}(u) = \frac{\sin(\pi u)}{\pi u}.$$

2.5.4 Convolution in higher dimension.

Convolution on $\mathbb{X}^d$, for $\mathbb{X} = \mathbb{R}$ or $\mathbb{X} = \mathbb{R}/2\pi\mathbb{Z}$, is defined as in one dimension:

$$f \star g(x) = \int_{\mathbb{X}^d} f(t)g(x-t)\mathrm{d}t.$$

Similarly, finite discrete convolution of vectors $f \in \mathbb{R}^{N_1\times N_2}$ extends formula (2.11) as

$$\forall\, n \in [\![0, N_1-1]\!] \times [\![0, N_2-1]\!], \quad (f \star g)_n \stackrel{\text{def.}}{=} \sum_{k_1=0}^{N_1-1} \sum_{k_2=0}^{N_2-1} f_k g_{n-k}$$

where additions and subtractions of vectors are performed modulo (N_1, N_2).

The Fourier-convolution identity $\mathcal{F}(f \star g) = \hat{f} \odot \hat{g}$ holds in each of these settings. In the finite case, it yields an $O(N \log(N))$ convolution algorithm even when f and g have large support.

2.6 Application to ODEs and PDEs

2.6.1 On Continuous Domains

We give the main calculations, leaving analytic details of existence and regularity aside.

On $\mathbb{X} = \mathbb{R}$ or $\mathbb{T}$, one has

$$\mathcal{F}(f^{(k)})(\omega) = (i\omega)^k \hat{f}(\omega).$$

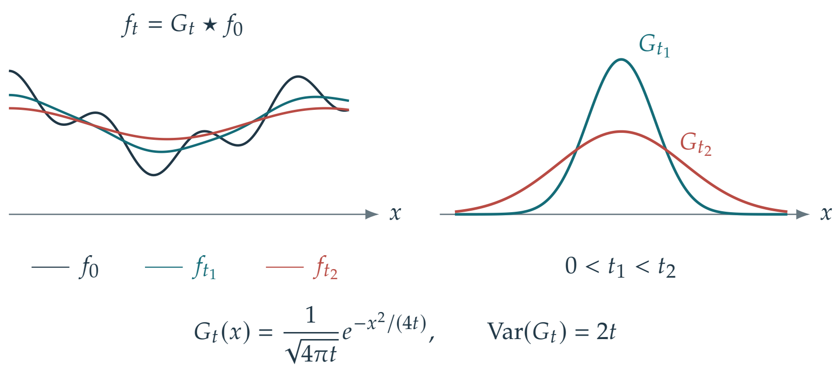

Figure 2.18. Heat diffusion as a convolution.

Intuitively, $f^{(k)} = f \star \delta^{(k)}$ where $\delta^{(k)}$ is a distribution with Fourier transform $\hat{\delta}^{(k)}(\omega) = (\mathrm{i}\omega)^k$. Similarly on $\mathbb{X} = \mathbb{R}^d$ (see Section 2.5 for the definition of the Fourier transform in dimension d), one has

$$\mathcal{F}(\Delta f)(\omega) = -\|\omega\|^2 \hat{f}(\omega) \tag{2.15}$$

The same identity holds on $\mathbb{T}^d$, with ω replaced by $n \in \mathbb{Z}^d$. Fourier transforms and Fourier coefficients provide a powerful tool for studying linear differential equations with constant coefficients by turning them into algebraic equations.

As a typical example, we consider the heat equation

$$\frac{\partial f_t}{\partial t} = \Delta f_t \quad \Longrightarrow \quad \forall\, \omega, \quad \frac{\partial \hat{f}_t(\omega)}{\partial t} = -\|\omega\|^2 \hat{f}(\omega).$$

Hence $\hat{f}_t(\omega) = \hat{f}_0(\omega) e^{-\|\omega\|^2 t}$. Applying the inverse Fourier transform and the convolution theorem gives

$$f_t = G_t \star f_0 \quad \text{where} \quad G_t = \frac{1}{(4\pi t)^{d/2}} e^{-\frac{\|x\|^2}{4t}}$$

On $\mathbb{R}^d$, G_t is a Gaussian density with standard deviation $\sqrt{2t}$ in each direction. On the torus, periodize this kernel and adjust its mass to the chosen Haar normalization.

2.6.2 Finite Domain and Discretization

On $\mathbb{Z}/N\mathbb{Z}$, corresponding to a discrete periodic domain, consider the forward first difference and the centered second difference:

$$D_1 f \overset{\text{def.}}{=} N(f_{n+1} - f_n)_n = f \star d_1 \quad \text{where} \quad d_1 = N[-1, 0, \dots, 0, 1]^\top \in \mathbb{R}^N, \tag{2.16}$$

$$D_2 f = -D_1^\top D_1 f \overset{\text{def.}}{=} N^2(f_{n+1} + f_{n-1} - 2f_n)_n = f \star d_2 \quad \text{where} \quad d_2 = -d_1 \star \bar{d}_1 = N^2[-2, 1, 0, \dots, 0, 1]^\top \in \mathbb{R}^N. \tag{2.17}$$

Thanks to Proposition 2.5, one can equivalently compute

$$\mathcal{F}(D_2 f) = \hat{d}_2 \odot \hat{f} \quad \text{where} \quad (\hat{d}_2)_k = N^2(e^{\frac{2i\pi}{N}k} + e^{-\frac{2i\pi}{N}k} - 2) = -4N^2 \sin\left(\frac{\pi k}{N}\right)^2. \tag{2.18}$$

Here $\bar{d}_1[n] = d_1[-n]$ denotes periodic reflection. For fixed nonzero k and $N \to \infty$, one has $(\hat{d}_2)_k \sim -(2\pi k)^2$ which matches the scaling of (2.15).

2.7 A Bit of Group Theory

The reference for this section is [2].

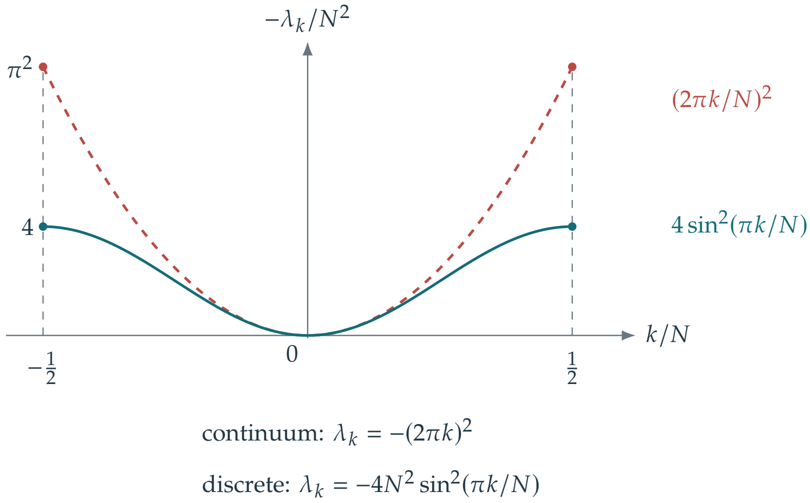

Figure 2.19. Comparison of the spectra of Δ and D_2.

2.7.1 Characters

For $(G,+)$ a commutative group, a character is a group morphism $\chi : (G,+) \to (\mathbb{C}^*,\cdot)$, i.e. it satisfies

$$\forall\,(n,m) \in G, \quad \chi(n+m) = \chi(n)\chi(m).$$

The characters form the dual group $(\hat{G},\odot)$ under pointwise multiplication $(\chi_1 \odot \chi_2)(n) \overset{\text{def.}}{=} \chi_1(n)\chi_2(n)$. The inverse of a character χ is $\chi^{-1}(n) = \chi(-n)$.

For a finite group G with $|G| = N$, we have $N \times n = 0$ for every $n \in G$. Consequently, $\chi(n)^N = \chi(Nn) = \chi(0) = 1$: characters take values among the roots of unity,

$$\chi(n) \in \left\{ e^{\frac{2i\pi}{N}k} \;;\; 0 \leqslant k \leqslant N-1 \right\}. \tag{2.19}$$

Thus $\hat{G}$ is a finite group (since there are finitely many maps between two finite sets) and $\chi^{-1} = \bar{\chi}$. In the case of a cyclic group, the dual is simple to describe.

Proposition 2.7. *For $G = \mathbb{Z}/N\mathbb{Z}$, the dual is $\hat{G} = (\varphi_k)_{k=0}^{N-1}$, where $\varphi_k = (e^{\frac{2i\pi}{N}nk})_n$. The map $k \mapsto \varphi_k$ defines a noncanonical isomorphism $G \sim \hat{G}$.*

Proof. The φ_k are indeed characters.

Conversely, for any $\chi \in \hat{G}$, according to (2.19), $\chi(1) = e^{\frac{2i\pi}{N}k}$ for some k. Then

$$\chi(n) = \chi(1)^n = e^{\frac{2i\pi}{N}kn} = \varphi_k(n).$$

These characters are distinct because the values $\varphi_k(1)$ are distinct. Hence $|G| = |\hat{G}|$, and the displayed map is an isomorphism. □

This proposition thus shows that characters of cyclic groups are exactly the orthogonal discrete Fourier basis defined in (2.8).

Commutative groups. The structure theorem for finitely generated abelian groups expresses every such group as a product of cyclic groups:

$$G \sim (\mathbb{Z}/N_1\mathbb{Z}) \times \ldots \times (\mathbb{Z}/N_d\mathbb{Z}) \times \mathbb{Z}^Q. \tag{2.20}$$

If G is finite, then $Q = 0$ and $N = N_1 \cdots N_d$. In this case, G is simply a discrete d-dimensional "rectangle" with periodic boundary conditions.

For two finite groups (G_1, G_2) one has

$$\widehat{G_1 \times G_2} = \hat{G}_1 \otimes \hat{G}_2 = \left\{ \chi_1 \otimes \chi_2 \;;\; (\chi_1,\chi_2) \in \hat{G}_1 \times \hat{G}_2 \right\}. \tag{2.21}$$

Here $\otimes$ is the tensor product of two functions

$$\forall (n_1, n_2) \in G_1 \times G_2, \quad (\chi_1 \otimes \chi_2)(n_1, n_2) \overset{\text{def.}}{=} \chi_1(n_1)\chi_2(n_2).$$

The tensor product $\chi_1 \otimes \chi_2$ is a character. Conversely, every character has the factorization $\chi = \chi(\cdot, 0) \otimes \chi(0, \cdot)$, since $(n_1, n_2) = (n_1, 0) + (0, n_2)$.

Combining this factorization with the structure theorem proves the following isomorphism.

Proposition 2.8. *If G is commutative and finite then $\hat{G} \sim G$.*

Proof. The structure theorem (2.20) with $Q = 0$ and the product-duality identity (2.21) give

$$\hat{G} \sim \hat{G}_1 \otimes \ldots \otimes \hat{G}_d$$

where $G_\ell \overset{\text{def.}}{=} \mathbb{Z}/N_\ell\mathbb{Z}$. One then observes that $\hat{G}_1 \otimes \hat{G}_2 \sim \hat{G}_1 \times \hat{G}_2$. Proposition 2.7 then completes the proof, since $\hat{G}_k \sim G_k$. □

The isomorphism $\hat{G} \sim G$ is noncanonical: it depends on a choice of generators and corresponding roots of unity. As in vector-space duality, the double-dual isomorphism $\hat{\hat{G}} \sim G$ is canonical and is given by evaluation:

$$g \in G \longmapsto e_g \in \hat{\hat{G}} \quad \text{where} \quad \left(e_g : \chi \in \hat{G} \mapsto \chi(g) \in \mathbb{C}^*.\right)$$

Discrete Fourier transform from the Character Viewpoint. This construction also identifies the Fourier basis explicitly. The characters in $\hat{G}_\ell$ are the discrete Fourier atoms (2.8), of the form

$$(e^{\frac{2i\pi}{N_\ell} k_\ell n_\ell})_{n_\ell=0}^{N_\ell-1} \quad \text{for some} \quad 0 \leqslant k_\ell < N_\ell.$$

Identifying G and $G_1 \times \ldots \times G_d$, their tensor products show that the characters in $\hat{G}$ form exactly the orthogonal multidimensional Fourier basis (2.13).

2.7.2 More General cases

Infinite groups. For an infinite finitely generated abelian group, the structure theorem has $Q > 0$. For locally compact abelian groups, the dual consists of continuous characters with values on the unit circle; invariant Haar measure determines integration. For $G = \mathbb{Z}$, the characters have a continuous frequency parameter,

$$\hat{\mathbb{Z}} = \left\{\varphi_\omega : n \mapsto e^{in\omega} \in \mathbb{C}^{\mathbb{Z}} \;;\; \omega \in \mathbb{R}/2\pi\mathbb{Z}\right\}$$

so that $\hat{\mathbb{Z}} \sim \mathbb{R}/2\pi\mathbb{Z}$. The case $G = \mathbb{Z}^Q$ follows by tensorization. The $(\varphi_\omega)_\omega$ are "orthogonal" in the sense that $\langle \varphi_\omega, \varphi_{\omega'} \rangle_{\mathbb{Z}} = 2\pi\delta_{2\pi}(\omega - \omega')$ can be understood as a Dirac kernel (this is similar to the Poisson formula), where $\langle u, v \rangle_{\mathbb{Z}} \overset{\text{def.}}{=} \sum_n u_n \bar{v}_n$. Expanding a sequence $(c_n)_{n\in\mathbb{Z}}$ in characters amounts to forming the Fourier series $\sum_n c_n e^{-in\omega}$.

Similarly, for $G = \mathbb{R}/2\pi\mathbb{Z}$, one has $\hat{G} = \mathbb{Z}$, with orthonormal characters $\varphi_n = e^{i\cdot n}$, so that the decomposition of functions in $L^2(G)$ is the computation of Fourier coefficients. On a compact group, normalized invariant measure provides the natural L^2 inner product.

Non-commutative groups. For a noncommutative group, one-dimensional characters do not recover G. For example, the symmetric group Σ_N on $N \geqslant 2$ letters has only the characters $\hat{G} = \{1, \varepsilon\}$, where $\varepsilon(\sigma) = (-1)^q$ is the signature and q is the number of transpositions in a decomposition of $\sigma \in \Sigma_N$.

To study noncommutative groups, replace homomorphisms $\chi : G \to \mathbb{C}^*$ by homomorphisms $\rho : G \to \mathrm{GL}(\mathbb{C}^{n_\rho})$ for some n_ρ. These are called representations of G. For $(g, g') \in G$, writing the group operation multiplicatively on G, the defining identity is $\rho(gg') = \rho(g) \circ \rho(g')$. When $n_\rho = 1$, the identification $\mathrm{GL}(\mathbb{C}) \sim \mathbb{C}^*$ recovers the one-dimensional characters. For any representation ρ, the function $\chi(g) \overset{\text{def.}}{=} \mathrm{tr}(\rho(g))$, where tr denotes trace, is called its character. This trace function is generally not a group homomorphism.

For a finite group G, every subspace V invariant under all $\rho(g)$ admits an invariant complement W, so that $\mathbb{C}^{n_\rho} = V \oplus W$ with W also invariant. In a basis adapted to these subspaces, the representation matrices are block diagonal. To construct the complement, use the inner product

$$\langle x, y \rangle \overset{\text{def.}}{=} \sum_{g\in G} \langle \rho(g)x, \rho(g)y \rangle_{\mathbb{C}^{n_\rho}}$$

and take the orthogonal complement $V^\perp$ for this inner product. In an orthonormal basis of $\mathbb{C}^{n_\rho}$ for this product, all matrices $\rho(g)$ are unitary. A representation is irreducible if it has no nonzero proper invariant subspace. These are the elementary representations from which larger ones are built by block-diagonal assembly. We identify irreducible representations up to isomorphism: (ρ,ρ') are isomorphic if $\rho'(g) = U^{-1}\rho(g)U$ for every g, with a single change-of-basis matrix $U \in \mathrm{GL}(\mathbb{C}^d)$. Their isomorphism classes form the unitary dual, denoted $\hat{G}$, which is generally a set rather than a group.

For the symmetric group, there is an explicit description of the set of irreducible representations. For example, $G = \Sigma_3$ has $|G| = 6$ elements. It has two one-dimensional representations, the trivial representation and the signature, and one two-dimensional irreducible representation obtained by identifying G with the isometries of an equilateral triangle in $\mathbb{R}^2$. Their dimensions satisfy $6 = |G| = \sum n_\rho^2 = 1 + 1 + 2^2$.

The dimensions n_ρ of the irreducible representations $\rho \in \hat{G}$ satisfy $\sum_{\rho\in\hat{G}} n_\rho^2 = N$. After choosing unitary representatives, their matrix entries form an orthogonal basis for functions $f : G \to \mathbb{C}$. This basis is noncanonical because it depends on a choice of basis within each representation. The associated Fourier transform of a function $f : G \to \mathbb{C}$ is defined as

$$\forall\, \rho \in \hat{G}, \quad \hat{f}(\rho) \overset{\text{def.}}{=} \sum_{g\in G} f(g)\rho(g) \in \mathbb{C}^{n_\rho\times n_\rho}.$$

With this sign convention, these are pairings with conjugated matrix entries. Their orthogonality makes the transform invertible. The next proposition gives the inverse transform.

Proposition 2.9. *One has*

$$\forall\, g \in G, \quad f(g) = \frac{1}{|G|}\sum_{\rho\in\hat{G}} n_\rho \operatorname{tr}(\hat{f}(\rho)\rho(g^{-1})).$$

Proof. The proof uses the regular-character identity, valid for every $g \in G$

$$A_g \overset{\text{def.}}{=} \sum_{\rho\in\hat{G}} n_\rho \operatorname{tr}(\rho(g)) = |G|\delta_g$$

where $\delta_g = 0$ for $g \neq \mathrm{Id}_G$ and $\delta_{\mathrm{Id}_G} = 1$. A proof is given in the book by Diaconis, p. 13: decompose the character of the regular representation into irreducible characters and use their orthogonality. The regular representation acts by permuting, according to G, the canonical basis of $\mathbb{R}^{|G|}$. One then has

$$\begin{aligned}\frac{1}{|G|}\sum_{\rho\in\hat{G}} n_\rho \operatorname{tr}(\hat{f}(\rho)\rho(g^{-1})) &= \frac{1}{|G|}\sum_{\rho\in\hat{G}} n_\rho \operatorname{tr}(\sum_u f(u)\rho(u)\rho(g^{-1})) \\ &= \sum_u f(u)\frac{1}{|G|}\sum_{\rho\in\hat{G}} n_\rho \operatorname{tr}(\rho(ug^{-1})) = \sum_u f(u)A_{ug^{-1}}/|G| = f(g).\end{aligned}$$

□

One can define the convolution of two functions $f, h : G \to \mathbb{C}$ as

$$(f \star h)(a) \overset{\text{def.}}{=} \sum_{bc=a} f(b)h(c)$$

Convolution need not be commutative on this group. The Fourier transform block-diagonalizes the convolution operators, as stated next.

Proposition 2.10. *Denoting $\mathcal{F}(f) = \hat{f}$*

$$\mathcal{F}(f \star h)(\rho) = \hat{f}(\rho) \times \hat{h}(\rho)$$

where $\times$ denotes matrix multiplication.

Proof.

$$\mathcal{F}(f\star h)(\rho) = \sum_x\sum_y f(y)h(y^{-1}x)\rho(yy^{-1}x) = \sum_x\sum_y f(y)\rho(y)h(y^{-1}x)\rho(y^{-1}x) = \sum_y f(y)\rho(y)\sum_z h(z)\rho(z)$$

where we made the change of variable $z = y^{-1}x$. □

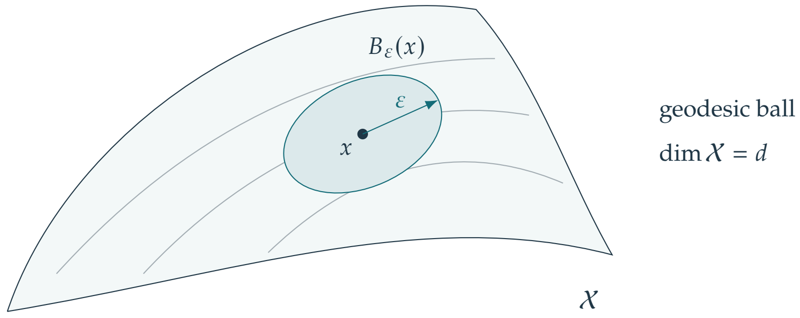

$$\frac{1}{\mathrm{Vol}(B_\varepsilon(x))}\int_{B_\varepsilon(x)} f(y)\,\mathrm{d}\mu(y) = f(x) + \frac{\varepsilon^2}{2(d+2)}\Delta f(x) + o(\varepsilon^2)$$

Figure 2.20. Computing the Laplacian on a surface.

Fast algorithms for computing $\hat{f}$ exist for some groups, including permutation groups. Their structure is considerably more involved than for $G = \mathbb{Z}/N\mathbb{Z}$.

For infinite compact groups, the theory uses a discrete infinite family of irreducible representations. The rotation group SO(3) is a useful example. Its representations can be described explicitly through the spherical harmonics introduced in Section 2.8.2. There is one representation of dimension $2\ell + 1$ for each frequency index ℓ.

2.8 A Bit of Spectral Theory

On a general domain $\mathbb{X}$, the group-theoretic approach applies when $\mathbb{X} = G$ is a group or when a group acts transitively on $\mathbb{X}$. Another approach defines Fourier-like basis functions as eigenfunctions of a differential operator, following the characterization in Section 2.6. The Laplacian is particularly useful: it is second order, rotation invariant, and has natural counterparts on surfaces and graphs.

2.8.1 On a Surface or a Manifold

Let $\mathbb{X}$ be a smooth compact connected Riemannian manifold of dimension d, without boundary. The Laplace–Beltrami operator can be characterized, for smooth f, by the small-ball mean-value expansion

$$(\Delta f)(x) = \lim_{\varepsilon\to 0} \frac{2(d+2)}{\varepsilon^2}\left(\frac{1}{\mathrm{Vol}(B_\varepsilon(x))}\int_{B_\varepsilon(x)} f(y)\,\mathrm{d}\mu(y) - f(x)\right).$$

Here μ is Riemannian volume and $B_\varepsilon(x)$ is the geodesic ball centered at x. The factor ε^{-2} is essential.

The operator Δ is unbounded and self-adjoint on its natural domain in $L^2(\mathbb{X})$; its resolvent is compact. It has a complete orthonormal family of eigenfunctions $(\varphi_n)_{n\geqslant 0}$, with eigenvalues, counted with multiplicity,

$$0 = \lambda_0 > \lambda_1 \geqslant \lambda_2 \geqslant \cdots \longrightarrow -\infty.$$

The constant eigenfunction is $\varphi_0 = \mathrm{Vol}(\mathbb{X})^{-1/2}$, and the inner product is $\langle f,\, g\rangle_{\mathbb{X}} = \int_{\mathbb{X}} f(x)\overline{g(x)}\,\mathrm{d}\mu(x)$. Boundary conditions must be specified if the manifold has a boundary.

On the flat unit torus $\mathbb{X} = (\mathbb{R}/\mathbb{Z})^d$,

$$\Delta f = \sum_{s=1}^{d} \frac{\partial^2 f}{\partial x_s^2}, \qquad \varphi_n(x) = e^{2\pi\mathrm{i}\langle n,\, x\rangle}, \qquad \lambda_n = -4\pi^2\|n\|^2, \quad n \in \mathbb{Z}^d.$$

2.8.2 Spherical Harmonics

Consider the $(d-1)$-dimensional sphere $\mathbb{S}^{d-1} = \left\{x \in \mathbb{R}^d \;;\; \|x\|_{\mathbb{R}^d} = 1\right\}$. The Laplacian has explicit eigenfunctions. For $d = 3$, they are indexed by $n = (\ell, m)$:

$$\forall\, \ell \in \mathbb{N}, \quad \forall\, m = -\ell, \ldots, \ell, \quad \varphi_{\ell,m}(\theta,\varphi) = e^{\mathrm{i}m\varphi} P_\ell^m(\cos(\theta))$$

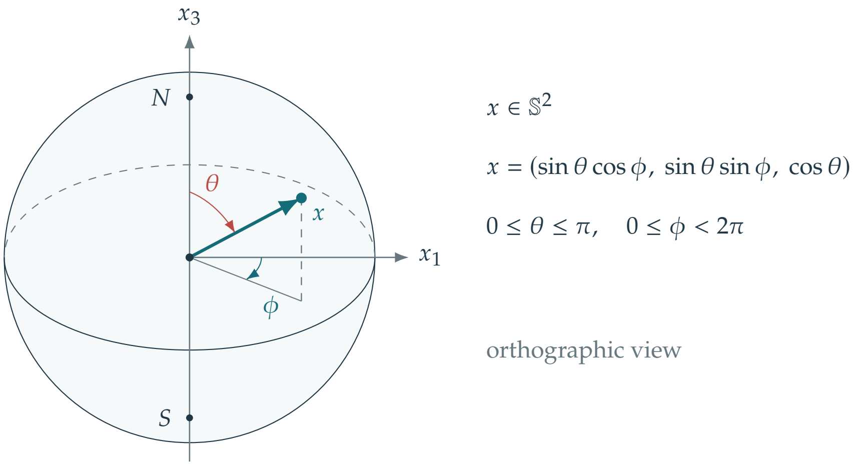

Figure 2.21. Spherical coordinates.

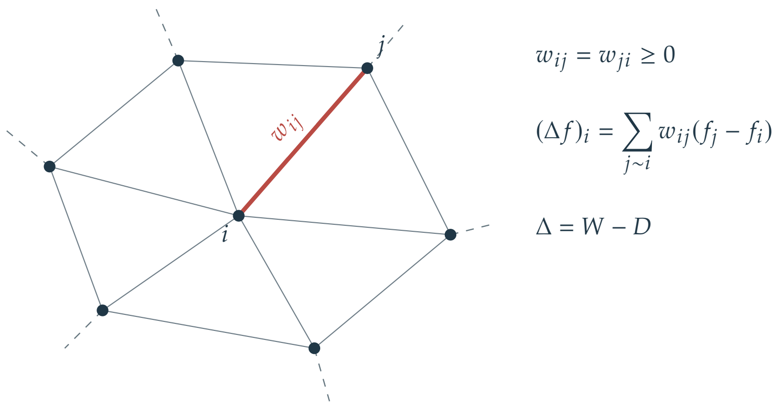

Figure 2.22. Weighted graph.

The corresponding eigenvalue is $\lambda_{\ell,m} = -\ell(\ell+1)$. Here P_ℓ^m are associated Legendre functions (polynomials multiplied by $(1-x^2)^{|m|/2}$). Normalization makes the eigenfunctions orthonormal. We use spherical coordinates $x = (\sin\theta\cos\varphi, \sin\theta\sin\varphi, \cos\theta) \in \mathbb{S}^2$ for $(\theta, \varphi) \in [0,\pi]\times[0,2\pi]$. The index ℓ plays the role of the magnitude of a two-dimensional Fourier frequency. For a fixed ℓ, the space $V_\ell = \mathrm{span}(\varphi_{\ell,m})$ is an eigenspace of Δ, and is also invariant under rotation.

2.8.3 On a Graph

Let $\mathbb{X}$ be a graph with N vertices indexed by $\{1,\ldots,N\}$. Its geometry is encoded by the weight matrix $W = (w_{i,j})_{1\leqslant i,j\leqslant N}$. The notation $i \sim j$ means that (i,j) is an edge, for $(i,j) \in \mathbb{X}^2$. Set the weights to zero on nonedges and assume nonnegative, symmetric weights: $w_{i,j} = w_{j,i}$.

The graph Laplacian $\Delta : \mathbb{R}^N \to \mathbb{R}^N$ sums weighted differences between neighboring values

$$\forall f \in \mathbb{R}^N, \quad (\Delta f)_i \overset{\text{def.}}{=} \sum_{j\sim i} w_{i,j} f_j - \Big(\sum_{j\sim i} w_{i,j}\Big) f_i \quad \Longrightarrow \quad \Delta = W - D$$

where $D \overset{\text{def.}}{=} \mathrm{diag}_i(\sum_{j\sim i} w_{i,j})$. In particular, note that $\Delta\mathbb{1} = 0$.

For instance, if $\mathbb{X} = \mathbb{Z}/N\mathbb{Z}$ with the graph $i \sim i-1$ and $i \sim i+1$ (modulo N), then Δ is the finite difference Laplacian operator $\Delta = N^{-2} D_2$ for unit edge weights defined in (2.17). This extends to any dimension by tensorization.

Proposition 2.11. *Define the graph-gradient operator $G : f \in \mathbb{R}^N \mapsto (\sqrt{w_{i,j}}(f_i - f_j))_{(i,j)\in E}$, with one component per undirected edge. Then*

$$-\Delta = G^\top G \quad \Longrightarrow \quad \forall f \in \mathbb{R}^N, \quad \langle \Delta f, f\rangle_{\mathbb{R}^N} = -\langle Gf, Gf\rangle_{\mathbb{R}^P}.$$

where P is the number of undirected edges $E = \{(i, j)\ ;\ i \sim j, i < j\}$.

Proof. By symmetry of W and counting each undirected edge once,

$$\begin{aligned}
\|Gf\|^2 &= \sum_{i<j} w_{ij}(f_i - f_j)^2 \\
&= \sum_i f_i^2 \sum_j w_{ij} - \sum_{i,j} w_{ij} f_i f_j \\
&= \langle Df, f\rangle - \langle Wf, f\rangle = -\langle \Delta f, f\rangle.
\end{aligned}$$

Equality of these quadratic forms proves $G^\top G = -\Delta$. □

Thus Δ is symmetric negative semidefinite and has an orthonormal eigenbasis. One can take $\varphi_1 = N^{-1/2}\mathbb{1}$, with $0 = \lambda_1 \geqslant \lambda_2 \geqslant \cdots \geqslant \lambda_N$. The multiplicity of zero is the number of connected components of the graph formed by the positive-weight edges; for a connected graph with $N > 1$, $\lambda_2 < 0$. On a uniform periodic grid, the Fourier basis is an eigenbasis.

The quadratic energy $\sum_{i<j} w_{ij}(f_i - f_j)^2$ is the discrete analogue of the Dirichlet energy and provides a natural measure of smoothness on a graph.

2.8.4 Further Applications

Polynomial multiplication provides another application (`Code text_polymult.m`). Analogous transforms can also be defined over suitable finite fields.

References

[1] Stephane Mallat. *A wavelet tour of signal processing: the sparse way*. Academic press, 2008.

[2] Gabriel Peyré. *L'algèbre discrète de la transformée de Fourier*. Ellipses, 2004.