A Wavelet Tour of Signal Processing

Stéphane Mallat

October 9, 2008

À la mémoire de mon père, Alexandre.
Pour ma mère, Francine.

Preface to the Sparse Edition

I can not help but find striking resemblances between scientific communities and schools of fish. We interact in conferences and through articles, we move together while a global trajectory emerges from individual contributions. Some of us like to be at the center of the school, others prefer to wander around, and few swim in multiple directions in front. To avoid dying by starvation in a progressively narrower and specialized domain, a scientific community needs to move on. Computational harmonic analysis is still well alive because it went beyond wavelets. Writing such a book is about decoding the trajectory of the school, and gathering the pearls that have been uncovered on the way. Wavelets are not any more the central topic, despite the original title. It is just an important tool, as the Fourier transform is. Sparse representation and processing are now at the core.

In the 80's, many researchers were focused on building time-frequency decompositions, trying to avoid the uncertainty barrier, and hoping to discover the ultimate representation. Along the way came the construction of wavelet orthogonal bases, which opened new perspectives through collaborations with physicists and mathematicians. Designing orthogonal bases with Xlets became a popular sport, with compression and noise reduction applications. Connections with approximations and sparsity also became more apparent. The search for sparsity has taken over, leading to new grounds, where orthonormal bases are replaced by redundant dictionaries of waveforms. Geometry is now also becoming more apparent through sparse approximation supports in dictionaries.

During these last 7 years, I also encountered the industrial world. With a lot of naiveness, some bandlets and more mathematics, we created a start-up with Christophe Bernard, Jérome Kalifa and Erwan Le Pennec. It took us some time to learn that in 3 months good engineering should produce robust algorithms that operate in real time, as opposed to the 3 years we were used to have for writing new ideas with promising perspectives. Yet, we survived because mathematics is a major source of industrial innovations for signal processing. Semi-conductor technology offers amazing computational power and flexibility. However, ad-hoc algorithms often do not scale easily and mathematics accelerates the trial and error development process. Sparsity decreases computations, memory and data communications. Although it brings beauty, mathematical understanding is not a luxury. It is required by increasingly sophisticated information processing devices.

New Additions Putting sparsity at the center of the book implied rewriting many parts and adding sections. Chapter 12 and Chapter 13 are new. They introduce sparse representations in redundant dictionaries, and inverse problems, super-resolution and compressive sensing. Here is a small catalogue of new elements in this third edition.
- Radon transform and tomography.
- Lifting for wavelets on surfaces, bounded domains and fast computations.
- JPEG-2000 image compression.
- Block thresholding for denoising.
- Geometric representations with adaptive triangulations, curvelets and bandlets.
- Sparse approximations in redundant dictionaries with pursuits algorithms.
- Noise reduction with model selection, in redundant dictionaries.
- Exact recovery of sparse approximation supports in dictionaries.
- Multichannel signal representations and processing.
- Dictionary learning.
- Inverse problems and super-resolution.

- Compressive sensing.
- Source separation.

Teaching   This book is intended as a graduate textbook. Its evolution is also the result of teaching courses in electrical engineering and applied mathematics. A new web site provides softwares for reproducible experimentations, exercise solutions, together with teaching material such as slides with figures, and Matlab softwares for numerical classes: *http://wavelet-tour.com.*

More exercises have been added at the end of each chapter, ordered by level of difficulty. Level[1] exercises are direct applications of the course. Level[2] requires more thinking. Level[3] includes some technical derivations. Level[4] are projects at the interface of research, that are possible topics for a final course project or an independent study. More exercises and projects can be found in the web site.

Sparse Course Programs   The Fourier transform and analog to digital conversion through linear sampling approximations provide a common ground for all courses (Chapters 2 and 3). It introduces basic signal representations, and reviews important mathematical and algorithmic tools needed afterwards. Many trajectories are then possible to explore and teach sparse signal processing. The following list gives several topics that can orient the course structure, with elements that can be covered along the way.

Sparse representations with bases and applications
- Principles of linear and non-linear approximations in bases (Chapter 9).
- Lipschitz regularity and wavelet coefficients decay (Chapter 6).
- Wavelet bases (Chapter 7).
- Properties of linear and non-linear wavelet basis approximations (Chapter 9).
- Image wavelet compression (Chapter 10).
- Linear and non-linear diagonal denoising (Chapter 11).

Sparse time-frequency representations
- Time-frequency wavelet and windowed Fourier ridges for audio processing (Chapter 4).
- Local cosine bases (Chapter 8).
- Linear and non-linear approximations in bases (Chapter 9).
- Audio compression (Chapter 10).
- Audio denoising and block thresholding (Chapter 11).
- Compression and denoising in redundant time-frequency dictionaries, with best bases or pursuit algorithms (Chapter 12).

Sparse signal estimation
- Bayes versus minimax, and linear versus non-linear estimations (Chapter 11).
- Wavelet bases (Chapter 7).
- Linear and non-linear approximations in bases (Chapter 9).
- Thresholding estimation (Chapter 11).
- Minimax optimality (Chapter 11).
- Model selection for denoising in redundant dictionaries (Chapter 12).
- Compressive sensing (Chapter 13).

Sparse compression and information theory
- Wavelet orthonormal bases (Chapter 7).
- Linear and non-linear approximations in bases (Chapter 9).
- Compression and sparse transform codes in bases (Chapter 10).
- Compression in redundant dictionaries (Chapter 12).
- Compressive sensing (Chapter 13).
- Source separation (Chapter 13).

Dictionary representations and inverse problems
- Frames and Riesz bases (Chapter 5).
- Linear and non-linear approximations in bases (Chapter 9).
- Ideal redundant dictionary approximations (Chapter 12).
- Pursuit algorithms and dictionary incoherence (Chapter 12).
- Linear and thresholding inverse estimators (Chapter 13).
- Super-resolution and source separation (Chapter 13).
- Compressive sensing (Chapter 13).

Geometric sparse processing
- Time-frequency spectral lines and ridges (Chapter 4).
- Frames and Riesz bases (Chapter 5).
- Multiscale edge representations with wavelet maxima (Chapter 6).
- Sparse approximation supports in bases (Chapter 9).
- Approximations with geometric regularity, curvelets and bandlets (Chapters 9 and 12).
- Sparse signal compression and geometric bit budget (Chapters 10 and 12).
- Exact recovery of sparse approximation supports (Chapter 12).
- Super-resolution (Chapter 13).

Contents

Notations

$\langle f, g \rangle$	Inner product (A.6).
$\|f\|$	Euclidean or Hilbert space norm.
$\|f\|_1$	$\mathbf{L^1}$ or $\mathbf{l^1}$ norm.
$\|f\|_\infty$	$\mathbf{L}^\infty$ norm.
$f[n] = O(g[n])$	Order of: there exists K such that $f[n] \leqslant Kg[n]$.
$f[n] = o(g[n])$	Small order of: $\lim_{n \to +\infty} \frac{f[n]}{g[n]} = 0$.
$f[n] \sim g[n]$	Equivalent to: $f[n] = O(g[n])$ and $g[n] = O(f[n])$.
$A < +\infty$	A is finite.
$A \gg B$	A is much bigger than B.
z^*	Complex conjugate of $z \in \mathbb{C}$.
$\lfloor x \rfloor$	Largest integer $n \leqslant x$.
$\lceil x \rceil$	Smallest integer $n \geqslant x$.
$(x)_+$	$\max(x, 0)$.
$n \bmod N$	Remainder of the integer division of n modulo N.

Sets

$\mathbb{N}$	Positive integers including 0.
$\mathbb{Z}$	Integers.
$\mathbb{R}$	Real numbers.
$\mathbb{R}^+$	Positive real numbers.
$\mathbb{C}$	Complex numbers.
$|\Lambda|$	Number of elements in a set Λ.

Signals

$f(t)$	Continuous time signal.
$f[n]$	Discrete signal.
$\delta(t)$	Dirac distribution (A.27).
$\delta[n]$	Discrete Dirac (3.32).
$\mathbf{1}_{[a,b]}$	Indicator function which is 1 in $[a, b]$ and 0 outside.

Spaces

$\mathbf{C}_0$	Uniformly continuous functions (7.207).
$\mathbf{C}^p$	p times continuously differentiable functions.
$\mathbf{C}^\infty$	Infinitely differentiable functions.
$\mathbf{W}^s(\mathbb{R})$	Sobolev s times differentiable functions (9.8).
$\mathbf{L^2}(\mathbb{R})$	Finite energy functions $\int |f(t)|^2 \, dt < +\infty$.
$\mathbf{L^p}(\mathbb{R})$	Functions such that $\int |f(t)|^p \, dt < +\infty$.
$\ell^2(\mathbb{Z})$	Finite energy discrete signals $\sum_{n=-\infty}^{+\infty} |f[n]|^2 < +\infty$.
$\ell^p(\mathbb{Z})$	Discrete signals such that $\sum_{n=-\infty}^{+\infty} |f[n]|^p < +\infty$.
$\mathbb{C}^N$	Complex signals of size N.
$\mathbf{U} \oplus \mathbf{V}$	Direct sum of two vector spaces.
$\mathbf{U} \otimes \mathbf{V}$	Tensor product of two vector spaces (A.16).

NullU   Null space of an operator U.
ImU   Image space of an operator U.

Operators
Id   Identity.
$f'(t)$ Derivative $\frac{df(t)}{dt}$.
$f^{(p)}(t)$ Derivative $\frac{d^p f(t)}{dt^p}$ of order p .
$\vec{\nabla} f(x,y)$ Gradient vector (6.51).
$f \star g(t)$ Continuous time convolution (2.2).
$f \star g[n]$ Discrete convolution (3.33).
$f \circledast g[n]$ Circular convolution (3.73)

Transforms
$\hat{f}(\omega)$ Fourier transform (2.6), (3.39).
$\hat{f}[k]$ Discrete Fourier transform (3.49).
$Sf(u,s)$ Short-time windowed Fourier transform (4.11).
$P_S f(u,\xi)$ Spectrogram (4.12).
$Wf(u,s)$ Wavelet transform (4.31).
$P_W f(u,\xi)$ Scalogram (4.55).
$P_V f(u,\xi)$ Wigner-Ville distribution (4.120).

Probability
X Random variable.
$\mathsf{E}\{X\}$ Expected value.
$\mathcal{H}(X)$ Entropy (10.4).
$\mathcal{H}_d(X)$ Differential entropy (10.20).
$\mathrm{Cov}(X_1, X_2)$ Covariance (A.19).
$F[n]$ Random vector.
$R_F[k]$ Autocovariance of a stationary process (A.23).

I

Sparse Representations

Signals carry overwhelming amounts of data in which relevant information is often harder to find than a needle in a haystack. Processing is faster and simpler in a sparse representation where few coefficients reveal the information we are looking for. Such representations can be constructed by decomposing signals over elementary waveforms chosen in a family called a dictionary. But the search for the Holy Grail of an ideal sparse transform adapted to all signals is a hopeless quest. The discovery of wavelet orthogonal bases and local time-frequency dictionaries has opened the door to a huge jungle of new transforms. Adapting sparse representations to signal properties, and deriving efficient processing operators, is therefore a necessary survival strategy.

An orthogonal basis is a dictionary of minimum size, that can yield a sparse representation if designed to concentrate the signal energy over a set of few vectors. This set gives a geometric signal description. Efficient signal compression and noise reduction algorithms are then implemented with diagonal operators, computed with fast algorithms. But this is not always optimal.

In natural languages, a richer dictionary helps to build shorter and more precise sentences. Similarly, dictionaries of vectors that are larger than bases are needed to build sparse representations of complex signals. But choosing is difficult, and requires more complex algorithms. Sparse representations in redundant dictionaries can improve pattern recognition, compression and noise reduction, but also the resolution of new inverse problems. This includes super-resolution, source separation and compressive sensing.

This first chapter is a sparse book representation, providing the story line and main ideas. It gives a sense of orientation, to choose a path for traveling in the book.

1.1 Computational Harmonic Analysis

Fourier and wavelet bases are the starting point of our journey. They decompose signals over oscillatory waveforms that reveal many signal properties, and provide a path to sparse representations. Discretized signals often have a very large size $N \geqslant 10^6$, and thus can only be processed by fast algorithms, typically implemented with $O(N \log N)$ operations and memories. Fourier and wavelet transforms illustrate the deep connection between well structured mathematical tools and fast algorithms.

1.1.1 Fourier Kingdom

The Fourier transform is everywhere in physics and mathematics, because it diagonalizes time-invariant convolution operators. It rules over linear time-invariant signal processing, whose building blocks are *frequency filtering* operators.

Fourier analysis represents any finite energy function $f(t)$ as a sum of sinusoidal waves $e^{i\omega t}$:

$$f(t) = \frac{1}{2\pi} \int_{-\infty}^{+\infty} \hat{f}(\omega)\, e^{i\omega t}\, d\omega. \tag{1.1}$$

The amplitude $\hat{f}(\omega)$ of each sinusoidal wave $e^{i\omega t}$ is equal to its correlation with f, also called Fourier transform:

$$\hat{f}(\omega) = \int_{-\infty}^{+\infty} f(t)\, e^{-i\omega t}\, dt. \tag{1.2}$$

The more regular $f(t)$ the faster the decay of the sinusoidal wave amplitude $|\hat{f}(\omega)|$ when the frequency ω increases.

When $f(t)$ is defined only on an interval, say $[0,1]$, then the Fourier transform becomes a decomposition in a Fourier orthonormal basis $\{e^{i2\pi mt}\}_{m\in\mathbb{Z}}$ of $\mathbf{L}^2[0,1]$. If $f(t)$ is uniformly regular then its Fourier transform coefficients also have a fast decay when the frequency $2\pi m$ increases, so it can be well approximated with few low-frequency Fourier coefficients. The Fourier transform therefore defines a sparse representation of uniformly regular functions.

Over discrete signals, the Fourier transform is a decomposition in a discrete orthogonal Fourier basis $\{e^{i2\pi kn/N}\}_{0\leqslant k<N}$ of $\mathbb{C}^N$, which has properties similar to a Fourier transform on functions. Its embedded structure leads to a Fast Fourier Transform algorithms, which computes discrete Fourier coefficients with $O(N\log N)$ instead of N^2. This FFT algorithm is a corner stone of discrete signal processing.

As long as we are satisfied with linear time-invariant operators or uniformly regular signals, the Fourier transform provides simple answers to most questions. Its richness makes it suitable for a wide range of applications such as signal transmissions or stationary signal processing. However, to represent a transient phenomena—a word pronounced at a particular time, an apple located in the left corner of an image—the Fourier transform becomes a cumbersome tool that requires many coefficients to represent a localized event. Indeed, the support of $e^{i\omega t}$ covers the whole real line, so $\hat{f}(\omega)$ depends on the values $f(t)$ for all times $t\in\mathbb{R}$. This global "mix" of information makes it difficult to analyze or represent any local property of $f(t)$ from $\hat{f}(\omega)$.

1.1.2 Wavelet Bases

Wavelet bases, like Fourier bases, reveal the signal regularity through the amplitude of coefficients, and their structure leads to a fast computational algorithm. However, wavelets are well localized and few coefficients are needed to represent local transient structures. As opposed to a Fourier basis, a wavelet basis defines a sparse representation of piecewise regular signals, which may include transients and singularities. In images, large wavelet coefficients are located in the neighborhood of edges and irregular textures.

The story begins in 1910, when Haar [290] constructed a piecewise constant function

$$\psi(t) = \begin{cases} 1 & \text{if } 0 \leqslant t < 1/2 \\ -1 & \text{if } 1/2 \leqslant t < 1 \\ 0 & \text{otherwise} \end{cases}$$

whose dilations and translations generate an orthonormal basis

$$\left\{ \psi_{j,n}(t) = \frac{1}{\sqrt{2^j}}\, \psi\left(\frac{t - 2^j n}{2^j} \right) \right\}_{(j,n)\in\mathbb{Z}^2}$$

of the space $\mathbf{L}^2(\mathbb{R})$ of signals having a finite energy

$$\|f\|^2 = \int_{-\infty}^{+\infty} |f(t)|^2\, dt < +\infty.$$

Let us write $\langle f, g \rangle = \int_{-\infty}^{+\infty} f(t)\, g^*(t)\, dt$ the inner product in $\mathbf{L}^2(\mathbb{R})$. Any finite energy signal f can thus represented by its wavelet inner-product coefficients

$$\langle f, \psi_{j,n} \rangle = \int_{-\infty}^{+\infty} f(t)\, \psi_{j,n}(t)\, dt$$

and recovered by summing them in this wavelet orthonormal basis:

$$f = \sum_{j=-\infty}^{+\infty} \sum_{n=-\infty}^{+\infty} \langle f, \psi_{j,n} \rangle\, \psi_{j,n}. \tag{1.3}$$

Each Haar wavelet $\psi_{j,n}(t)$ has a zero average over its support $[2^j n, 2^j (n+1)]$. If f is locally regular and 2^j is small then it is nearly constant over this interval and the wavelet coefficient $\langle f, \psi_{j,n} \rangle$ is thus nearly zero. This means that large wavelet coefficients are located at sharp signal transitions only.

With a jump in time, the story continues in 1980, when Strömberg [448] found a piecewise linear function ψ that also generates an orthonormal basis and gives better approximations of smooth functions. Meyer was not aware of this result, and motivated by the work of Morlet and Grossmann over continuous wavelet transform, he tried to prove that there exists no regular wavelet ψ that generates an orthonormal basis. This attempt was a failure since he ended up constructing a whole family of orthonormal wavelet bases, with functions ψ that are infinitely continuously differentiable [374]. This was the fundamental impulse that lead to a widespread search for new orthonormal wavelet bases, which culminated in the celebrated Daubechies wavelets of compact support [193].

The systematic theory for constructing orthonormal wavelet bases was established by Meyer and Mallat through the elaboration of multiresolution signal approximations [361] presented in Chapter 7. It was inspired by original ideas developed in computer vision by Burt and Adelson [125] to analyze images at several resolutions. Digging more into the properties of orthogonal wavelets and multiresolution approximations brought to light a surprising relation with filter banks constructed with conjugate mirror filters, and a fast wavelet transform algorithm decomposing signals of size N with $O(N)$ operations [360].

Filter Banks   Motivated by speech compression, in 1976 Croisier, Esteban and Galand [188] introduced an invertible filter bank, which decomposes a discrete signal $f[n]$ in two signals of half its size, using a filtering and subsampling procedure. They showed that $f[n]$ can be recovered from these subsampled signals by canceling the aliasing terms with a particular class of filters called *conjugate mirror filters*. This breakthrough led to a 10-year research effort to build a complete filter bank theory. Necessary and sufficient conditions for decomposing a signal in subsampled components with a filtering scheme, and recovering the same signal with an inverse transform, were established by Smith and Barnwell [443], Vaidyanathan [469] and Vetterli [471].

The multiresolution theory of Mallat [361] and Meyer [43] proves that any conjugate mirror filter characterizes a wavelet ψ that generates an orthonormal basis of $\mathbf{L}^2(\mathbb{R})$, and that a fast discrete wavelet transform is implemented by cascading these conjugate mirror filters [360]. The equivalence between this continuous time wavelet theory and discrete filter banks led to a new fruitful interface between digital signal processing and harmonic analysis, first creating an interesting culture shock that is now well resolved.

Continuous Versus Discrete and Finite   Originally, many signal processing engineers were wondering what is the point of considering wavelets and signals as functions, since all computations are performed over discrete signals, with conjugate mirror filters. Why bother with the convergence of infinite convolution cascades if in practice we only compute a finite number of convolutions? Answering these important questions is necessary in order to understand why this book alternates between theorems on continuous time functions and discrete algorithms applied to finite sequences.

A short answer would be "simplicity". In $\mathbf{L}^2(\mathbb{R})$, a wavelet basis is constructed by dilating and translating a single function ψ. Several important theorems relate the amplitude of wavelet coefficients to the local regularity of the signal f. Dilations are not defined over discrete sequences, and discrete wavelet bases are therefore more complex to describe. The regularity of a discrete sequence is not well defined either, which makes it more difficult to interpret the amplitude of wavelet coefficients. A theory of continuous time functions gives asymptotic results for discrete sequences with sampling intervals decreasing to zero. This theory is useful because these asymptotic results are precise enough to understand the behavior of discrete algorithms.

But continuous time or space models are not sufficient for elaborating discrete signal processing algorithms. The transition between continuous and discrete signals must be done with great care, to maintain important properties such as orthogonality. Restricting the constructions to finite discrete signals adds another layer of complexity because of border problems. How these border issues affect numerical implementations is carefully addressed once the properties of the bases are well understood.

Wavelets for Images Wavelet orthonormal bases of images can be constructed from wavelet or-
thonormal bases of one-dimensional signals. Three mother wavelets $\psi^1(x)$, $\psi^2(x)$ and $\psi^3(x)$, with
$x = (x_1, x_2) \in \mathbb{R}^2$, are dilated by 2^j and translated by $2^j n$ with $n = (n_1, n_2) \in \mathbb{Z}^2$. It yields an
orthonormal basis of the space $\mathbf{L^2}(\mathbb{R}^2)$ of finite energy functions $f(x) = f(x_1, x_2)$:

$$\left\{ \psi_{j,n}^k(x) = \frac{1}{2^j} \psi^k \left(\frac{x - 2^j n}{2^j} \right) \right\}_{j \in \mathbb{Z}, n \in \mathbb{Z}^2, 1 \leqslant k \leqslant 3} .$$

The support of a wavelet $\psi_{j,n}^k$ is a square of width proportional to the scale 2^j. Two dimensional
wavelet bases are discretized to define orthonormal bases of images including N pixels. Wavelet
coefficients are calculated with a fast $O(N)$ algorithm described in Chapter 7.

Like in one dimension, a wavelet coefficient $\langle f, \psi_{j,n}^k \rangle$ has a small amplitude if $f(x)$ is regular
over the support of $\psi_{j,n}^k$. It is large near sharp transitions such as edges. Figure 1.1(b) is the array
of N wavelet coefficients. Each direction k and scale 2^j corresponds to a sub-image, which shows
in black the position of the largest coefficients above a threshold: $|\langle f, \psi_{j,n}^k \rangle| \geqslant T$.

(a)   (b)

(c)   (d)
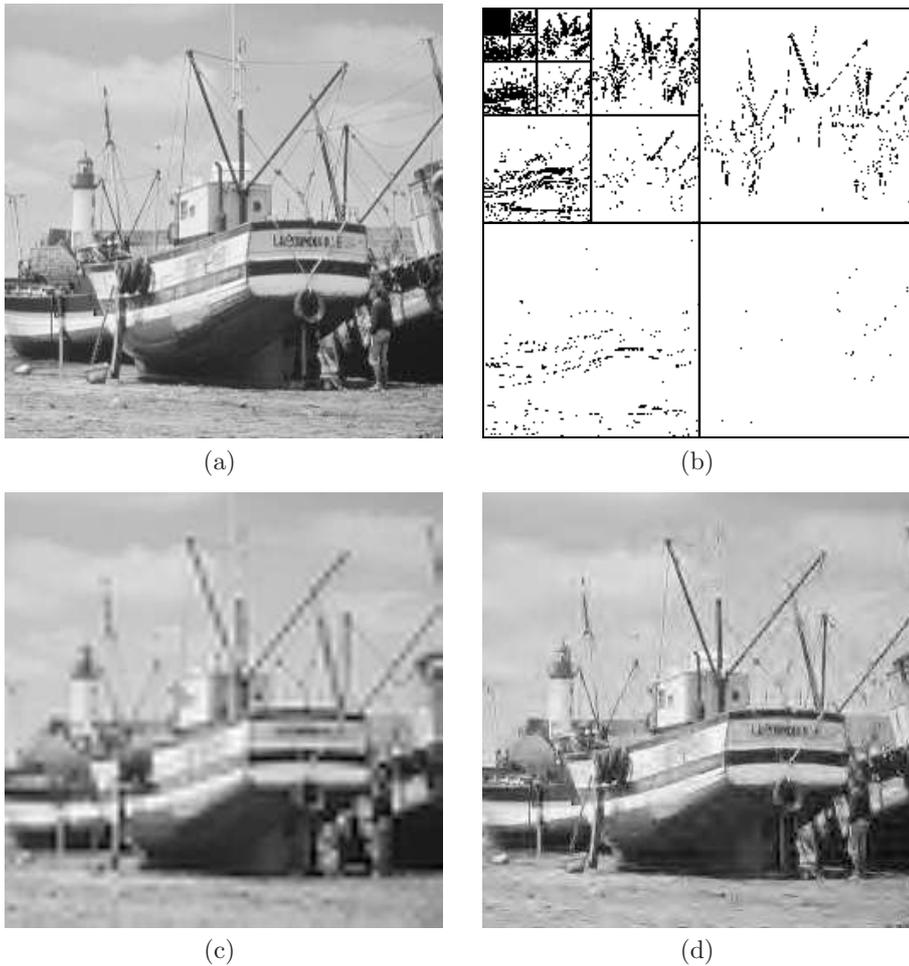
Figure 1.1: (a): Discrete image $f[n]$ of $N = 256^2$ pixels. (b): Array of N orthogonal wavelet
coefficients $\langle f, \psi_{j,n}^k \rangle$ for $k = 1, 2, 3$ and 4 scales 2^j. Black points correspond to $|\langle f, \psi_{j,n}^k \rangle| > T$. (c):
Linear approximation from the $N/16$ wavelet coefficients at the 3 largest scales. (d): Non-linear
approximation from the $M = N/16$ wavelet coefficients of largest amplitude shown in (b).

1.2 Approximation and Processing in Bases

Analog to digital signal conversion is the first step of digital signal processing. Chapter 3 explains that it amounts to projecting the signal over a basis of an approximation space. Most often, the resulting digital representation remains much too large and needs to be further reduced. A digital image includes typically more than 10^6 samples and a CD music recording has $40\,10^3$ samples per second. Sparse representations that reduce the number of parameters can be obtained by thresholding coefficients in an appropriate orthogonal basis. Efficient compression and noise reduction algorithms are then implemented with simple operators in this basis.

Stochastic versus Deterministic Signal Models A representation is optimized relatively to a signal class, corresponding to all potential signals encountered in an application. This requires building signal models that carries available prior information.

A signal f can be modeled as a realization of a random process F, whose probability distribution is known a priori. A Bayesian approach then tries to minimize the expected approximation error. Linear approximations are simpler because they only depend upon the covariance. Chapter 9 shows that optimal linear approximations are obtained in a basis of principal components that are the eigenvectors of the covariance matrix. However, the expected error of non-linear approximations depends upon the full probability distribution of F. This distribution is most often not know for complex signals such as images or sounds, because their transient structures are not well modeled as realizations of known processes such as Gaussian processes.

To optimize non-linear representations, weaker but sufficiently powerful deterministic models can be elaborated. A deterministic model specifies a set Θ where the signal belongs. This set is defined by any prior information, for example on the time-frequency localization of transients in musical recordings or on the geometric regularity of edges in images. Simple models can also define Θ as a ball in a functional space, with a specific regularity norm such as a total variation norm. A stochastic model is richer because it provides the probability distribution in Θ. When this distribution is not available, the average error can not be calculated and is replaced by the maximum error over Θ. Optimizing the representation then amounts to minimize this maximum error, which is called a *minimax* optimization.

1.2.1 Sampling with Linear Approximations

Analog to digital signal conversion is most often implemented with a linear approximation operator that filters and samples the input analog signal. From these samples, a linear digital to analog converter recovers a projection of the original analog signal over an approximation space whose dimension depends upon the sampling density. Linear approximations project signals in spaces of lowest possible dimensions to reduce computations and storage cost, while controlling the resulting error.

Sampling Theorems Let us consider finite energy signals $\|\bar{f}\|^2 = \int |\bar{f}(x)|^2\,dx$ of finite support, which is normalized to $[0,1]$ or $[0,1]^2$ for images. A sampling process implements a filtering of $\bar{f}(x)$ with a low-pass impulse response $\bar{\phi}_s(x)$, and a uniform sampling to output a discrete signal:

$$f[n] = \bar{f} \star \bar{\phi}_s(ns) \quad \text{for} \quad 0 \leqslant n < N.$$

In two dimensions, $n = (n_1, n_2)$ and $x = (x_1, x_2)$. These filtered samples can also be written as inner products:

$$\bar{f} \star \bar{\phi}_s(ns) = \int f(u)\,\bar{\phi}_s(ns - u)\,du = \langle f(x), \phi_s(x - ns)\rangle$$

with $\phi_s(x) = \bar{\phi}_s(-x)$. Chapter 3 explains that ϕ_s are chosen, like in the classical Shannon-Whittaker sampling theorem, so that family of functions $\{\phi_s(x - ns)\}_{1 \leqslant n \leqslant N}$ is a basis of an appropriate approximation space $\mathbf{U}_N$. The best linear approximation of $\bar{f}$ in $\mathbf{U}_N$ recovered from these samples is the orthogonal projection $\bar{f}_N$ of f in $\mathbf{U}_N$, and if the basis is orthonormal then

$$\bar{f}_N(x) = \sum_{n=0}^{N-1} f[n]\,\phi_s(x - ns). \tag{1.4}$$

A sampling theorem states that if $\bar{f} \in \mathbf{U}_N$ then $\bar{f} = \bar{f}_N$ so (1.4) recovers $\bar{f}(x)$ from the measured samples. Most often, $\bar{f}$ does not belong to this approximation space. It is called *aliasing* in the context of the Shannon-Whittaker sampling, where $\mathbf{U}_N$ is the space of functions having a frequency support restricted to the N lower frequencies. The approximation error $\|\bar{f} - \bar{f}_N\|^2$ must then be controled.

Linear Approximation Error   The approximation error is computed by finding an orthogonal basis $\mathcal{B} = \{\bar{g}_m(x)\}_{0 \leqslant m < +\infty}$ of the whole analog signal space $\mathbf{L}^2[0,1]^2$, whose first N vectors $\{\bar{g}_m(x)\}_{0 \leqslant m < N}$ is an orthogonal basis of $\mathbf{U}_N$. The orthogonal projection on $\mathbf{U}_N$ can thus be rewritten:

$$\bar{f}_N(x) = \sum_{m=0}^{N-1} \langle \bar{f}, \bar{g}_m \rangle \, \bar{g}_m(x) \ .$$

Since $\bar{f} = \sum_{m=0}^{+\infty} \langle \bar{f}, \bar{g}_m \rangle \, \bar{g}_m$, the approximation error is the energy of the removed inner products:

$$\varepsilon_l(N, f) = \|\bar{f} - \bar{f}_N\|^2 = \sum_{m=N}^{+\infty} |\langle \bar{f}, \bar{g}_m \rangle|^2.$$

This error decreases quickly when N increases if the coefficient amplitudes $|\langle \bar{f}, \bar{g}_m \rangle|$ have a fast decay when the index m increases. The dimension N is adjusted to the desired approximation error. Figure 1.1(a) shows a discrete image $f[n]$ approximated with $N = 256^2$ pixels. Figure 1.1(c) displays a lower resolution image $f_{N/16}$ projected on a space $\mathbf{U}_{N/16}$ of dimension $N/16$, generated by $N/16$ large scale wavelets. It is calculated by setting to zero all the wavelet coefficients at the first two smaller scales. The approximation error is $\|f - f_{N/16}\|^2 / \|f\|^2 = 14 \, 10^{-3}$. Reducing the resolution introduces more blur and errors. A linear approximation space $\mathbf{U}_N$ corresponds to a uniform grid that approximates precisely uniformly regular signals. Since images $\bar{f}$ are often not uniformly regular, it is necessary to measure it at a high resolution N. This is why digital camera have a resolution that increases as technology improves.

1.2.2   Sparse Non-linear Approximations

Linear approximations reduce the space dimensionality but can introduce important errors when reducing the resolution, if the signal is not uniformly regular, as shown by Figure 1.1(c). To improve such approximations, more coefficients should be kept where it is needed, not in regular regions but near sharp transitions and edges. This requires defining an irregular sampling adapted to the local signal regularity. This optimized irregular sampling has a simple equivalent solution through non-linear approximations in wavelet bases.

Non-linear approximations operate in two stages. First a linear approximation approximates the analog signal $\bar{f}$ with N samples that are written $f[n] = \bar{f} \star \bar{\phi}_s(ns)$. Then a non-linear approximation of $f[n]$ is computed to reduce the N coefficients $f[n]$ to $M \ll N$ coefficients in a sparse representation.

The discrete signal f can be considered as a vector of $\mathbb{C}^N$. Inner products and norms in $\mathbb{C}^N$ are written

$$\langle f, g \rangle = \sum_{n=0}^{N-1} f[n] \, g^*[n] \quad \text{and} \quad \|f\|^2 = \sum_{n=0}^{N-1} |f[n]|^2 \ .$$

To obtain a sparse representation with a non-linear approximation, we choose a new orthonormal basis $\mathcal{B} = \{g_m[n]\}_{m \in \Gamma}$ of $\mathbb{C}^N$, which concentrates as much as possible the signal energy over few coefficients. Signal coefficients $\{\langle f, g_m \rangle\}_{m \in \Gamma}$ are computed from the N input sample values $f[n]$ with an orthogonal change of basis that takes N^2 operations in non-structured bases. In a wavelet or Fourier bases, fast algorithms require respectively $O(N)$ and $O(N \log_2 N)$ operations.

Approximation by Thresholding   For $M < N$, an approximation f_M is computed by selecting the "best" $M < N$ vectors within $\mathcal{B}$. The orthogonal projection of f on the space $\mathbf{V}_\Lambda$ generated by M vectors $\{g_m\}_{m \in \Lambda}$ in $\mathcal{B}$ is

$$f_\Lambda = \sum_{m \in \Lambda} \langle f, g_m \rangle \, g_m \ . \tag{1.5}$$

Since $f = \sum_{m \in \Gamma} \langle f, g_m \rangle g_m$, the resulting error is

$$\|f - f_\Lambda\|^2 = \sum_{m \notin \Lambda} |\langle f, g_m \rangle|^2 \ . \tag{1.6}$$

We write $|\Lambda|$ the size of the set Λ. The best $M = |\Lambda|$ term approximation which minimizes $\|f - f_\Lambda\|^2$ is thus obtained by selecting the M coefficients of largest amplitude. These coefficients are above a threshold T that depends on M:

$$f_M = f_{\Lambda_T} = \sum_{m \in \Lambda_T} \langle f, g_m \rangle g_m \quad \text{with} \quad \Lambda_T = \{m \in \Gamma \ : \ |\langle f, g_m \rangle| \geqslant T\} \ . \tag{1.7}$$

This approximation is non-linear because the approximation set Λ_T changes with f. The resulting approximation error is:

$$\varepsilon_n(M, f) = \|f - f_M\|^2 = \sum_{m \notin \Lambda_T} |\langle f, g_m \rangle|^2 \ . \tag{1.8}$$

Figure 1.1(b) shows that the approximation support Λ_T of an image in a wavelet orthonormal basis, which depends upon the geometry of edges and textures. Keeping large wavelet coefficients is equivalent to constructing an adaptive approximation grid specified by the scale-space support Λ_T. It increases the approximation resolution where the signal is irregular. The geometry of Λ_T gives the spatial distribution on sharp image transitions and edges, and shows how they propagate across scales. Chapter 6 proves that it gives important information on their sharpness and local Lipschitz regularity. This example illustrates how an approximation support provides "geometric" information on f, relatively to a dictionary, which is a wavelet basis in this example.

Figure 1.1(d) gives the non-linear wavelet approximation f_M recovered from the $M = N/16$ largest amplitude wavelet coefficients, with an error $\|f - f_M\|^2 / \|f\|^2 = 5\,10^{-3}$. This error is nearly 3 times smaller than the the linear approximation error obtained with the same number of wavelet coefficients, and the image quality is much better.

An analog signal can be recovered from the discrete non-linear approximation f_M:

$$\bar{f}_M(x) = \sum_{n=0}^{N-1} f_M[n]\, \phi_s(x - ns).$$

Since all projections are orthogonal, the overall approximation error on the original analog signal $\bar{f}(x)$ is the sum of the analog sampling error and the discrete non-linear error:

$$\|\bar{f} - \bar{f}_M\|^2 = \|\bar{f} - \bar{f}_N\|^2 + \|f - f_M\|^2 = \varepsilon_l(N, f) + \varepsilon_n(M, f).$$

In practice, N is imposed by the resolution of the signal acquisition hardware, and M is typically adjusted so that $\varepsilon_n(M, f) \geqslant \varepsilon_l(N, f)$.

Sparsity with Regularity   Sparse representations are obtained in a basis that takes advantage of some form of regularity of the input signals, which creates many small amplitude coefficients. Since wavelets have a localized support, functions with isolated singularities produce few large amplitude wavelet coefficients in the neighborhood of these singularities. Non-linear wavelet approximation produce a small error over spaces of functions that do not have "too many" sharp transitions and singularities. Chapter 9 shows that functions having a bounded total variation norm are useful models for images with non-fractal (finite length) edges.

Edges often define regular geometric curves. Wavelets detect the location of edges but their square support can not take advantage of their potential geometric regularity. More sparse representations are defined in dictionaries of curvelets or bandlets, that have elongated support in multiple directions that can be adapted to this geometrical regularity. In such dictionaries, the approximation support Λ_T is smaller but provides explicit information about the local geometrical properties of edges, such as their orientation. In this context, geometry does not just apply to multidimensional signals. Audio signals such as musical recordings also have a complex geometric regularity in time-frequency dictionaries.

1.2.3 Compression

Storage limitations and fast transmission through narrow band-width channels requires to compress signals while minimizing the degradation. Transform codes compress signals by coding a sparse representation. Chapter 10 introduces the information theory needed to understand these codes and optimize their performance.

In a compression framework, the analog signal has already been discretized into a signal $f[n]$ of size N. This discrete signal is decomposed in an orthonormal basis $\mathcal{B} = \{g_m\}_{m \in \Gamma}$ of $\mathbb{C}^N$:

$$f = \sum_{m \in \Gamma} \langle f, g_m \rangle \, g_m \, .$$

Coefficients $\langle f, g_m \rangle$ are approximated by quantized values $Q(\langle f, g_m \rangle)$. If Q is a uniform quantizer of step Δ then $|x - Q(x)| \leqslant \Delta/2$, and if $|x| < \Delta/2$ then $Q(x) = 0$. The signal $\tilde{f}$ restored from quantized coefficients is

$$\tilde{f} = \sum_{m \in \Gamma} Q(\langle f, g_m \rangle) \, g_m \, .$$

An entropy code records these coefficients with R bits. The goal is to minimize the signal distortion-rate $d(R, f) = \|\tilde{f} - f\|^2$.

The coefficients not quantized to zero correspond to the set $\Lambda_T = \{m \in \Gamma \; : \; |\langle f, g_m \rangle| \geqslant T\}$ with $T = \Delta/2$. For sparse signals, Chapter 10 shows that the bit budget R is dominated by the number of bits to code Λ_T in Γ, which is nearly proportional to its size $|\Lambda_T|$. This means that the "information" of a sparse representation is mostly geometric. Moreover, the distortion is dominated by the non-linear approximation error $\|f - f_{\Lambda_T}\|^2$, for $f_{\Lambda_T} = \sum_{m \in \Lambda_T} \langle f, g_m \rangle g_m$. Compression is thus a sparse approximation problem. For a given distortion $d(R, f)$, minimizing R requires to reduce $|\Lambda_T|$ and optimize the sparsity.

The number of bits to code Λ_T can take advantage of any prior information on the geometry. Figure 1.1(b) shows that large wavelet coefficients are not randomly distributed. They have a tendency to be aggregated towards larger scales, and at fine scales they are regrouped along edge curves or in texture regions. Using such prior geometric model is an source of gain in coders such as JPEG-2000.

Chapter 10 describes the implementation of audio transform codes. Image transform codes in block cosine bases and wavelet bases are introduced, together with the JPEG and JPEG-2000 compression standards.

1.2.4 Denoising

Signal acquisition devices add noises, that can be reduced by estimators using prior information on signal properties. Signal processing has long remained mostly Bayesian and linear. Non-linear smoothing algorithms existed in statistics, but these procedures were often ad-hoc and complex. Two statisticians, Donoho and Johnstone [220], changed the game by proving that simple thresholding in sparse representations can yield nearly optimal non-linear estimators. This was the beginning of a considerable refinement of non-linear estimation algorithms that is still on-going.

Let us consider digital measurements that add a random noise $W[n]$ to the original signal $f[n]$

$$X[n] = f[n] + W[n] \;\; \text{for} \;\; 0 \leqslant n < N.$$

The signal f is estimated by transforming the noisy data X with an operator D:

$$\tilde{F} = D \, X \, .$$

The risk of the estimator $\tilde{F}$ of f is the average error, calculated with respect to the probability distribution of the noise W:

$$r(D, f) = \mathsf{E}\{\|f - DX\|^2\} \, .$$
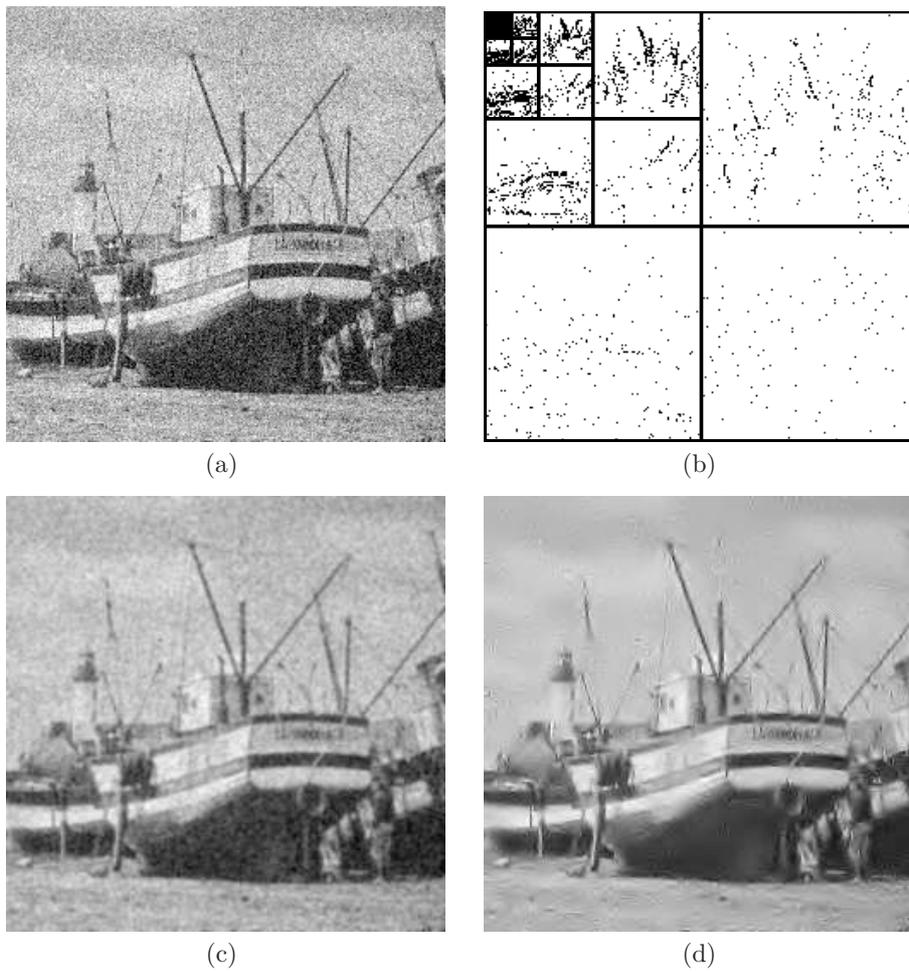
Figure 1.2: (a): Noisy image X. (b): Noisy wavelet coefficients above threshold, $|\langle X, \psi_{j,n} \rangle| \geqslant T$. (c): Linear estimation $X \star h$. (d): Non-linear estimator recovered from thresholded wavelet coefficients, over several translated bases.

Bayes Versus Minimax   To optimize the estimation operator D, one must take advantage of prior information available about the signal f. In a Bayes framework, f is considered as a realization of a random vector F and the Bayes risk is the expected risk calculated with respect to the prior probability distribution π of the random signal model F:

$$r(D, \pi) = \mathsf{E}_\pi\{r(D, F)\} .$$

Optimizing D among all possible operators yields the *minimum Bayes risk*:

$$r_n(\pi) = \inf_{all\ D} r(D, \pi) .$$

In the 1940's, Wald brought a new perspective on statistics, through a decision theory partly imported from the theory of games. This point of view uses deterministic models, where signals are elements of a set Θ, without specifying their probability distribution in this set. To control the risk for any $f \in \Theta$, we compute the maximum risk

$$r(D, \Theta) = \sup_{f \in \Theta} r(D, f) .$$

The *minimax risk* is the lower bound computed over all operators D:

$$r_n(\Theta) = \inf_{all\ D} r(D, \Theta).$$

In practice, the goal is to find an operator D that is simple to implement and which yields a risk close the minimax lower bound.

Thresholding estimators   It is tempting to restrict calculations to linear operators D, because of their simplicity. Optimal linear Wiener estimators are introduced in Chapter 11. Figure 1.2(a) is an image contaminated by a Gaussian white noise. Figure 1.2(b) shows an optimized linear filtering estimation $\tilde{F} = X \star h[n]$, which is therefore diagonal in a Fourier basis $\mathcal{B}$. This convolution operator averages the noise but also blurs the image and keeps low-frequency noise by keeping the image low-frequencies.

If f has a sparse representation in a dictionary, then projecting X on the vectors of this sparse support can considerably improve linear estimators. The difficulty is to identify the sparse support of f from the noisy data X. Donoho and Johnstone [220] proved that in an orthonormal basis, a simple thresholding of noisy coefficients does the trick. Noisy signal coefficients in an orthonormal basis $\mathcal{B} = \{g_m\}_{m \in \Gamma}$ are

$$\langle X, g_m \rangle = \langle f, g_m \rangle + \langle W, g_m \rangle \ \ \text{for} \ \ m \in \Gamma .$$

Thresholding these noisy coefficients yields an orthogonal projection estimator

$$\tilde{F} = X_{\tilde{\Lambda}_T} = \sum_{m \in \tilde{\Lambda}_T} \langle X, g_m \rangle g_m \quad \text{with} \quad \tilde{\Lambda}_T = \{m \in \Gamma \ : \ |\langle X, g_m \rangle| \geqslant T\} . \tag{1.9}$$

The set $\tilde{\Lambda}_T$ is an estimate of an approximation support of f. It is hopefully close to the optimal approximation support $\Lambda_T = \{m \in \Gamma \ : \ |\langle f, g_m \rangle| \geqslant T\}$. Figure 1.2(b) shows the estimated approximation set $\tilde{\Lambda}_T$ of noisy wavelet coefficients $|\langle X, \psi_{j,n} \rangle| \geqslant T$, that can be compared to the optimal approximation support Λ_T shown in Figure 1.1(b). The estimation in Figure 1.2(d) from wavelet coefficients in $\tilde{\Lambda}_T$ has considerably reduced the noise in regular regions while keeping the sharpness of edges by preserving large wavelet coefficients. This estimation is improved with a translation invariant procedure that averages this estimator over several translated wavelet bases. Thresholding wavelet coefficients implements an adaptive smoothing, which averages the data X with a kernel that depends on the estimated regularity of the original signal f.

Donoho and Johnstone proved that for a Gaussian white noise of variance σ^2, choosing $T = \sigma\sqrt{2\log_e N}$ yields a risk $\mathsf{E}\{\|f - \tilde{F}\|^2\}$ of the order of $\|f - f_{\Lambda_T}\|^2$, up to a $\log_e N$ factor. This spectacular result shows that the estimated support $\tilde{\Lambda}_T$ does nearly as well as the optimal unknown support Λ_T. The resulting risk is small if the representation is sparse and precise.

The set $\tilde{\Lambda}_T$ in Figure 1.2(b) "looks" different from the Λ_T in Figure 1.1(b) because it has more isolated points. This indicates that some prior information on the geometry of Λ_T could be used to improve the estimation. For audio noise reduction, thresholding estimators are applied in sparse representations provided by time-frequency bases. Similar isolated time-frequency coefficients produce a highly annoying "musical noise". Musical noise is removed with a block thresholding that regularizes the geometry of the estimated support $\tilde{\Lambda}_T$, and avoids leaving isolated points. Block thresholding also improves wavelet estimators.

If W is a Gaussian noise and signals in Θ have a sparse representation in $\mathcal{B}$, then Chapter 11 proves that thresholding estimators can produce a nearly minimax risk. In particular, wavelet thresholding estimators have a nearly minimax risk for large classes of piecewise smooth signals, including bounded variation images.

1.3   Time-Frequency Dictionaries

Motivated by quantum mechanics, in 1946 the physicist Gabor [266] proposed to decompose signals over dictionaries of elementary waveforms, that he called time-frequency atoms, which have a minimal spread in a time-frequency plane. By showing that such decompositions are closely related to our perception of sounds, and that they exhibit important structures in speech and music recordings, Gabor demonstrated the importance of localized time-frequency signal processing. Beyond sounds, large classes of signals have sparse decompositions as sums of time-frequency atoms selected from appropriate dictionaries. The key issue is to understand how to construct dictionaries with time-frequency atoms adapted to signal properties.
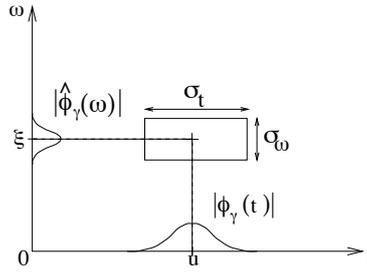
Figure 1.3: Heisenberg box representing an atom ϕ_γ.

1.3.1   Heisenberg Uncertainty

A time-frequency dictionary $\mathcal{D} = \{\phi_\gamma\}_{\gamma \in \Gamma}$ is composed of waveforms of unit norm $\|\phi_\gamma\| = 1$, which are well localized in time and in frequency. The time localization u of ϕ_γ and its spread around u are defined by

$$u = \int t\,|\phi_\gamma(t)|^2\,dt \quad \text{and} \quad \sigma_{t,\gamma}^2 = \int |t-u|^2\,|\phi_\gamma(t)|^2\,dt \;.$$

Similarly, the frequency localization and spread of $\hat{\phi}_\gamma$ are defined by

$$\xi = (2\pi)^{-1} \int \omega\,|\hat{\phi}_\gamma(\omega)|^2\,d\omega \quad \text{and} \quad \sigma_{\omega,\gamma}^2 = (2\pi)^{-1} \int |\omega-\xi|^2\,|\hat{\phi}_\gamma(\omega)|^2\,d\omega \;.$$

The Fourier Parseval formula

$$\langle f, \phi_\gamma \rangle = \int_{-\infty}^{+\infty} f(t)\,\phi_\gamma^*(t)\,dt = \frac{1}{2\pi} \int_{-\infty}^{+\infty} \hat{f}(\omega)\,\hat{\phi}_\gamma^*(\omega)\,d\omega \;. \tag{1.10}$$

shows that $\langle f, \phi_\gamma \rangle$ depends mostly on the values $f(t)$ and $\hat{f}(\omega)$ where $\phi_\gamma(t)$ and $\hat{\phi}_\gamma(\omega)$ are nonnegligible, and hence for (t,ω) in a rectangle centered at (u,ξ), of size $\sigma_{t,\gamma} \times \sigma_{\omega,\gamma}$. This rectangle is illustrated by Figure 1.3 in this time-frequency plane (t,ω). It can be interpreted as a "quantum

of information" over an elementary resolution cell. The uncertainty principle theorem proves in Chapter 2 that this rectangle has a minimum surface which limits the joint time-frequency resolution:

$$\sigma_{t,\gamma}\,\sigma_{\omega,\gamma} \geqslant \frac{1}{2}\ . \tag{1.11}$$

Constructing a dictionary of time-frequency atoms can thus be thought as covering the time-frequency plan with resolution cells whose time width $\sigma_{t,\gamma}$ and frequency width $\sigma_{\omega,\gamma}$ may vary, but with a surface larger than $1/2$. Windowed Fourier and wavelet transforms are two important examples.

1.3.2   Windowed Fourier Transform

A windowed Fourier dictionary is constructed by translating in time and frequency a time window $g(t)$, of unit norm $\|g\| = 1$, centered at $t = 0$:

$$\mathcal{D} = \left\{ g_{u,\xi}(t) = g(t-u)\,\mathrm{e}^{i\xi t} \right\}_{(u,\xi)\in\mathbb{R}^2}\ .$$

The atom $g_{u,\xi}$ is translated by u in time and by ξ in frequency. The time and frequency spread of $g_{u,\xi}$ is independent of u and ξ. This means that each atom $g_{u,\xi}$ corresponds to a Heisenberg rectangle whose size $\sigma_t \times \sigma_\omega$ is independent from its position (u,ξ), as shown by Figure 1.4.
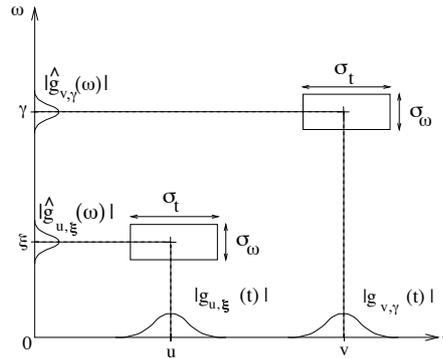
Figure 1.4: Time-frequency boxes ("Heisenberg rectangles") representing the energy spread of two windowed Fourier atoms.

The windowed Fourier transform projects f on each dictionary atom $g_{u,\xi}$:

$$Sf(u,\xi) = \langle f, g_{u,\xi} \rangle = \int_{-\infty}^{+\infty} f(t)\,g(t-u)\,\mathrm{e}^{-i\xi t}\,dt. \tag{1.12}$$

It can be interpreted as a Fourier transform of f at the frequency ξ, localized by the window $g(t-u)$ in the neighborhood of u. This windowed Fourier transform is highly redundant and represents one-dimensional signals by a time-frequency image in (u,ξ). It is thus necessary to understand how to select much fewer time-frequency coefficients that represent the signal efficiently.

When listening to music, we perceive sounds that have a frequency that varies in time. Chapter 4 shows that a spectral line of f creates high amplitude windowed Fourier coefficients $Sf(u,\xi)$ at frequencies $\xi(u)$ that depend on the time u. These spectral components are detected and characterized by ridge points, which are local maxima in this time-frequency plane. Ridge points define a time-frequency approximation support Λ of f, whose geometry depends on the time-frequency evolution of the signal spectral components. Modifying the sound duration or audio transpositions are implemented by modifying the geometry of the ridge support in time-frequency.

A windowed Fourier transform decomposes signals over waveforms that have the same time and frequency resolution. It is thus effective as long as the signal does not include structures having different time-frequency resolution, some being very localized in time and others very localized in frequency. Wavelets address this issue by changing the time and frequency resolution.

1.3.3 Continuous Wavelet Transform

In reflection seismology, Morlet knew that the waveforms sent underground have a duration that is too long at high frequencies to separate the returns of fine, closely-spaced geophysical layers. These waveforms are called *wavelets* in geophysics. Instead of emitting pulses of equal duration, he thus thought of sending shorter waveforms at high frequencies. Such waveforms could simply be obtained by scaling the mother wavelet, hence the name of this transform. Although Grossmann was working in theoretical physics, he recognized in Morlet's approach some ideas that were close to his own work on coherent quantum states. Nearly forty years after Gabor, Morlet and Grossmann reactivated a fundamental collaboration between theoretical physics and signal processing, which led to the formalization of the continuous wavelet transform [287]. Yet, these ideas were not totally new to mathematicians working in harmonic analysis, or to computer vision researchers studying multiscale image processing. It was thus only the beginning of a rapid catalysis that brought together scientists with very different backgrounds.

A wavelet dictionary is constructed from a mother wavelet ψ of zero average

$$\int_{-\infty}^{+\infty} \psi(t)\, dt = 0,$$

which is dilated with a scale parameter s, and translated by u:

$$\mathcal{D} = \left\{ \psi_{u,s}(t) = \frac{1}{\sqrt{s}}\, \psi\left(\frac{t-u}{s}\right) \right\}_{u \in \mathbb{R}, s > 0}. \tag{1.13}$$

The continuous wavelet transform of f at any scale s and position u is the projection of f on the corresponding wavelet atom:

$$Wf(u,s) = \langle f, \psi_{u,s} \rangle = \int_{-\infty}^{+\infty} f(t)\, \frac{1}{\sqrt{s}}\, \psi^*\left(\frac{t-u}{s}\right)\, dt. \tag{1.14}$$

It represents one-dimensional signals by highly redundant time-scale images in (u, s).
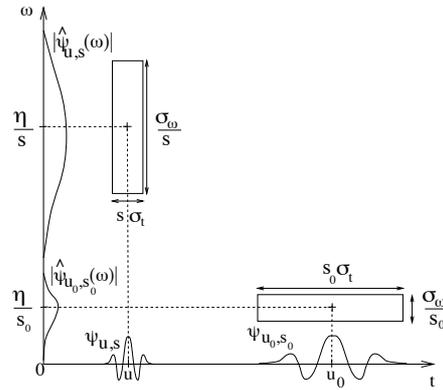
Figure 1.5: Heisenberg time-frequency boxes of two wavelets $\psi_{u,s}$ and ψ_{u_0,s_0}. When the scale s decreases, the time support is reduced but the frequency spread increases and covers an interval that is shifted towards high frequencies.

Varying Time-Frequency Resolution As opposed to windowed Fourier atoms, wavelets have a time-frequency resolution that change. The wavelet $\psi_{u,s}$ has a time support centered at u whose size is proportional to s. Let us choose a wavelet ψ whose Fourier transform $\hat{\psi}(\omega)$ is non-zero in a positive frequency interval centered at η. The Fourier transform $\hat{\psi}_{u,s}(\omega)$ is dilated by $1/s$ and is hence localized in a positive frequency interval centered at $\xi = \eta/s$, whose size is scaled by $1/s$. In the time-frequency plane, the Heisenberg box of a wavelet atom $\psi_{u,s}$ is therefore a rectangle centered at $(u, \eta/s)$, with time and frequency widths respectively proportional to s and $1/s$. When

s varies, the time and frequency width of this time-frequency resolution cell changes but its area remains constant, as illustrated by Figure 1.5.

Large amplitude wavelet coefficients can detect and measure short high frequency variations because they are well localized in time at high frequencies. At low frequencies their time resolution is lower but they have a better frequency resolution. This modification of time and frequency resolution is well adapted to represent sounds with sharp attacks, or radar signals whose frequency may vary quickly at high frequencies.

Multiscale Zooming   A wavelet dictionary is also well adapted to analyze the scaling evolution of transients with zooming procedures across scales. Suppose now that ψ is real. Since it has a zero average, a wavelet coefficient $Wf(u, s)$ measures the variation of f in a neighborhood of u whose size is proportional to s. Sharp signal transitions create large amplitude wavelet coefficients.

Signal singularities have specific scaling invariance characterized by Lipschitz exponents. Chapter 6 relates the pointwise regularity of f to the asymptotic decay of the wavelet transform amplitude $|Wf(u, s)|$, when s goes to zero. Singularities are detected by following across scales the local maxima of the wavelet transform.

In images, wavelet local maxima indicate the position of edges, which are sharp variations of the image intensity. It defines a scale-space approximation support of f from which precise image approximations are reconstructed. At different scales, the geometry of this local maxima support provides contours of image structures of varying sizes. This multiscale edge detection is particularly effective for pattern recognition in computer vision [145].

The zooming capability of the wavelet transform not only locates isolated singular events, but can also characterize more complex multifractal signals having non-isolated singularities. Mandelbrot [40] was the first to recognize the existence of multifractals in most corners of nature. Scaling one part of a multifractal produces a signal that is statistically similar to the whole. This self-similarity appears in the continuous wavelet transform, which modifies the analyzing scale. From global measurements of the wavelet transform decay, Chapter 6 measures the singularity distribution of multifractals. This is particularly important in analyzing their properties and testing multifractal models in physics or in financial time series.

1.3.4   Time-Frequency Orthonormal Bases

Orthonormal bases of time-frequency atoms remove all redundancy and define stable representations. A wavelet orthonormal basis is an example of time-frequency basis obtained by scaling a wavelet ψ with dyadic scales $s = 2^j$ and translating it by $2^j n$, which is written $\psi_{j,n}$. In the time-frequency plane, the Heisenberg resolution box of $\psi_{j,n}$ is a dilation by 2^j and translation by $2^j n$ of the Heisenberg box of ψ. A wavelet orthonormal is thus a subdictionary of the continuous wavelet transform dictionary, which yields a perfect tiling of the time-frequency plane illustrated in Figure 1.6.
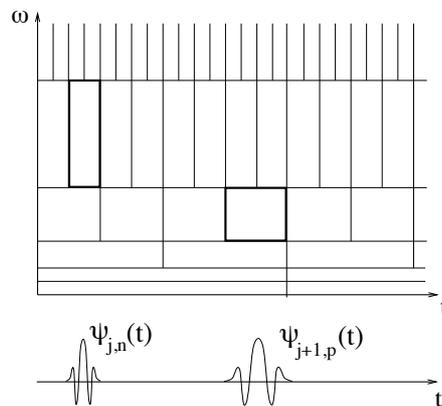
Figure 1.6: The time-frequency boxes of a wavelet basis define a tiling of the time-frequency plane.

One can construct many other orthonormal bases of time-frequency atoms, corresponding to different tilings of the time-frequency plane. Wavelet packet and local cosine bases are two important examples constructed in Chapter 8, with time-frequency atoms that split respectively the frequency and the time axis in intervals of varying sizes.
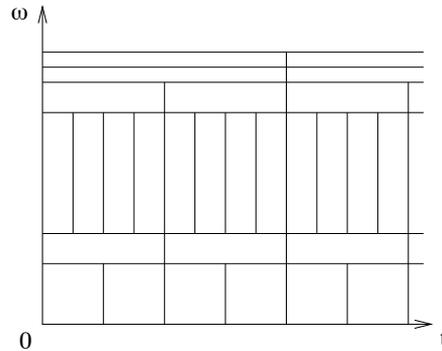
Figure 1.7: A wavelet packet basis divides the frequency axis in separate intervals of varying sizes. A tiling is obtained by translating in time the wavelet packets covering each frequency interval.

Wavelet Packet Bases Wavelet bases divide the frequency axis into intervals of 1 octave bandwidth. Coifman, Meyer and Wickerhauser [181] have generalized this construction with bases that split the frequency axis in intervals whose bandwidths may be adjusted. Each frequency interval is covered by the Heisenberg time-frequency boxes of wavelet packet functions translated in time, in order to cover the whole plane, as shown by Figure 1.7.

Like for wavelets, wavelet packet coefficients are obtained with a filter bank of conjugate mirror filters, that splits the frequency axis in several frequency intervals. Different frequency segmentations correspond to different wavelet packet bases. For images, a filter bank divides the image frequency support in squares of dyadic sizes, that can be adjusted.

Local Cosine Bases Local cosine orthonormal bases are constructed by dividing the time axis instead of the frequency axis. The time axis is segmented in successive intervals $[a_p, a_{p+1}]$. The local cosine bases of Malvar [367] are obtained by designing smooth windows $g_p(t)$ that cover each interval $[a_p, a_{p+1}]$, and multiplying them by cosine functions $\cos(\xi t + \phi)$ of different frequencies. This is yet another idea that was independently studied in physics, signal processing and mathematics. Malvar's original construction was for discrete signals. At the same time, the physicist Wilson [486] was designing a local cosine basis with smooth windows of infinite support, to analyze the properties of quantum coherent states. Malvar bases were also rediscovered and generalized by the harmonic analysts Coifman and Meyer [180]. These different views of the same bases brought to light mathematical and algorithmic properties that opened new applications.

A multiplication by $\cos(\xi t + \phi)$ translates the Fourier transform $\hat{g}_p(\omega)$ of $g_p(t)$ by $\pm \xi$. Over positive frequencies, the time-frequency box of the modulated window $g_p(t) \cos(\xi t + \phi)$ is therefore equal to the time-frequency box of g_p translated by ξ along frequencies. Figure 1.8 shows the time-frequency tiling corresponding to such a local cosine basis. For images, a two-dimensional cosine basis is constructed by dividing the image support in squares of varying sizes.

1.4 Sparsity in Redundant Dictionaries

In natural languages, large dictionaries are needed to refine ideas with short sentences, and they evolve with their usage. Eskimos have 8 different words to describe snow quality where as a single word is used in a Parisian dictionary. Similarly, large signal dictionaries of vectors are needed to construct sparse representations of complex signals. However, computing and optimizing a signal approximation by choosing the best M dictionary vectors is much more difficult.
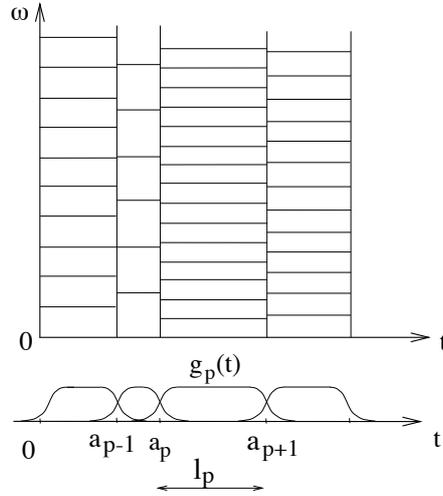
Figure 1.8: A local cosine basis divides the time axis with smooth windows $g_p(t)$, and translates these windows in frequency.

1.4.1   Frame Analysis and Synthesis

Suppose that a sparse family of vectors $\{\phi_p\}_{p\in\Lambda}$ has been selected to approximate a signal f. An approximation can be recovered as an orthogonal projection in the space $\mathbf{V}_\Lambda$ generated by these vectors. We are then facing one of the following two problems.

In a *dual synthesis* problem, the orthogonal projection f_Λ of f in $\mathbf{V}_\Lambda$ must be computed from dictionary coefficients $\{\langle f, \phi_p\rangle\}_{p\in\Lambda}$, provided by an analysis operator. This is the case when a signal transform $\{\langle f, \phi_p\rangle\}_{p\in\Gamma}$ is calculated in some large dictionary, and a subset of inner products are selected. Such inner products may correspond to coefficients above a threshold or local maxima values.

In a *dual analysis* problem, the decomposition coefficients of f_Λ must be computed on a family of selected vectors $\{\phi_p\}_{p\in\Lambda}$. This problem appears when sparse representation algorithms select vectors as opposed to inner products. This is the case for pursuit algorithms, which compute approximation supports in highly redundant dictionaries.

The frame theory gives energy equivalence conditions to solve both problems with stable operators. A family $\{\phi_p\}_{p\in\Lambda}$ is a frame of the space $\mathbf{V}$ it generates if there exists $B \geqslant A > 0$ such that

$$\forall h \in \mathbf{V} \quad, \quad A\|h\|^2 \leqslant \sum_{m\in\Lambda} |\langle h, \phi_p\rangle|^2 \leqslant B\|h\|^2 \ .$$

The representation is stable since any perturbation of frame coefficients implies a modification of similar magnitude on h. Chapter 5 derives the existence of a dual frame $\{\tilde{\phi}_p\}_{p\in\Lambda}$ that solves both the dual synthesis and analysis problems:

$$f_\Lambda = \sum_{p\in\Lambda} \langle f, \phi_p\rangle \, \tilde{\phi}_p = \sum_{p\in\Lambda} \langle f, \tilde{\phi}_p\rangle \, \phi_p \ . \tag{1.15}$$

Algorithms are provided to calculate these decompositions. The dual frame is also stable:

$$\forall f \in \mathbf{V} \quad, \quad B^{-1}\|f\|^2 \leqslant \sum_{m\in\Gamma} |\langle f, \tilde{\phi}_p\rangle|^2 \leqslant B^{-1}\|f\|^2 \ .$$

The frame bounds A and B are redundancy factors. If the vectors $\{\phi_p\}_{p\in\Gamma}$ are normalized and linearly independent then $A \leqslant 1 \leqslant B$. Such a dictionary is called a *Riesz basis* of $\mathbf{V}$ and the dual frame is biorthogonal:

$$\forall (p, p') \in \Lambda^2 \quad, \quad \langle \phi_p, \tilde{\phi}_{p'}\rangle = \delta[p - p'] \ .$$

Whent the basis is orthonormal then both bases are equal. Analysis and synthesis problems are then identical.

The frame theory is also be used to construct redundant dictionaries that define complete, stable and redundant signal representations, where $\mathbf{V}$ is then the whole signal space. The frame bounds measure the redundancy of such dictionaries. Chapter 5 studies the construction of windowed Fourier and wavelet frame dictionaries, by sampling their time, frequency and scaling parameters, while controlling frame bounds. In two dimensions, directional wavelet frames include wavelets sensitive to directional structures such as textures or edges in images.

To improve the sparsity of images having edges along regular geometric curves, Candès and Donoho [133] have introduced curvelet frames, with elongated waveforms having different directions, positions and scales. Images with piecewise regular edges have representations that are asymptotically more sparse by thresholding curvelet coefficients than wavelet coefficients.

1.4.2 Ideal Dictionary Approximations

In a redundant dictionary $\mathcal{D} = \{\phi_p\}_{p \in \Gamma}$ we would like to find the best approximation support Λ with $M = |\Lambda|$ vectors, which minimize the error $\|f - f_\Lambda\|^2$. Chapter 12 proves that it is equivalent find Λ_T which minimizes the corresponding approximation Lagrangian

$$\mathcal{L}_0(T, f, \Lambda) = \|f - f_\Lambda\|^2 + T^2 |\Lambda| , \qquad (1.16)$$

for some multiplier T.

Compression and denoising are two applications of redundant dictionary approximations. When compressing signals by quantizing dictionary coefficients, the distortion-rate varies like the Lagrangian (1.16), with a multiplier T that depends on the quantization step. Optimizing the coder is thus equivalent to minimizing this approximation Lagrangian. For sparse representations, most of the bits are devoted to coding the geometry of the sparse approximation set Λ_T in Γ.

Estimators reducing noise from observations $X = f + W$ are also optimized by finding a best orthogonal projector over a set of dictionary vectors. The *model selection* theory of Barron, Birgé and Massart [95] proves that finding $\tilde{\Lambda}_T$ which minimizes this same Lagrangian $\mathcal{L}_0(T, X, \Lambda)$ defines an estimator whose risk is of the same order as the minimum approximation error $\|f - f_{\Lambda_T}\|^2$, up to a logarithmic factor. This is similar to the optimality result obtained for thresholding estimators in an orthonormal basis.

The bad news is that minimizing the approximation Lagrangian $\mathcal{L}_0$ is an NP-hard problem, and is therefore computationally intractable. It is thus necessary to find algorithms that are sufficiently fast, and which computes sub-optimal but "good enough" solutions.

Dictionaries of Orthonormal Bases To reduce the complexity of optimal approximations, the search can be reduced to sub-families of orthogonal dictionary vectors. In a dictionary of orthonormal bases, any family of orthogonal dictionary vectors can be complemented to form an orthogonal basis $\mathcal{B}$ included in $\mathcal{D}$. As a result the best approximation of f from orthogonal vectors in $\mathcal{B}$ is obtained by thresholding the coefficients of f in a "best basis" in $\mathcal{D}$.

For tree-dictionaries of orthonormal bases obtained by a recursive split of orthogonal vector spaces, the fast dynamic programming algorithm of Coifman and Wickerhauser [181] finds such a best basis with $O(P)$ operations, where P is the dictionary size.

Wavelet packet and local cosine bases are examples of tree-dictionaries of time-frequency orthonormal bases of size $P = N \log_2 N$. A best basis is a time-frequency tiling that best match the signal time-frequency structures.

To approximate edges that are geometrically regular, wavelet are not as efficient as curvelets, but wavelets provide more sparse representations of singularities that are not distributed along geometrically regular curves. Bandlet dictionaries, introduced by Le Pennec, Mallat and Peyré [341, 364], are dictionaries of orthonormal bases, that can adapt to the variability of the image geometric regularity. Minimax optimal asymptotic rates are derived for compression and denoising.

1.4.3 Pursuit in Dictionaries

Approximating signals only from orthogonal vectors brings rigidity which limits the ability to optimize the representation. Pursuit algorithms remove this constraint with flexible procedures that

search for sparse though not necessarily optimal dictionary approximations. These approximations are computed by optimizing the choice of dictionary vectors $\{\phi_p\}_{p\in\Lambda}$.

Matching Pursuit   Matching pursuit algorithms introduced by Mallat and Zhang [365] are greedy algorithms that optimize approximations by selecting one by one the dictionary vectors. The vector in $\phi_{p_0} \in \mathcal{D}$ that best approximates a signal f is

$$\phi_{p_0} = \underset{p\in\Gamma}{\operatorname{argmax}} |\langle f, \phi_p \rangle|$$

and the residual approximation error is

$$Rf = f - \langle f, \phi_{p_0} \rangle \phi_{p_0} \ .$$

A matching pursuit further approximates the residue Rf by selecting another best vector ϕ_{p_1} from the dictionary, and continues this process over the next order residues $R^m f$, which produces a signal decomposition:

$$f = \sum_{m=0}^{M-1} \langle R^m f, \phi_{p_m} \rangle \phi_{p_m} + R^M f \ .$$

The approximation from the M selected vectors $\{\phi_{p_m}\}_{0\leqslant m<M}$ can be refined with an orthogonal back projection on the space generated by these vectors. An orthogonal matching pursuit further improves this decomposition by orthogonalizing progressively the projection directions ϕ_{p_m}, during the decompositon. The resulting decompositions are applied to compression, denoising and pattern recognition of various types of signals, images and videos.

Basis Pursuit   Approximating f with a minimum number of non-zero coefficients $a[p]$ in a dictionary $\mathcal{D}$ is equivalent to minimize the $\mathbf{l^0}$ norm $\|a\|_0$, which gives the number of non-zero coefficients. This $\mathbf{l^0}$ norm is highly non convex which explains why the resulting minimization is NP-hard. Donoho and Chen [157] thus proposed to replace the $\mathbf{l^0}$ norm by the $\mathbf{l^1}$ norm $\|a\|_1 = \sum_{p\in\Gamma} |a[p]|$ which is convex. The resulting basis pursuit algorithm computes a synthesis operator

$$f = \sum_{p\in\Gamma} a[p]\,\phi_p \quad \text{which minimizes} \quad \|a\|_1 = \sum_{p\in\Gamma} |a[p]| \ . \tag{1.17}$$

This optimal solution is calculated with a linear programming algorithm. A basis pursuit is computationally more intensive than a matching pursuit, but is a more global optimization that yields representations that can be more sparse.

In approximation, compression or denoising applications, f is recovered with an error which bounded by a precision parameter ε. The optimization (1.18) is thus relaxed by finding a synthesis such that

$$\|f - \sum_{p\in\Gamma} a[p]\,\phi_p\| \leqslant \varepsilon \quad \text{which minimizes} \quad \|a\|_1 = \sum_{p\in\Gamma} |a[p]| \ . \tag{1.18}$$

This is a convex minimization problem whose solution can be found by minimizing the corresponding $\mathbf{l^1}$ Lagrangian

$$\mathcal{L}_1(T, f, a) = \|f - \sum_{p\in\Gamma} a[p]\,\phi_p\|^2 + T\,\|a\|_1 \ ,$$

where T is a Lagrange multiplier that depends on ε. This is called an $\mathbf{l^1}$ Lagrangian pursuit in this book. A solution $\tilde{a}[p]$ is computed with iterative algorithms that are guaranteed to converge. The number of non-zero coordinates of $\tilde{a}$ typically decreases as T increases.

Incoherence for Support Recovery   Matching pursuit and $\mathbf{l^1}$ Lagrangian pursuits are optimal if they recover the approximation support Λ_T which minimizes the approximation Lagrangian

$$\mathcal{L}_0(T, f, \Lambda) = \|f - f_\Lambda\|^2 + T^2\,|\Lambda| \ ,$$

where f_Λ is the orthogonal projection of f in the space $\mathbf{V}_\Lambda$ generated by $\{\phi_p\}_{p\in\Lambda}$. This is not always true and depends on Λ_T. An *Exact Recovery Criteria* proved by Tropp [464] guarantees that pursuit algorithms do recover the optimal support Λ_T if

$$ERC(\Lambda_T) = \max_{q\notin\Lambda_T} \sum_{p\in\Lambda_T} |\langle \tilde{\phi}_p, \phi_q \rangle| < 1 \;, \tag{1.19}$$

where $\{\tilde{\phi}_p\}_{p\in\Lambda_T}$ is the biorthogonal basis of $\{\phi_p\}_{p\in\Lambda_T}$ in $\mathbf{V}_{\Lambda_T}$. This criteria implies that dictionary vectors ϕ_q outside Λ_T should have a small inner-product with vectors in Λ_T.

This recovery is stable relatively to noise perturbations if $\{\phi_p\}_{p\in\Lambda}$ has Riesz bounds that are not too far from 1. These vectors should be nearly orthogonal and hence have small inner-products. These small inner-product conditions are interpreted as a form of incoherence. A stable recovery of Λ_T is possible if vectors in Λ_T are incoherent with respect to other dictionary vectors and are incoherent between themselves. This opens the possibility of performing super-resolution with sparse representations, but it depends upon the geometric configuration of Λ_T in Γ.

1.5 Inverse Problems

Most digital measurement devices, such as cameras, microphones or medical imaging systems, can be modeled as a linear transformation of an incoming analog signal, plus a noise due to intrinsic measurement fluctuations or to electronic noises. This linear transformation can be decomposed into a stable analog to digital linear conversion followed by a discrete operator U that carries the specific transfer function of the measurement device. The resulting measured data can be written

$$Y[q] = U f[q] + W[q]$$

where $f \in \mathbb{C}^N$ is the high resolution signal we want to recover, and $W[q]$ is the measurement noise. For a camera whose optics is out of focus, the operator U is a low-pass convolution producing a blur. For a magnetic resonance imaging system, U is a Radon transform integrating the signal along rays and the number Q of measurements is smaller than N. In such problems, U is not invertible and recovering an estimate of f is an *ill-posed* inverse problem.

Inverse problems are among the most difficult signal processing problems, with considerable applications. When data acquisition is difficult, costly or dangerous or when the signal is degraded, super-resolution is important to recover the highest possible resolution information. This applies to satellite observations, seismic exploration, medical imaging, radars, camera phones or degraded Internet videos displayed on high resolution screens. Separating mixed information sources from fewer measurements is yet another super-resolution problem in telecommunication or audio recognition.

Incoherence, sparsity and geometry play a crucial role in the solution of these ill-defined inverse problems. With a sensing matrix U with random coefficients, Candès and Tao [138] and Donoho [216] proved that super-resolution becomes stable for signals having a sufficiently sparse representation in a dictionary. This remarkable result opens the door to new compressive sensing devices and algorithms, that recover high resolution signals from few randomized linear measurements.

1.5.1 Diagonal Inverse Estimation

In an ill-posed inverse problem

$$Y = U f + W$$

the image space $\mathbf{ImU} = \{Uh \;:\; h \in \mathbb{C}^N\}$ of U is of dimension Q smaller than the high resolution space N where f belongs. Inverse problems include two difficulties. In the image space $\mathbf{ImU}$ where U is invertible, its inverse may amplify the noise W which then needs to be reduced by an efficient denoising procedure. In the null space $\mathbf{NullU}$, all signals h are set to zero $Uh = 0$ and thus disappear in the measured data Y. Recovering the projection of f in $\mathbf{NullU}$ thus requires using some strong prior information. A super-resolution estimator recovers an estimation of f in a space of dimension larger than Q and hopefully equal to N, but this is not always possible.

Singular Value Decompositions   Let $f = \sum_{m \in \Gamma} a[m] \, g_m$ be the representation of f in an orthonormal basis $\mathcal{B} = \{g_m\}_{m \in \Gamma}$. An approximation must be recovered from

$$Y = \sum_{m \in \Gamma} a[m] \, U g_m + W \ .$$

If $\mathcal{B}$ is chosen to diagonalize U^*U then U transforms a subset of Q vectors $\{g_m\}_{m \in \Gamma_Q}$ of $\mathcal{B}$ into an orthogonal basis $\{U g_m\}_{m \in \Gamma_Q}$ of $\mathbf{Im U}$, and sets to zero all other vectors. A singular value decomposition estimates the coefficients $a[m]$ of f by projecting Y on this orthogonal basis and by renormalizing the resulting coefficients

$$\forall m \in \Gamma \quad , \quad \tilde{a}[m] = \frac{\langle Y, U g_m \rangle}{\|U g_m\|^2 + h_m^2} \ ,$$

where h_m^2 are regularization parameters. Such estimators recover non-zero coefficients in a space of dimension Q and thus brings no super-resolution. If U is a convolution operator then $\mathcal{B}$ is the Fourier basis and a singular value estimation implements a regularized inverse convolution.

Diagonal Thresholding Estimation   The basis which diagonalizes U^*U rarely provides a sparse signal representation. For example, a Fourier basis which diagonalizes convolution operators does not approximate well signals including singularities.

Donoho [213] introduced more flexibility by looking for a basis $\mathcal{B}$ providing a sparse signal representation, and where a subset of Q vectors $\{g_m\}_{m \in \Gamma_Q}$ are transformed by U in a Riesz basis $\{U g_m\}_{m \in \Gamma_Q}$ of $\mathbf{Im U}$, while the others are set to zero. With an appropriate renormalization, $\{\tilde{\lambda}_m^{-1} U g_m\}_{m \in \Gamma_Q}$ has a biorthogonal basis $\{\tilde{\phi}_m\}_{m \in \Gamma_Q}$ which is normalized $\|\tilde{\phi}_m\| = 1$. The sparse coefficients of f in $\mathcal{B}$ can then be estimated with a thresholding

$$\forall m \in \Gamma_Q \quad , \quad \tilde{a}[m] = \rho_{T_m}(\tilde{\lambda}_m^{-1}\langle Y, \tilde{\phi}_m \rangle) \quad \text{with} \quad \rho_T(x) = x \, \mathbf{1}_{|x|>T} \ ,$$

for thresholds T_m appropriately defined. For classes of signals that are sparse in $\mathcal{B}$, such thresholding estimators may yield a nearly minimax risk, but provides no super-resolution since this non-linear projector remains in a space of dimension Q. This result applies to classes of convolution operators U in wavelet or wavelet packet bases. Diagonal inverse estimators are computationally efficient and potentially optimal in cases where super-resolution is not possible.

1.5.2   Super-Resolution and Compressive Sensing

Suppose that f has a sparse representation in some dictionary $\mathcal{D} = \{g_p\}_{p \in \Gamma}$ of P normalized vectors. The P vectors of the transformed dictionary $\mathcal{D}_U = U\mathcal{D} = \{U g_p\}_{p \in \Gamma}$ belong to the space $\mathbf{Im U}$ of dimension $Q < P$ and thus define a redundant dictionary. Vectors in the approximation support Λ of f are not restricted a priori to a particular subspace of $\mathbb{C}^N$. Super-resolution is possible if the approximation support Λ of f in $\mathcal{D}$ can be estimated by decomposing the noisy data Y over $\mathcal{D}_U$. It depends upon the properties of the approximation support Λ of f in Γ.

Geometric conditions for super-resolution   Let $w_\Lambda = f - f_\Lambda$ be the approximation error of a sparse representation $f_\Lambda = \sum_{p \in \Lambda} a[p] \, g_p$ of f. The observed signal can be written

$$Y = Uf + W = \sum_{p \in \Lambda} a[p] \, U g_p + U w_\Lambda + W \ .$$

If the support Λ can be identified by finding a sparse approximation of Y in $\mathcal{D}_U$

$$Y_\Lambda = \sum_{p \in \Lambda} \tilde{a}[p] \, U g_p$$

then we can recover a super-resolution estimation of f

$$\tilde{F} = \sum_{p \in \Lambda} \tilde{a}[p] \, g_p \ .$$

This shows that super-resolution is possible if the approximation support Λ can be identified by decomposing Y in the redundant transformed dictionary $\mathcal{D}_U$. If the exact recovery criteria is satisfies $ERC(\Lambda) < 1$ and if $\{Ug_p\}_{p\in\Lambda}$ is a Riesz basis then Λ can be recovered using pursuit algorithms, with controlled error bounds.

For most operator U not all sparse approximation sets can be recovered. It is necessary to impose some further geometric conditions on Λ in Γ, which makes super-resolution difficult and often unstable. Numerical applications to sparse spike deconvolution, tomography, super-resolution zooming and inpainting illustrate these results.

Compressive Sensing with Randomness Candès, Tao [138] and Donoho [216] proved that stable super-resolution is possible for any sufficiently sparse signal f, if U is an operator with random coefficients. Compressive sensing then becomes possible, by recovering a close approximation of $f \in \mathbb{C}^N$ from $Q \ll N$ linear measurements [132].

A recovery is stable for a sparse approximation set $|\Lambda| \leqslant M$ only if the corresponding dictionary family $\{Ug_m\}_{m\in\Lambda}$ is a Riesz basis of the space it generates. The *M-restricted isometry conditions* of Candès, Tao and Donoho [216] imposes uniform Riesz bounds for all sets $\Lambda \subset \Gamma$ with $|\Lambda| \leqslant M$:

$$\forall c \in \mathbb{C}^{|\Lambda|} \quad , \quad (1-\delta_M)\,\|c\|^2 \leqslant \|\sum_{m\in\Lambda} c[p]\,Ug_p\|^2 \leqslant (1+\delta_M)\,\|c\|^2 \; . \tag{1.20}$$

This is a strong incoherence condition on the P vectors of $\{Ug_m\}_{m\in\Gamma}$, which imposes that any subset of less than M vectors is nearly uniformly distributed on the unit sphere of $\mathbf{Im}U$. For an orthogonal basis $\mathcal{D} = \{g_m\}_{m\in\Gamma}$, this is possible for $M \leqslant C\,Q\,(\log N)^{-1}$, if U is a matrix with independent Gaussian random coefficients. A pursuit algorithm then provides a stable approximation of any $f \in C^N$ having a sparse approximation from vectors in $\mathcal{D}$.

These results open a new compressive sensing approach to signal acquistion and representation. Instead of first discretizing linearly the signal at a high resolution N and then computing a non-linear representation over M coefficients in some dictionary, compressive sensing measures directly M randomized linear coefficients. A reconstructed signal is then recovered by a non-linear algorithm, producing an error that can be of the same order of magnitude as as the error obtained by the more classic two step approximation process, with a more economic acquisiton process. These results remain valid for several types of random matrices U. Examples of applications to single pixel camera, video super-resolution, new analog to digital converters and MRI imaging are described.

Blind Source Separation Sparsity in redundant dictionaries also provide efficient strategies to separate a family of signals $\{f_s\}_{0\leqslant s<S}$ which are linearly mixed in $K \leqslant S$ observed signals with noise:

$$Y_k[n] = \sum_{s=0}^{S-1} u_{k,s}\,f_s[n] + W_k[n] \;\; \text{for} \;\; 0 \leqslant n < N \;\; \text{and} \;\; 0 \leqslant k < K \; .$$

Most often the mixing matrix $U = \{u_{k,s}\}_{0\leqslant k<K,0\leqslant s<S}$ is unknown. Source separation is a super-resolution problem since $S\,N$ data values must be recovered from $Q = K\,N \leqslant S\,N$ measurements. Not knowing the operator U makes it even more complicated.

If each source f_s has a sparse approximation support Λ_s in a dictionary $\mathcal{D}$, with $\sum_{s=0}^{S-1} |\Lambda_s| \ll N$, then it is likely that the sets $\{\Lambda_s\}_{0\leqslant s<s}$ are nearly disjoint. In this case, the operator U, the supports Λ_s and the sources f_s are approximated by computing sparse approximations of the observed data Y_k in $\mathcal{D}$. The distribution of these coefficients identifies the coefficients of the mixing matrix U, and the nearly disjoint source supports. Time-frequency separation of sounds illustrate these results.

1.6 Travel Guide

Reproducible Computational Science This book covers the whole spectrum from theorems on functions of continuous variables to fast discrete algorithms and their applications. Section 1.1.2 argues that models based on continuous time functions give useful asymptotic results for understanding the behavior of discrete algorithms. Yet, a mathematical analysis alone is often unable

to predict fully the behavior and suitability of algorithms for specific signals. Experiments are necessary and such experiments should be reproducible, just like experiments in other fields of sciences [123].

The reproducibility of experiments requires having the complete software and full source code for inspection, modification and application under varied parameter settings. Following this perspective, computational algorithms presented in this book are available as Matlab subroutines or in other software packages. Figures can be reproduced and the source code is available. Softwares, demonstrations and exercise solutions are available on the Web site *http://wavelet-tour.com.*

Book Road Map   Some redundancy is introduced between sections to avoid imposing a linear progression through the book. The preface describes several possible programs for a sparse signal processing course.

All theorems are explained in the text and reading the proofs is not necessary to understand the results. Most theorems of the book are proved in detail, and important techniques are included. Exercises at the end of each chapter give examples of mathematical, algorithmic and numerical applications, ordered by level of difficulty from **1** to **4**, and solutions can be found in *http://wavelet-tour.com.*

The book begins with Chapters 2 and 3, which review the Fourier transform and linear discrete signal processing. They provide the necessary background for readers with no signal processing background. Important properties of linear operators, projectors and vector spaces can be found in Appendix A. Local time-frequency transforms and dictionaries are presented in Chapter 4. The wavelet and windowed Fourier transforms are introduced and compared. The measurement of instantaneous frequencies illustrates the limitations of time-frequency resolution. Dictionary stability and redundancy is introduced in Chapter 5 through the frame theory, with examples of windowed Fourier, wavelet and curvelet frames. Chapter 6 explains the relations between wavelet coefficient amplitude and local signal regularity. It is applied to the detection of singularities, edges and to the analysis of multifractals.

Wavelet bases and fast filter bank algorithms are important tools presented in Chapter 7. An overdose of orthonormal bases can strike the reader while studying the construction and properties of wavelet packets and local cosine bases in Chapter 8. It is thus important to read Chapter 9, which describes sparse approximations in bases. Signal compression and denoising applications of Chapters 10 and 11 give life to most theoretical and algorithmic results of the book. These chapters offer a practical perspective on the relevance of these linear and non-linear signal processing algorithms. Chapter 12 introduce sparse decompositions in redundant dictionaries and their applications. The resolution of inverse problems is studied in Chapter 13, with super-resolution, compressive sensing and source separation.

II

Fourier Kingdom

The story begins in 1807 when Fourier presents a memoir to the Institut de France, where he claims that any periodic function can be represented as a series of harmonically related sinusoids. This idea had a profound impact in mathematical analysis, physics and engineering, but it took one and a half centuries to understand the convergence of Fourier series and complete the theory of Fourier integrals.

Fourier was motivated by the study of heat diffusion, which is governed by a linear differential equation. However, the Fourier transform diagonalizes all linear time-invariant operators, which are the building blocks of signal processing. It is therefore not only the starting point of our exploration but the basis of all further developments.

2.1 Linear Time-Invariant Filtering

Classical signal processing operations such as signal transmission, stationary noise removal or predictive coding are implemented with linear time-invariant operators. The time invariance of an operator L means that if the input $f(t)$ is delayed by τ, $f_\tau(t) = f(t - \tau)$, then the output is also delayed by τ:

$$g(t) = Lf(t) \quad \Rightarrow \quad g(t - \tau) = Lf_\tau(t). \tag{2.1}$$

For numerical stability, the operator L must have a weak form of continuity, which means that Lf is modified by a small amount if f is slightly modified. This weak continuity is formalized by the theory of distributions [60, 63], which guarantees that we are on a safe ground without further worrying about it.

2.1.1 Impulse Response

Linear time-invariant systems are characterized by their response to a Dirac impulse, defined in Appendix A.7. If f is continuous, its value at t is obtained by an "integration" against a Dirac located at t. Let $\delta_u(t) = \delta(t - u)$:

$$f(t) = \int_{-\infty}^{+\infty} f(u)\, \delta_u(t)\, du.$$

The continuity and linearity of L imply that

$$Lf(t) = \int_{-\infty}^{+\infty} f(u)\, L\delta_u(t)\, du.$$

Let h be the impulse response of L:

$$h(t) = L\delta(t).$$

The time-invariance proves that $L\delta_u(t) = h(t-u)$ and hence

$$Lf(t) = \int_{-\infty}^{+\infty} f(u)\, h(t-u)\, du = \int_{-\infty}^{+\infty} h(u) f(t-u)\, du = h \star f\,(t). \tag{2.2}$$

A time-invariant linear filter is thus equivalent to a convolution with the impulse response h. The continuity of f is not necessary. This formula remains valid for any signal f for which the convolution integral converges.

Let us recall a few useful properties of convolution products:

- Commutativity

$$f \star h(t) = h \star f(t). \tag{2.3}$$

- Differentiation

$$\frac{d}{dt}(f \star h)(t) = \frac{df}{dt} \star h(t) = f \star \frac{dh}{dt}(t). \tag{2.4}$$

- Dirac convolution

$$f \star \delta_\tau(t) = f(t-\tau). \tag{2.5}$$

Stability and Causality   A filter is said to be *causal* if $Lf(t)$ does not depend on the values $f(u)$ for $u > t$. Since

$$Lf(t) = \int_{-\infty}^{+\infty} h(u)\, f(t-u)\, du,$$

this means that $h(u) = 0$ for $u < 0$. Such impulse responses are said to be *causal*.

The *stability* property guarantees that $Lf(t)$ is bounded if $f(t)$ is bounded. Since

$$|Lf(t)| \leqslant \int_{-\infty}^{+\infty} |h(u)|\, |f(t-u)|\, du \leqslant \sup_{u \in \mathbb{R}} |f(u)| \int_{-\infty}^{+\infty} |h(u)|\, du,$$

it is sufficient that $\int_{-\infty}^{+\infty} |h(u)|\, du < +\infty$. One can verify that this condition is also necessary if h is a function. We thus say that h is *stable* if it is integrable.

Example 2.1. *An amplification and delay system is defined by*

$$Lf(t) = \lambda\, f(t-\tau).$$

The impulse response of this filter is $h(t) = \lambda\, \delta(t-\tau)$.

Example 2.2. *A uniform averaging of f over intervals of size T is calculated by*

$$Lf(t) = \frac{1}{T} \int_{t-T/2}^{t+T/2} f(u)\, du.$$

This integral can be rewritten as a convolution of f with the impulse response $h = 1/T\, \mathbf{1}_{[-T/2,T/2]}$.

2.1.2   Transfer Functions

Complex exponentials $e^{i\omega t}$ are eigenvectors of convolution operators. Indeed

$$Le^{i\omega t} = \int_{-\infty}^{+\infty} h(u)\, e^{i\omega(t-u)}\, du,$$

which yields

$$Le^{i\omega t} = e^{it\omega} \int_{-\infty}^{+\infty} h(u)\, e^{-i\omega u}\, du = \hat{h}(\omega)\, e^{i\omega t}.$$

The eigenvalue

$$\hat{h}(\omega) = \int_{-\infty}^{+\infty} h(u)\, e^{-i\omega u}\, du$$

is the Fourier transform of h at the frequency ω. Since complex sinusoidal waves $e^{i\omega t}$ are the eigenvectors of time-invariant linear systems, it is tempting to try to decompose any function f as a sum of these eigenvectors. We are then able to express Lf directly from the eigenvalues $\hat{h}(\omega)$. The Fourier analysis proves that under weak conditions on f, it is indeed possible to write it as a Fourier integral.

2.2 Fourier Integrals

To avoid convergence issues, the Fourier integral is first defined over the space $\mathbf{L}^1(\mathbb{R})$ of integrable functions [53]. It is then extended to the space $\mathbf{L}^2(\mathbb{R})$ of finite energy functions [22].

2.2.1 Fourier Transform in $\mathbf{L}^1(\mathbb{R})$

The Fourier integral

$$\hat{f}(\omega) = \int_{-\infty}^{+\infty} f(t)\, \mathrm{e}^{-i\omega t}\, dt \tag{2.6}$$

measures "how much" oscillations at the frequency ω there is in f. If $f \in \mathbf{L}^1(\mathbb{R})$ this integral does converge and

$$|\hat{f}(\omega)| \leqslant \int_{-\infty}^{+\infty} |f(t)|\, dt < +\infty. \tag{2.7}$$

The Fourier transform is thus bounded, and one can verify that it is a continuous function of ω (Exercise 2.1). If $\hat{f}$ is also integrable, the following theorem gives the inverse Fourier transform.

Theorem 2.1 (Inverse Fourier Transform). *If $f \in \mathbf{L}^1(\mathbb{R})$ and $\hat{f} \in \mathbf{L}^1(\mathbb{R})$ then*

$$f(t) = \frac{1}{2\pi} \int_{-\infty}^{+\infty} \hat{f}(\omega)\, \mathrm{e}^{i\omega t}\, d\omega. \tag{2.8}$$

Proof. Replacing $\hat{f}(\omega)$ by its integral expression yields

$$\frac{1}{2\pi} \int_{-\infty}^{+\infty} \hat{f}(\omega) \exp(i\omega t)\, d\omega = \frac{1}{2\pi} \int_{-\infty}^{+\infty} \left(\int_{-\infty}^{+\infty} f(u) \exp[i\omega(t - u)]\, du \right) d\omega.$$

We cannot apply the Fubini Theorem A.2 directly because $f(u) \exp[i\omega(t - u)]$ is not integrable in $\mathbb{R}^2$. To avoid this technical problem, we multiply by $\exp(-\varepsilon^2 \omega^2 / 4)$ which converges to 1 when ε goes to 0. Let us define

$$I_\varepsilon(t) = \frac{1}{2\pi} \int_{-\infty}^{+\infty} \left(\int_{-\infty}^{+\infty} f(u) \exp\left(\frac{-\varepsilon^2 \omega^2}{4} \right) \exp[i\omega(t - u)]\, du \right) d\omega. \tag{2.9}$$

We compute I_ε in two different ways using the Fubini theorem. The integration with respect to u gives

$$I_\varepsilon(t) = \frac{1}{2\pi} \int_{-\infty}^{+\infty} \hat{f}(\omega) \exp\left(\frac{-\varepsilon^2 \omega^2}{4} \right) \exp(i\omega t)\, d\omega.$$

Since

$$\left| \hat{f}(\omega) \exp\left(\frac{-\varepsilon^2 \omega^2}{4} \right) \exp[i\omega(t - u)] \right| \leqslant |\hat{f}(\omega)|$$

and since $\hat{f}$ is integrable, we can apply the dominated convergence Theorem A.1, which proves that

$$\lim_{\varepsilon \to 0} I_\varepsilon(t) = \frac{1}{2\pi} \int_{-\infty}^{+\infty} \hat{f}(\omega) \exp(i\omega t)\, d\omega. \tag{2.10}$$

Let us now compute the integral (2.9) differently by applying the Fubini theorem and integrating with respect to ω:

$$I_\varepsilon(t) = \int_{-\infty}^{+\infty} g_\varepsilon(t - u)\, f(u)\, du, \tag{2.11}$$

with

$$g_\varepsilon(x) = \frac{1}{2\pi} \int_{-\infty}^{+\infty} \exp(ix\omega) \exp\left(\frac{-\varepsilon^2 \omega^2}{4} \right)\, d\omega.$$

A change of variable $\omega' = \varepsilon \omega$ shows that $g_\varepsilon(x) = \varepsilon^{-1} g_1(\varepsilon^{-1} x)$, and it is proved in (2.32) that $g_1(x) = \pi^{-1/2}\, \mathrm{e}^{-x^2}$. The Gaussian g_1 has an integral equal to 1 and a fast decay. The squeezed Gaussians g_ε have an integral that remains equal to 1, and thus they converge to a Dirac δ when ε goes to 0. By inserting (2.11) one can thus verify that

$$\lim_{\varepsilon \to 0} \int_{-\infty}^{+\infty} |I_\varepsilon(t) - f(t)|\, dt = \lim_{\varepsilon \to 0} \iint g_\varepsilon(t - u)\, |f(u) - f(t)|\, du\, dt = 0.$$

Inserting (2.10) proves (2.8). ∎ ∎

The inversion formula (2.8) decomposes f as a sum of sinusoidal waves $e^{i\omega t}$ of amplitude $\hat{f}(\omega)$. By using this formula, we can show (Exercise 2.1) that the hypothesis $\hat{f} \in \mathbf{L}^1(\mathbb{R})$ implies that f must be continuous. The reconstruction (2.8) is therefore not proved for discontinuous functions. The extension of the Fourier transform to the space $\mathbf{L}^2(\mathbb{R})$ will address this issue.

The most important property of the Fourier transform for signal processing applications is the convolution theorem. It is another way to express the fact that sinusoidal waves $e^{it\omega}$ are eigenvalues of convolution operators.

Theorem 2.2 (Convolution). *Let $f \in \mathbf{L}^1(\mathbb{R})$ and $h \in \mathbf{L}^1(\mathbb{R})$. The function $g = h \star f$ is in $\mathbf{L}^1(\mathbb{R})$ and*

$$\hat{g}(\omega) = \hat{h}(\omega)\,\hat{f}(\omega). \tag{2.12}$$

Proof.

$$\hat{g}(\omega) = \int_{-\infty}^{+\infty} \exp(-it\omega)\left(\int_{-\infty}^{+\infty} f(t-u)\,h(u)\,du\right) dt.$$

Since $|f(t-u)||h(u)|$ is integrable in $\mathbb{R}^2$, we can apply the Fubini Theorem A.2, and the change of variable $(t, u) \to (v = t - u, u)$ yields

$$
\begin{aligned}
\hat{g}(\omega) &= \int_{-\infty}^{+\infty}\int_{-\infty}^{+\infty} \exp[-i(u+v)\omega]\,f(v)\,h(u)\,du\,dv \\
&= \left(\int_{-\infty}^{+\infty} \exp(-iv\omega)\,f(v)\,dv\right)\left(\int_{-\infty}^{+\infty} \exp(-iu\omega)\,h(u)\,du\right),
\end{aligned}
$$

which verifies (2.12).   ∎   ∎

The response $Lf = g = f \star h$ of a linear time-invariant system can be calculated from its Fourier transform $\hat{g}(\omega) = \hat{f}(\omega)\,\hat{h}(\omega)$ with the inverse Fourier formula

$$g(t) = \frac{1}{2\pi} \int_{-\infty}^{+\infty} \hat{g}(\omega)\,e^{i\omega t}\,d\omega, \tag{2.13}$$

which yields

$$Lf(t) = \frac{1}{2\pi} \int_{-\infty}^{+\infty} \hat{h}(\omega)\,\hat{f}(\omega)\,e^{i\omega t}\,d\omega. \tag{2.14}$$

Each frequency component $e^{it\omega}$ of amplitude $\hat{f}(\omega)$ is amplified or attenuated by $\hat{h}(\omega)$. Such a convolution is thus called a *frequency filtering*, and $\hat{h}$ is the *transfer function* of the filter.

The following table summarizes important properties of the Fourier transform, often used in calculations. Most of these formulas are proved with a change of variable in the Fourier integral.

Property	Function $f(t)$	Fourier Transform $\hat{f}(\omega)$	
Inverse	$\hat{f}(t)$	$2\pi\,f(-\omega)$	(2.15)
Convolution	$f_1 \star f_2(t)$	$\hat{f}_1(\omega)\,\hat{f}_2(\omega)$	(2.16)
Multiplication	$f_1(t)\,f_2(t)$	$\dfrac{1}{2\pi}\,\hat{f}_1 \star \hat{f}_2(\omega)$	(2.17)
Translation	$f(t-u)$	$e^{-iu\omega}\,\hat{f}(\omega)$	(2.18)
Modulation	$e^{i\xi t}\,f(t)$	$\hat{f}(\omega - \xi)$	(2.19)
Scaling	$f(t/s)$	$|s|\,\hat{f}(s\,\omega)$	(2.20)
Time derivatives	$f^{(p)}(t)$	$(i\omega)^p\,\hat{f}(\omega)$	(2.21)
Frequency derivatives	$(-it)^p\,f(t)$	$\hat{f}^{(p)}(\omega)$	(2.22)
Complex conjugate	$f^*(t)$	$\hat{f}^*(-\omega)$	(2.23)
Hermitian symmetry	$f(t) \in \mathbb{R}$	$\hat{f}(-\omega) = \hat{f}^*(\omega)$	(2.24)

2.2.2 Fourier Transform in $\mathbf{L}^2(\mathbb{R})$

The Fourier transform of the indicator function $f = \mathbf{1}_{[-1,1]}$ is

$$\hat{f}(\omega) = \int_{-1}^{1} \mathrm{e}^{-i\omega t}\, dt = \frac{2\sin\omega}{\omega}.$$

This function is not integrable because f is not continuous, but its square is integrable. The inverse Fourier transform Theorem 2.1 thus does not apply. This motivates the extension of the Fourier transform to the space $\mathbf{L}^2(\mathbb{R})$ of functions f with a finite energy $\int_{-\infty}^{+\infty} |f(t)|^2\, dt < +\infty$. By working in the Hilbert space $\mathbf{L}^2(\mathbb{R})$, we also have access to all the facilities provided by the existence of an inner product. The inner product of $f \in \mathbf{L}^2(\mathbb{R})$ and $g \in \mathbf{L}^2(\mathbb{R})$ is

$$\langle f, g \rangle = \int_{-\infty}^{+\infty} f(t)\, g^*(t)\, dt,$$

and the resulting norm in $\mathbf{L}^2(\mathbb{R})$ is

$$\|f\|^2 = \langle f, f \rangle = \int_{-\infty}^{+\infty} |f(t)|^2\, dt.$$

The following theorem proves that inner products and norms in $\mathbf{L}^2(\mathbb{R})$ are conserved by the Fourier transform up to a factor of 2π. Equations (2.25) and (2.26) are called respectively the *Parseval* and *Plancherel* formulas.

Theorem 2.3. *If f and h are in $\mathbf{L}^1(\mathbb{R}) \cap \mathbf{L}^2(\mathbb{R})$ then*

$$\int_{-\infty}^{+\infty} f(t)\, h^*(t)\, dt = \frac{1}{2\pi} \int_{-\infty}^{+\infty} \hat{f}(\omega)\, \hat{h}^*(\omega)\, d\omega. \tag{2.25}$$

For $h = f$ it follows that

$$\int_{-\infty}^{+\infty} |f(t)|^2\, dt = \frac{1}{2\pi} \int_{-\infty}^{+\infty} |\hat{f}(\omega)|^2\, d\omega. \tag{2.26}$$

Proof. Let $g = f \star \bar{h}$ with $\bar{h}(t) = h^*(-t)$. The convolution Theorem 2.2 and property (2.23) show that $\hat{g}(\omega) = \hat{f}(\omega)\, \hat{h}^*(\omega)$. The reconstruction formula (2.8) applied to $g(0)$ yields

$$\int_{-\infty}^{+\infty} f(t)\, h^*(t)\, dt = g(0) = \frac{1}{2\pi} \int_{-\infty}^{+\infty} \hat{g}(\omega)\, d\omega = \frac{1}{2\pi} \int_{-\infty}^{+\infty} \hat{f}(\omega)\, \hat{h}^*(\omega)\, d\omega.$$

∎ ∎

Density Extension in $\mathbf{L}^2(\mathbb{R})$ If $f \in \mathbf{L}^2(\mathbb{R})$ but $f \notin \mathbf{L}^1(\mathbb{R})$, its Fourier transform cannot be calculated with the Fourier integral (2.6) because $f(t)\, \mathrm{e}^{i\omega t}$ is not integrable. It is defined as a limit using the Fourier transforms of functions in $\mathbf{L}^1(\mathbb{R}) \cap \mathbf{L}^2(\mathbb{R})$.

Since $\mathbf{L}^1(\mathbb{R}) \cap \mathbf{L}^2(\mathbb{R})$ is dense in $\mathbf{L}^2(\mathbb{R})$, one can find a family $\{f_n\}_{n \in \mathbb{Z}}$ of functions in $\mathbf{L}^1(\mathbb{R}) \cap \mathbf{L}^2(\mathbb{R})$ that converges to f:

$$\lim_{n \to +\infty} \|f - f_n\| = 0.$$

Since $\{f_n\}_{n \in \mathbb{Z}}$ converges, it is a Cauchy sequence, which means that $\|f_n - f_p\|$ is arbitrarily small if n and p are large enough. Moreover, $f_n \in \mathbf{L}^1(\mathbb{R})$, so its Fourier transform $\hat{f}_n$ is well defined. The Plancherel formula (2.26) proves that $\{\hat{f}_n\}_{n \in \mathbb{Z}}$ is also a Cauchy sequence because

$$\|\hat{f}_n - \hat{f}_p\| = \sqrt{2\pi}\, \|f_n - f_p\|$$

is arbitrarily small for n and p large enough. A Hilbert space (Appendix A.2) is complete, which means that all Cauchy sequences converge to an element of the space. Hence, there exists $\hat{f} \in \mathbf{L}^2(\mathbb{R})$ such that

$$\lim_{n \to +\infty} \|\hat{f} - \hat{f}_n\| = 0.$$

By definition, $\hat{f}$ is the Fourier transform of f. This extension of the Fourier transform to $\mathbf{L}^2(\mathbb{R})$ satisfies the convolution theorem, the Parseval and Plancherel formulas, as well as all properties (2.15-2.24).

Diracs Diracs are often used in calculations; their properties are summarized in Appendix A.7. A Dirac δ associates to a function its value at $t = 0$. Since $e^{i\omega t} = 1$ at $t = 0$ it seems reasonable to define its Fourier transform by

$$\hat{\delta}(\omega) = \int_{-\infty}^{+\infty} \delta(t)\, e^{-i\omega t}\, dt = 1. \tag{2.27}$$

This formula is justified mathematically by the extension of the Fourier transform to tempered distributions [60, 63].

2.2.3 Examples

The following examples often appear in Fourier calculations. They also illustrate important Fourier transform properties.

• The *indicator function* $f = \mathbf{1}_{[-T,T]}$ is discontinuous at $t = \pm T$. Its Fourier transform is therefore not integrable:

$$\hat{f}(\omega) = \int_{-T}^{T} e^{-i\omega t}\, dt = \frac{2\sin(T\omega)}{\omega}. \tag{2.28}$$

• An *ideal low-pass filter* has a transfer function $\hat{\phi} = \mathbf{1}_{[-\xi,\xi]}$ that selects low frequencies over $[-\xi, \xi]$. The impulse response is calculated with the inverse Fourier integral (2.8):

$$\phi(t) = \frac{1}{2\pi} \int_{-\xi}^{\xi} e^{i\omega t}\, d\omega = \frac{\sin(\xi t)}{\pi t}. \tag{2.29}$$

• A *passive electronic circuit* implements analog filters with resistances, capacities and inductors. The input voltage $f(t)$ is related to the output voltage $g(t)$ by a differential equation with constant coefficients:

$$\sum_{k=0}^{K} a_k\, f^{(k)}(t) = \sum_{k=0}^{M} b_k\, g^{(k)}(t). \tag{2.30}$$

Suppose that the circuit is not charged for $t < 0$, which means that $f(t) = g(t) = 0$. The output g is a linear time-invariant function of f and can thus be written $g = f \star \phi$. Computing the Fourier transform of (2.30) and applying (2.22) proves that

$$\hat{\phi}(\omega) = \frac{\hat{g}(\omega)}{\hat{f}(\omega)} = \frac{\sum_{k=0}^{K} a_k\, (i\omega)^k}{\sum_{k=0}^{M} b_k\, (i\omega)^k}. \tag{2.31}$$

It is therefore a rational function of $i\omega$. An ideal low-pass transfer function $\mathbf{1}_{[-\xi,\xi]}$ thus cannot be implemented by an analog circuit. It must be approximated by a rational function. Chebyshev or Butterworth filters are often used for this purpose [13].

• A *Gaussian* $f(t) = \exp(-t^2)$ is a $\mathbf{C}^\infty$ function with a fast asymptotic decay. Its Fourier transform is also a Gaussian:

$$\hat{f}(\omega) = \sqrt{\pi}\, \exp(-\omega^2/4). \tag{2.32}$$

This Fourier transform is computed by showing with an integration by parts that $\hat{f}(\omega) = \int_{-\infty}^{+\infty} \exp(-t^2)\, e^{-i\omega t} dt$ is differentiable and satisfies the differential equation

$$2\, \hat{f}'(\omega) + \omega\, \hat{f}(\omega) = 0. \tag{2.33}$$

The solution of this equation is a Gaussian $\hat{f}(\omega) = K \exp(-\omega^2/4)$, and since $\hat{f}(0) = \int_{-\infty}^{+\infty} \exp(-t^2)\, dt = \sqrt{\pi}$, we obtain (2.32).

• A Gaussian *chirp* $f(t) = \exp[-(a - ib)t^2]$ has a Fourier transform calculated with a similar differential equation:

$$\hat{f}(\omega) = \sqrt{\frac{\pi}{a - ib}}\, \exp\left(\frac{-(a + ib)\omega^2}{4(a^2 + b^2)}\right). \tag{2.34}$$

- A translated *Dirac* $\delta_\tau(t) = \delta(t - \tau)$ has a Fourier transform calculated by evaluating $e^{-i\omega t}$ at $t = \tau$:

$$\hat{\delta}_\tau(\omega) = \int_{-\infty}^{+\infty} \delta(t - \tau) e^{-i\omega t} \, dt = e^{-i\omega\tau}. \tag{2.35}$$

- The *Dirac comb* is a sum of translated Diracs

$$c(t) = \sum_{n=-\infty}^{+\infty} \delta(t - nT)$$

that is used to uniformly sample analog signals. Its Fourier transform is derived from (2.35):

$$\hat{c}(\omega) = \sum_{n=-\infty}^{+\infty} e^{-inT\omega}. \tag{2.36}$$

The Poisson formula proves that it is also equal to a Dirac comb with a spacing equal to $2\pi/T$.

Theorem 2.4 (Poisson Formula). *In the sense of distribution equalities (A.29),*

$$\sum_{n=-\infty}^{+\infty} e^{-inT\omega} = \frac{2\pi}{T} \sum_{k=-\infty}^{+\infty} \delta\left(\omega - \frac{2\pi k}{T}\right). \tag{2.37}$$

Proof. The Fourier transform $\hat{c}$ in (2.36) is periodic with period $2\pi/T$. To verify the Poisson formula, it is therefore sufficient to prove that the restriction of $\hat{c}$ to $[-\pi/T, \pi/T]$ is equal to $2\pi/T \, \delta$. The formula (2.37) is proved in the sense of a distribution equality (A.29) by showing that for any test function $\hat{\theta}(\omega)$ with a support included in $[-\pi/T, \pi/T]$,

$$\langle \hat{c}, \hat{\theta} \rangle = \lim_{N\to+\infty} \int_{-\infty}^{+\infty} \sum_{n=-N}^{N} \exp(-inT\omega) \, \hat{\theta}(\omega) \, d\omega = \frac{2\pi}{T} \, \hat{\theta}(0).$$

The sum of the geometric series is

$$\sum_{n=-N}^{N} \exp(-inT\omega) = \frac{\sin[(N + 1/2)T\omega]}{\sin[T\omega/2]}. \tag{2.38}$$

Hence

$$\langle \hat{c}, \hat{\theta} \rangle = \lim_{N\to+\infty} \frac{2\pi}{T} \int_{-\pi/T}^{\pi/T} \frac{\sin[(N + 1/2)T\omega]}{\pi\omega} \frac{T\omega/2}{\sin[T\omega/2]} \hat{\theta}(\omega) \, d\omega. \tag{2.39}$$

Let

$$\hat{\psi}(\omega) = \begin{cases} \hat{\theta}(\omega) \frac{T\omega/2}{\sin[T\omega/2]} & \text{if } |\omega| \leqslant \pi/T \\ 0 & \text{if } |\omega| > \pi/T \end{cases}$$

and $\psi(t)$ be the inverse Fourier transform of $\hat{\psi}(\omega)$. Since $2\omega^{-1}\sin(a\omega)$ is the Fourier transform of $\mathbf{1}_{[-a,a]}(t)$, the Parseval formula (2.25) implies

$$\begin{aligned} \langle \hat{c}, \hat{\theta} \rangle &= \lim_{N\to+\infty} \frac{2\pi}{T} \int_{-\infty}^{+\infty} \frac{\sin[(N + 1/2)T\omega]}{\pi\omega} \hat{\psi}(\omega) \, d\omega \\ &= \lim_{N\to+\infty} \frac{2\pi}{T} \int_{-(N+1/2)T}^{(N+1/2)T} \psi(t) \, dt. \end{aligned} \tag{2.40}$$

When N goes to $+\infty$ the integral converges to $\hat{\psi}(0) = \hat{\theta}(0)$. ∎ ∎

2.3 Properties

2.3.1 Regularity and Decay

The global regularity of a signal f depends on the decay of $|\hat{f}(\omega)|$ when the frequency ω increases. The differentiability of f is studied. If $\hat{f} \in \mathbf{L}^1(\mathbb{R})$, then the Fourier inversion formula (2.8) implies

that f is continuous and bounded:

$$|f(t)| \leqslant \frac{1}{2\pi} \int_{-\infty}^{+\infty} |e^{i\omega t} \hat{f}(\omega)| \, d\omega = \frac{1}{2\pi} \int_{-\infty}^{+\infty} |\hat{f}(\omega)| \, d\omega < +\infty \,. \tag{2.41}$$

The next theorem applies this property to obtain a sufficient condition that guarantees the differentiability of f at any order p.

Theorem 2.5. *A function f is bounded and p times continuously differentiable with bounded derivatives if*

$$\int_{-\infty}^{+\infty} |\hat{f}(\omega)| \, (1 + |\omega|^p) \, d\omega < +\infty \,. \tag{2.42}$$

Proof. The Fourier transform of the k^{th} order derivative $f^{(k)}(t)$ is $(i\omega)^k \hat{f}(\omega)$. Applying (2.41) to this derivative proves that

$$|f^{(k)}(t)| \leqslant \int_{-\infty}^{+\infty} |\hat{f}(\omega)| \, |\omega|^k \, d\omega.$$

Condition (2.42) implies that $\int_{-\infty}^{+\infty} |\hat{f}(\omega)||\omega|^k \, d\omega < +\infty$ for any $k \leqslant p$, so $f^{(k)}(t)$ is continuous and bounded. ∎

This result proves that if there exist a constant K and $\varepsilon > 0$ such that

$$|\hat{f}(\omega)| \leqslant \frac{K}{1 + |\omega|^{p+1+\varepsilon}}, \quad \text{then} \quad f \in \mathbf{C}^p.$$

If $\hat{f}$ has a compact support then (2.42) implies that $f \in \mathbf{C}^\infty$.

The decay of $|\hat{f}(\omega)|$ depends on the worst singular behavior of f. For example, $f = \mathbf{1}_{[-T,T]}$ is discontinuous at $t = \pm T$, so $|\hat{f}(\omega)|$ decays like $|\omega|^{-1}$. In this case, it could also be important to know that $f(t)$ is regular for $t \neq \pm T$. This information cannot be derived from the decay of $|\hat{f}(\omega)|$. To characterize local regularity of a signal f it is necessary to decompose it over waveforms that are well localized in time, as opposed to sinusoidal waves $e^{i\omega t}$. Section 6.1.3 explains that wavelets are particularly well adapted to this purpose.

2.3.2 Uncertainty Principle

Can we construct a function f whose energy is well localized in time and whose Fourier transform $\hat{f}$ has an energy concentrated in a small frequency neighborhood? The Dirac $\delta(t - u)$ has a support restricted to $t = u$ but its Fourier transform $e^{-iu\omega}$ has an energy uniformly spread over all frequencies. We know that $|\hat{f}(\omega)|$ decays quickly at high frequencies only if f has regular variations in time. The energy of f must therefore be spread over a relatively large domain.

To reduce the time spread of f, we can scale it by $s < 1$ while maintaining constant its total energy. If

$$f_s(t) = \frac{1}{\sqrt{s}} f\left(\frac{t}{s}\right) \quad \text{then} \quad \|f_s\|^2 = \|f\|^2.$$

The Fourier transform $\hat{f}_s(\omega) = \sqrt{s} \, \hat{f}(s\omega)$ is dilated by $1/s$ so we lose in frequency localization what we gained in time. Underlying is a trade-off between time and frequency localization.

Time and frequency energy concentrations are restricted by the Heisenberg uncertainty principle. This principle has a particularly important interpretation in quantum mechanics as an uncertainty as to the position and momentum of a free particle. The state of a one-dimensional particle is described by a wave function $f \in \mathbf{L}^2(\mathbb{R})$. The probability density that this particle is located at t is $\frac{1}{\|f\|^2} |f(t)|^2$. The probability density that its momentum is equal to ω is $\frac{1}{2\pi\|f\|^2} |\hat{f}(\omega)|^2$. The average location of this particle is

$$u = \frac{1}{\|f\|^2} \int_{-\infty}^{+\infty} t \, |f(t)|^2 \, dt, \tag{2.43}$$

and the average momentum is

$$\xi = \frac{1}{2\pi \|f\|^2} \int_{-\infty}^{+\infty} \omega \, |\hat{f}(\omega)|^2 \, d\omega. \tag{2.44}$$

The variances around these average values are respectively

$$\sigma_t^2 = \frac{1}{\|f\|^2} \int_{-\infty}^{+\infty} (t-u)^2 \, |f(t)|^2 \, dt \tag{2.45}$$

and

$$\sigma_\omega^2 = \frac{1}{2\pi \|f\|^2} \int_{-\infty}^{+\infty} (\omega - \xi)^2 \, |\hat{f}(\omega)|^2 \, d\omega. \tag{2.46}$$

The larger σ_t, the more uncertainty there is concerning the position of the free particle; the larger σ_ω, the more uncertainty there is concerning its momentum.

Theorem 2.6 (Heisenberg Uncertainty). *The temporal variance and the frequency variance of $f \in \mathbf{L}^2(\mathbb{R})$ satisfy*

$$\sigma_t^2 \, \sigma_\omega^2 \geqslant \frac{1}{4} \, . \tag{2.47}$$

This inequality is an equality if and only if there exist $(u, \xi, a, b) \in \mathbb{R}^2 \times \mathbb{C}^2$ such that

$$f(t) = a \, \exp[i\xi t - b(t-u)^2]. \tag{2.48}$$

Proof. The following proof due to Weyl [68] supposes that $\lim_{|t| \to +\infty} \sqrt{t} f(t) = 0$, but the theorem is valid for any $f \in \mathbf{L}^2(\mathbb{R})$. If the average time and frequency localization of f is u and ξ, then the average time and frequency location of $\exp(-i\xi t) f(t + u)$ is zero. It is thus sufficient to prove the theorem for $u = \xi = 0$. Observe that

$$\sigma_t^2 \, \sigma_\omega^2 = \frac{1}{2\pi \|f\|^4} \int_{-\infty}^{+\infty} |t \, f(t)|^2 \, dt \int_{-\infty}^{+\infty} |\omega \, \hat{f}(\omega)|^2 \, d\omega. \tag{2.49}$$

Since $i\omega \hat{f}(\omega)$ is the Fourier transform of $f'(t)$, the Plancherel identity (2.26) applied to $i\omega \hat{f}(\omega)$ yields

$$\sigma_t^2 \sigma_\omega^2 = \frac{1}{\|f\|^4} \int_{-\infty}^{+\infty} |t \, f(t)|^2 \, dt \int_{-\infty}^{+\infty} |f'(t)|^2 \, dt. \tag{2.50}$$

Schwarz's inequality implies

$$
\begin{aligned}
\sigma_t^2 \sigma_\omega^2 \quad &\geqslant \quad \frac{1}{\|f\|^4} \left[\int_{-\infty}^{+\infty} |t \, f'(t) \, f^*(t)| \, dt \right]^2 \\
&\geqslant \quad \frac{1}{\|f\|^4} \left[\int_{-\infty}^{+\infty} \frac{t}{2} \, [f'(t) \, f^*(t) + f'^*(t) \, f(t)] \, dt \right]^2 \\
&\geqslant \quad \frac{1}{4\|f\|^4} \left[\int_{-\infty}^{+\infty} t \, (|f(t)|^2)' \, dt \right]^2 .
\end{aligned}
$$

Since $\lim_{|t| \to +\infty} \sqrt{t} \, f(t) = 0$, an integration by parts gives

$$\sigma_t^2 \sigma_\omega^2 \geqslant \frac{1}{4\|f\|^4} \left[\int_{-\infty}^{+\infty} |f(t)|^2 \, dt \right]^2 = \frac{1}{4}. \tag{2.51}$$

To obtain an equality, Schwarz's inequality applied to (2.50) must be an equality. This implies that there exists $b \in \mathbb{C}$ such that

$$f'(t) = -2 \, b \, t \, f(t). \tag{2.52}$$

Hence, there exists $a \in \mathbb{C}$ such that $f(t) = a \, \exp(-bt^2)$. The other steps of the proof are then equalities so that the lower bound is indeed reached. When $u \neq 0$ and $\xi \neq 0$ the corresponding time and frequency translations yield (2.48). ■   ■

In quantum mechanics, this theorem shows that we cannot reduce arbitrarily the uncertainty as to the position and the momentum of a free particle. In signal processing, the modulated Gaussians (2.48) that have a minimum joint time-frequency localization are called Gabor chirps. As expected, they are smooth functions with a fast time asymptotic decay.

Compact Support   Despite the Heisenberg uncertainty bound, we might still be able to construct a function of compact support whose Fourier transform has a compact support. Such a function would be very useful in constructing a finite impulse response filter with a band-limited transfer function. Unfortunately, the following theorem proves that it does not exist.

Theorem 2.7. *If $f \neq 0$ has a compact support then $\hat{f}(\omega)$ cannot be zero on a whole interval. Similarly, if $\hat{f} \neq 0$ has a compact support then $f(t)$ cannot be zero on a whole interval.*

Proof. We prove only the first statement, since the second is derived from the first by applying the Fourier transform. If $\hat{f}$ has a compact support included in $[-b, b]$ then

$$f(t) = \frac{1}{2\pi} \int_{-b}^{b} \hat{f}(\omega) \exp(i\omega t) \, d\omega. \tag{2.53}$$

If $f(t) = 0$ for $t \in [c, d]$, by differentiating n times under the integral at $t_0 = (c + d)/2$, we obtain

$$f^{(n)}(t_0) = \frac{1}{2\pi} \int_{-b}^{b} \hat{f}(\omega) \, (i\omega)^n \, \exp(i\omega t_0) \, d\omega = 0. \tag{2.54}$$

Since

$$f(t) = \frac{1}{2\pi} \int_{-b}^{b} \hat{f}(\omega) \exp[i\omega(t - t_0)] \exp(i\omega t_0) \, d\omega, \tag{2.55}$$

developing $\exp[i\omega(t - t_0)]$ as an infinite series yields for all $t \in \mathbb{R}$

$$f(t) = \frac{1}{2\pi} \sum_{n=0}^{+\infty} \frac{[i(t - t_0)]^n}{n!} \int_{-b}^{b} \hat{f}(\omega) \, \omega^n \, \exp(i\omega t_0) \, d\omega = 0. \tag{2.56}$$

This contradicts our assumption that $f \neq 0$.   ∎   ∎

2.3.3   Total Variation

The total variation measures the total amplitude of signal oscillations. It plays an important role in image processing, where its value depends on the length of the image level sets. We show that a low-pass filter can considerably amplify the total variation by creating Gibbs oscillations.

Variations and Oscillations   If f is differentiable, its total variation is defined by

$$\|f\|_V = \int_{-\infty}^{+\infty} |f'(t)| \, dt \,. \tag{2.57}$$

If $\{x_p\}_p$ are the abscissa of the local extrema of f where $f'(x_p) = 0$, then

$$\|f\|_V = \sum_p |f(x_{p+1}) - f(x_p)| \,.$$

It thus measures the total amplitude of the oscillations of f. For example, if $f(t) = \exp(-t^2)$, then $\|f\|_V = 2$. If $f(t) = \sin(\pi t)/(\pi t)$, then f has a local extrema at $x_p \in [p, p+1]$ for any $p \in \mathbb{Z}$. Since $|f(x_{p+1}) - f(x_p)| \sim |p|^{-1}$, we derive that $\|f\|_V = +\infty$.

The total variation of non-differentiable functions can be calculated by considering the derivative in the general sense of distributions [60, 73]. This is equivalent to approximating the derivative by a finite difference on an interval h that goes to zero:

$$\|f\|_V = \lim_{h \to 0} \int_{-\infty}^{+\infty} \frac{|f(t) - f(t - h)|}{|h|} \, dt \,. \tag{2.58}$$

The total variation of discontinuous functions is thus well defined. For example, if $f = \mathbf{1}_{[a,b]}$ then (2.58) gives $\|f\|_V = 2$. We say that f has a *bounded variation* if $\|f\|_V < +\infty$.

Whether f' is the standard derivative of f or its generalized derivative in the sense of distributions, its Fourier transform is $\widehat{f'}(\omega) = i\omega \hat{f}(\omega)$. Hence

$$|\omega| \, |\hat{f}(\omega)| \leqslant \int_{-\infty}^{+\infty} |f'(t)| dt = \|f\|_V \,,$$

which implies that

$$|\hat{f}(\omega)| \leqslant \frac{\|f\|_V}{|\omega|} . \tag{2.59}$$

However, $|\hat{f}(\omega)| = O(|\omega|^{-1})$ is not a sufficient condition to guarantee that f has bounded variation. For example, if $f(t) = \sin(\pi t)/(\pi t)$, then $\hat{f} = \mathbf{1}_{[-\pi,\pi]}$ satisfies $|\hat{f}(\omega)| \leqslant \pi |\omega|^{-1}$ although $\|f\|_V = +\infty$. In general, the total variation of f cannot be evaluated from $|\hat{f}(\omega)|$.

Discrete Signals Let $f_N[n] = f \star \phi_N(n/N)$ be a discrete signal obtained with an averaging filter $\phi_N(t) = \mathbf{1}_{[0,N^{-1}]}(t)$, and a uniform sampling at intervals N^{-1}. The discrete total variation is calculated by approximating the signal derivative by a finite difference over the sampling distance $h = N^{-1}$, and replacing the integral (2.58) by a Riemann sum, which gives:

$$\|f_N\|_V = \sum_n |f_N[n] - f_N[n-1]| . \tag{2.60}$$

If n_p are the abscissa of the local extrema of f_N, then

$$\|f_N\|_V = \sum_p |f_N[n_{p+1}] - f_N[n_p]| .$$

The total variation thus measures the total amplitude of the oscillations of f. In accordance with (2.58), we say that the discrete signal has a *bounded variation* if $\|f_N\|_V$ is bounded by a constant independent of the resolution N.

Gibbs Oscillations Filtering a signal with a low-pass filter can create oscillations that have an infinite total variation. Let $f_\xi = f \star \phi_\xi$ be the filtered signal obtained with an ideal low-pass filter whose transfer function is $\hat{\phi}_\xi = \mathbf{1}_{[-\xi,\xi]}$. If $f \in \mathbf{L}^2(\mathbb{R})$, then f_ξ converges to f in $\mathbf{L}^2(\mathbb{R})$ norm: $\lim_{\xi \to +\infty} \|f - f_\xi\| = 0$. Indeed, $\hat{f}_\xi = \hat{f} \mathbf{1}_{[-\xi,\xi]}$ and the Plancherel formula (2.26) implies that

$$\|f - f_\xi\|^2 = \frac{1}{2\pi} \int_{-\infty}^{+\infty} |\hat{f}(\omega) - \hat{f}_\xi(\omega)|^2 \, d\omega = \frac{1}{2\pi} \int_{|\omega|>\xi} |\hat{f}(\omega)|^2 \, d\omega,$$

which goes to zero as ξ increases. However, if f is discontinuous in t_0, then we show that f_ξ has Gibbs oscillations in the neighborhood of t_0, which prevents $\sup_{t \in \mathbb{R}} |f(t) - f_\xi(t)|$ from converging to zero as ξ increases.

Let f be a bounded variation function $\|f\|_V < +\infty$ that has an isolated discontinuity at t_0, with a left limit $f(t_0^-)$ and right limit $f(t_0^+)$. It is decomposed as a sum of f_c, which is continuous in the neighborhood of t_0, plus a Heaviside step of amplitude $f(t_0^+) - f(t_0^-)$:

$$f(t) = f_c(t) + [f(t_0^+) - f(t_0^-)] \, u(t - t_0),$$

with

$$u(t) = \begin{cases} 1 & \text{if } t \geqslant 0 \\ 0 & \text{otherwise} \end{cases} . \tag{2.61}$$

Hence

$$f_\xi(t) = f_c \star \phi_\xi(t) + [f(t_0^+) - f(t_0^-)] \, u \star \phi_\xi(t - t_0). \tag{2.62}$$

Since f_c has bounded variation and is uniformly continuous in the neighborhood of t_0, one can prove (Exercise 2.15) that $f_c \star \phi_\xi(t)$ converges uniformly to $f_c(t)$ in a neighborhood of t_0. The following theorem shows that this is not true for $u \star \phi_\xi$, which creates Gibbs oscillations.

Theorem 2.8 (Gibbs). *For any $\xi > 0$,*

$$u \star \phi_\xi(t) = \int_{-\infty}^{\xi t} \frac{\sin x}{\pi x} \, dx. \tag{2.63}$$

Proof. The impulse response of an ideal low-pass filter, calculated in (2.29), is $\phi_\xi(t) = \sin(\xi t)/(\pi t)$. Hence

$$u \star \phi_\xi(t) = \int_{-\infty}^{+\infty} u(\tau) \frac{\sin \xi(t-\tau)}{\pi(t-\tau)} \, d\tau = \int_0^{+\infty} \frac{\sin \xi(t-\tau)}{\pi(t-\tau)} \, d\tau .$$

The change of variable $x = \xi(t-\tau)$ gives (2.63).   ■   ■
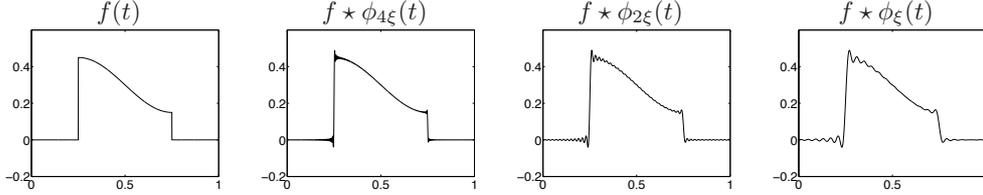
Figure 2.1: Gibbs oscillations created by low-pass filters with cut-off frequencies that decrease from left to right.

The function

$$s(\xi t) = \int_{-\infty}^{\xi t} \frac{\sin x}{\pi x} \, dx$$

is a sigmoid that increases from 0 at $t = -\infty$ to 1 at $t = +\infty$, with $s(0) = 1/2$. It has oscillations of period π/ξ, which are attenuated when the distance to 0 increases, but their total variation is infinite: $\|s\|_V = +\infty$. The maximum amplitude of the Gibbs oscillations occurs at $t = \pm \pi/\xi$, with an amplitude independent of ξ:

$$A = s(\pi) - 1 = \int_{-\infty}^{\pi} \frac{\sin x}{\pi x} dx - 1 \approx 0.045 .$$

Inserting (2.63) in (2.62) shows that

$$f(t) - f_\xi(t) = [f(t_0^+) - f(t_0^-)] \, s(\xi(t-t_0)) + \varepsilon(\xi, t) , \tag{2.64}$$

where $\lim_{\xi \to +\infty} \sup_{|t-t_0|<\alpha} |\varepsilon(\xi, t)| = 0$ in some neighborhood of size $\alpha > 0$ around t_0. The sigmoid $s(\xi(t-t_0))$ centered at t_0 creates a maximum error of fixed amplitude for all ξ. This is seen in Figure 2.1, where the Gibbs oscillations have an amplitude proportional to the jump $f(t_0^+) - f(t_0^-)$ at all frequencies ξ.

Image Total Variation   The total variation of an image $f(x_1, x_2)$ depends on the amplitude of its variations as well as the length of the contours along which they occur. Suppose that $f(x_1, x_2)$ is differentiable. The total variation is defined by

$$\|f\|_V = \int \int |\vec{\nabla} f(x_1, x_2)| \, dx_1 \, dx_2 , \tag{2.65}$$

where the modulus of the gradient vector is

$$|\vec{\nabla} f(x_1, x_2)| = \left(\left| \frac{\partial f(x_1, x_2)}{\partial x_1} \right|^2 + \left| \frac{\partial f(x_1, x_2)}{\partial x_2} \right|^2 \right)^{1/2} .$$

As in one dimension, the total variation is extended to discontinuous functions by taking the derivatives in the general sense of distributions. An equivalent norm is obtained by approximating the partial derivatives by finite differences:

$$|\Delta_h f(x_1, x_2)| = \left(\left| \frac{f(x_1, x_2) - f(x_1 - h, x_2)}{h} \right|^2 + \right.$$

$$\left. \left| \frac{f(x_1, x_2) - f(x_1, x_2 - h)}{h} \right|^2 \right)^{1/2} .$$

One can verify that

$$\|f\|_V \leqslant \lim_{h \to 0} \int \int |\Delta_h f(x_1, x_2)| \, dx_1 \, dx_2 \leqslant \sqrt{2} \, \|f\|_V. \tag{2.66}$$

The finite difference integral gives a larger value when $f(x_1, x_2)$ is discontinuous along a diagonal line in the (x_1, x_2) plane.

The total variation of f is related to the length of it level sets. Let us define

$$\Omega_y = \{(x_1, x_2) \in \mathbb{R}^2 \ : \ f(x_1, x_2) > y\} .$$

If f is continuous then the boundary $\partial \Omega_y$ of Ω_y is the level set of all (x_1, x_2) such that $f(x_1, x_2) = y$. Let $H^1(\partial \Omega_y)$ be the length of $\partial \Omega_y$. Formally, this length is calculated in the sense of the mono-dimensional Hausdorff measure. The following theorem relates the total variation of f to the length of its level sets.

Theorem 2.9 (Co-area Formula). *If $\|f\|_V < +\infty$ then*

$$\|f\|_V = \int_{-\infty}^{+\infty} H^1(\partial \Omega_y) \, dy. \tag{2.67}$$

Proof. The proof is a highly technical result that is given in [73]. We give an intuitive explanation when f is continuously differentiable. In this case $\partial \Omega_y$ is a differentiable curve $x(y, s) \in \mathbb{R}^2$, which is parameterized by the arc-length s. Let $\vec{\tau}(x)$ be the vector tangent to this curve in the plane. The gradient $\vec{\nabla} f(x)$ is orthogonal to $\vec{\tau}(x)$. The Frenet coordinate system along $\partial \Omega_y$ is composed of $\vec{\tau}(x)$ and of the unit vector $\vec{n}(x)$ parallel to $\vec{\nabla} f(x)$. Let ds and dn be the Lebesgue measures in the direction of $\vec{\tau}$ and $\vec{n}$. We have

$$|\vec{\nabla} f(x)| = \vec{\nabla} f(x) \cdot \vec{n} = \frac{dy}{dn} , \tag{2.68}$$

where dy is the differential of amplitudes across level sets. The idea of the proof is to decompose the total variation integral over the plane as an integral along the level sets and across level sets, which we write:

$$\|f\|_V = \int \int |\vec{\nabla} f(x_1, x_2)| \, dx_1 \, dx_2 = \int \int_{\partial \Omega_y} |\vec{\nabla} f(x(y, s))| \, ds \, dn. \tag{2.69}$$

By using (2.68) we can get

$$\|f\|_V = \int \int_{\partial \Omega_y} ds \, dy .$$

But $\int_{\partial \Omega_y} ds = H^1(\partial \Omega_y)$ is the length of the level set, which justifies (2.67). ∎ ∎

The co-area formula gives an important geometrical interpretation of the total image variation. Images are uniformly bounded so the integral (2.67) is calculated over a finite interval and is proportional to the average length of level sets. It is finite as long as the level sets are not fractal curves. Let $f = \alpha \mathbf{1}_\Omega$ be proportional to the indicator function of a set $\Omega \subset \mathbb{R}^2$ which has a boundary $\partial \Omega$ of length L. The co-area formula (2.9) implies that $\|f\|_V = \alpha \, L$. In general, bounded variation images must have step edges of finite length.

Discrete Images A camera measures light intensity with photoreceptors that perform an averaging and a uniform sampling over a grid that is supposed to be uniform. For a resolution N, the sampling interval is N^{-1}. The resulting image can be written $f_N[n_1, n_2] = f \star \phi_N(n_1/N, n_2/N)$, where $\phi_N = \mathbf{1}_{[0, N^{-1}]^2}$ and f is the averaged analog image. Its total variation is defined by approximating derivatives by finite differences and the integral (2.66) by a Riemann sum:

$$\|f_N\|_V = \frac{1}{N} \sum_{n_1} \sum_{n_2} \left(\ \left| f_N[n_1, n_2] - f_N[n_1 - 1, n_2] \right|^2 + \right. \tag{2.70}$$

$$\left. \left| f_N[n_1, n_2] - f_N[n_1, n_2 - 1] \right|^2 \right)^{1/2} .$$

In accordance with (2.66) we say that the image has bounded variation if $\|f_N\|_V$ is bounded by a constant independent of the resolution N. The co-area formula proves that it depends on the

(a)   (b)

Figure 2.2: (a): The total variation of this image remains nearly constant when the resolution N increases. (b): Level sets $\partial\Omega_y$ obtained by sampling uniformly the amplitude variable y.

length of the level sets as the image resolution increases. The $\sqrt{2}$ upper bound factor in (2.66) comes from the fact that the length of a diagonal line can be increased by $\sqrt{2}$ if it is approximated by a zig-zag line that remains on the horizontal and vertical segments of the image sampling grid. Figure 2.2(a) shows a bounded variation image and Figure 2.2(b) displays the level sets obtained by discretizing uniformly the amplitude variable y. The total variation of this image remains nearly constant as the resolution varies.

2.4   Two-Dimensional Fourier Transform

The Fourier transform in $\mathbb{R}^n$ is a straightforward extension of the one-dimensional Fourier transform. The two-dimensional case is briefly reviewed for image processing applications. The Fourier transform of a two-dimensional integrable function $f \in \mathbf{L}^1(\mathbb{R}^2)$ is

$$\hat{f}(\omega_1, \omega_2) = \int_{-\infty}^{+\infty} \int_{-\infty}^{+\infty} f(x_1, x_2) \exp[-i(\omega_1 x_1 + \omega_2 x_2)] \, dx_1 \, dx_2. \qquad (2.71)$$

In polar coordinates $\exp[i(\omega_1 x + \omega_2 y)]$ can be rewritten

$$\exp[i(\omega_1 x_1 + \omega_2 x_2)] = \exp[i\xi(x_1 \cos\theta + x_2 \sin\theta)]$$

with $\xi = \sqrt{\omega_1^2 + \omega_2^2}$. It is a plane wave that propagates in the direction of θ and oscillates at the frequency ξ. The properties of a two-dimensional Fourier transform are essentially the same as in one dimension. We summarize a few important results. We write $\omega = (\omega_1, \omega_2)$, $x = (x_1, x_2)$, $\omega \cdot x = \omega_1 x_1 + \omega_2 x_2$ and $\iint f(x_1, x_2) dx_1 \, dx_2 = \iint f(x) \, dx$.

- If $f \in \mathbf{L}^1(\mathbb{R}^2)$ and $\hat{f} \in \mathbf{L}^1(\mathbb{R}^2)$ then

$$f(x) = \frac{1}{4\pi^2} \iint \hat{f}(\omega) \exp[i(\omega \cdot x)] \, d\omega. \qquad (2.72)$$

- If $f \in \mathbf{L}^1(\mathbb{R}^2)$ and $h \in \mathbf{L}^1(\mathbb{R}^2)$ then the convolution

$$g(x) = f \star h(x) = \iint f(u) \, h(x - u) \, du$$

  has a Fourier transform

$$\hat{g}(\omega) = \hat{f}(\omega) \, \hat{h}(\omega). \qquad (2.73)$$

- The Parseval formula proves that

$$\iint f(x) \, g^*(x) \, dx = \frac{1}{4\pi^2} \iint \hat{f}(\omega) \, \hat{g}^*(\omega) \, d\omega \ . \qquad (2.74)$$

If $f = g$, we obtain the Plancherel equality

$$\iint |f(x)|^2 \, dx = \frac{1}{4\pi^2} \iint |\hat{f}(\omega)|^2 \, d\omega \, . \tag{2.75}$$

The Fourier transform of a finite energy function thus has finite energy. With the same density based argument as in one dimension, energy equivalence makes it possible to extend the Fourier transform to any function $f \in \mathbf{L}^2(\mathbb{R}^2)$.

- If $f \in \mathbf{L}^2(\mathbb{R}^2)$ is separable, which means that

$$f(x) = f(x_1, x_2) = g(x_1) \, h(x_2),$$

then its Fourier transform is

$$\hat{f}(\omega) = \hat{f}(\omega_1, \omega_2) = \hat{g}(\omega_1) \, \hat{h}(\omega_2),$$

where $\hat{h}$ and $\hat{g}$ are the one-dimensional Fourier transforms of g and h. For example, the indicator function

$$f(x_1, x_2) = \left\{ \begin{array}{ll} 1 & \text{if } |x_1| \leqslant T \, , \ |x_2| \leqslant T \\ 0 & \text{otherwise} \end{array} \right. = \mathbf{1}_{[-T,T]}(x_1) \times \mathbf{1}_{[-T,T]}(x_2)$$

is a separable function whose Fourier transform is derived from (2.28):

$$\hat{f}(\omega_1, \omega_2) = \frac{4 \sin(T\omega_1) \, \sin(T\omega_2)}{\omega_1 \, \omega_2}.$$

- If $f(x_1, x_2)$ is rotated by θ:

$$f_\theta(x_1, x_2) = f(x_1 \cos\theta - x_2 \sin\theta, x_1 \sin\theta + x_2 \cos\theta),$$

then its Fourier transform is rotated by θ:

$$\hat{f}_\theta(\omega_1, \omega_2) = \hat{f}(\omega_1 \cos\theta - \omega_2 \sin\theta, \omega_1 \sin\theta + \omega_2 \cos\theta). \tag{2.76}$$

Radon Transform A Radon transform computes integrals of $f \in \mathbf{L}^2(\mathbb{R}^2)$ along rays. It is provides a good model for some tomographic systems such as X-ray measurements in medical imaging. Inverting the Radon transform is then needed to reconstruct the 2D or 3D body from these integrals.

Let us write $\tau_\theta = (\cos\theta, \sin\theta)$. A ray $\Delta_{t,\theta}$ is a line defined by its equation

$$x \cdot \tau_\theta = x_1 \cos\theta + x_2 \sin\theta = t \, .$$

The projection p_θ of f along a parallel line of orientation θ is defined by

$$\forall \theta \in [0, \pi), \forall t \in \mathbb{R}, \quad p_\theta(t) = \int_{\Delta_{t,\theta}} f(x) \, ds = \iint f(x) \, \delta(x \cdot \tau_\theta - t) \, dx, \tag{2.77}$$

where δ is the Dirac distribution. The Radon transform maps $f(x)$ to $p_\theta(t)$ for $\theta \in [0, \pi)$. In medical imaging applications, a scanner is rotated around an object to compute the projection p_θ for many angles $\theta \in [0, \pi)$, as illustrated in Figure 2.3. The Fourier slice theorem relates the Fourier transform of p_θ to slices of the Fourier transform of f.

Theorem 2.10 (Fourier slice). *The Fourier transform of projections satisfies*

$$\forall \theta \in [0, \pi) \, , \ \forall \xi \in \mathbb{R} \quad \hat{p}_\theta(\xi) = \hat{f}(\xi \cos\theta, \xi \sin\theta).$$

Proof. The Fourier transform of the projection is

$$\hat{p}_\theta(\xi) = \int_{-\infty}^{+\infty} \left(\iint f(x) \, \delta(x \cdot \tau_\theta - t) \, dx \right) e^{-it\xi} dt$$

$$= \iint f(x) \exp\left(-i(x \cdot \tau_\theta)\xi\right) dx = \hat{f}(\xi\tau_\theta).$$

∎

An image f can be recovered from its projections p_θ thanks to the projection slice theorem. Indeed, the Fourier transform $\hat{f}$ is known along each ray of direction θ and f is thus obtained with the 2D inverse Fourier transform (2.71). The following back projection theorem gives an inversion formula.

Theorem 2.11 (Back projection). *The image f is recovered using a one-dimensional filter $h(t)$:*

$$f(x) = \frac{1}{2\pi} \int_0^\pi p_\theta * h(x \cdot \tau_\theta)\, d\theta \quad with \quad \hat{h}(\xi) = |\xi| \ .$$

Proof. The inverse Fourier transform (2.72) in polar coordinates $(\omega_1, \omega_2) = (\xi\cos\theta, \xi\sin\theta)$, with $d\omega_1\, d\omega_2 = \xi\, d\theta\, d\xi$, can be written

$$f(x) = \frac{1}{4\pi^2} \int_0^{+\infty} \int_0^{2\pi} \hat{f}(\xi\cos\theta, \xi\sin\theta) \exp\left(i(x \cdot \tau_\theta)\xi\right) \xi\, d\theta\, d\xi \ .$$

Using the Fourier slice Theorem 2.10 with $p_{\theta+\pi}(t) = p_\theta(-t)$, this is rewritten as

$$f(x) = \frac{1}{2\pi} \int_0^\pi \left(\frac{1}{2\pi} \int_{-\infty}^{+\infty} |\xi|\, \hat{p}_\theta(\xi) \exp\left(i(x \cdot \tau_\theta)\xi\right) d\xi \right) d\theta \ .$$

The inner integral is the inverse Fourier transform of $\hat{p}_\theta(\xi)\, |\xi|$ evaluated at $x \cdot \tau_\theta \in \mathbb{R}$. The convolution formula (2.73) shows that it is equal to $p_\theta * h(x \cdot \tau_\theta)$. ∎
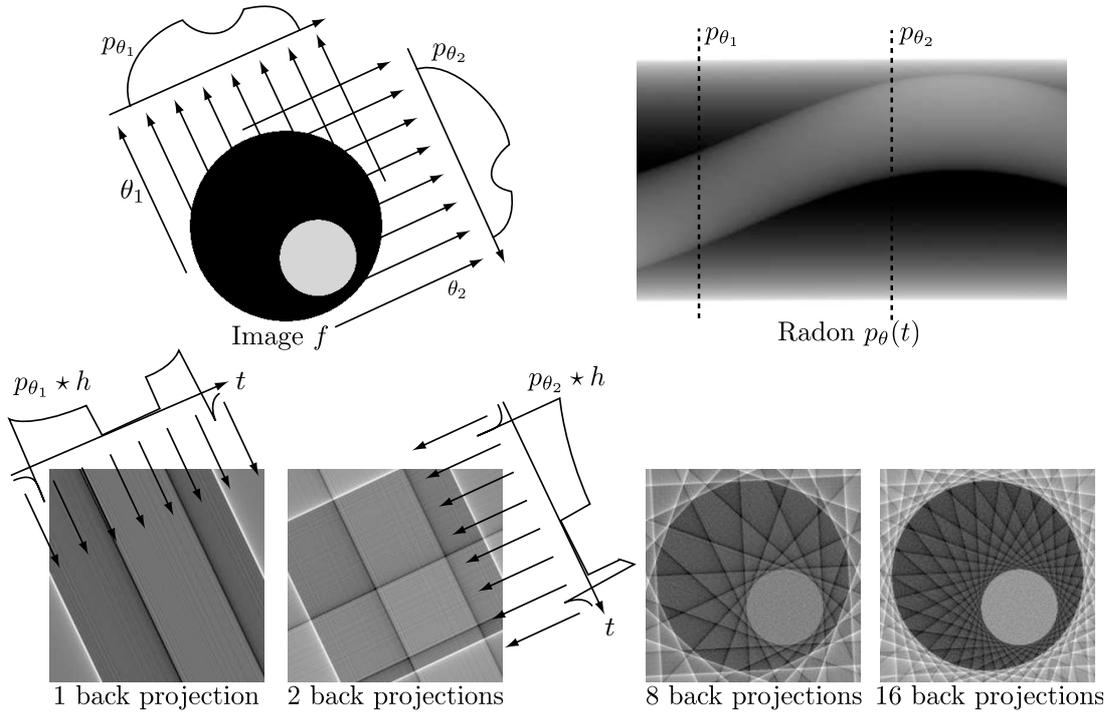
Figure 2.3: The Radon transform and its reconstruction with an increasing number of back projections.

In medical imaging applications, only a limited number of projections is available, and the Fourier transform $\hat{f}$ is thus partially known. In this case, an approximation of f can still be recovered by summing the corresponding filtered back projections $p_\theta * h(x \cdot \tau_\theta)$. Figure 2.3 describes this process, and shows the reconstruction of an image with a geometric object, using an increasing number of evenly spaced projections. Section 13.3 describes a non-linear super-resolution reconstruction algorithm, that recovers a more precise image by using a sparse representation.

2.5 Exercises

2.1. [1] Prove that if $f \in \mathbf{L}^1(\mathbb{R})$ then $\hat{f}(\omega)$ is a continuous function of ω, and if $\hat{f} \in \mathbf{L}^1(\mathbb{R})$ then $f(t)$ is continuous.

2.2. [1] Prove that a filter whose impulse response is a function $h(t)$ can be stable only if $\int |h(t)| \, dt < \infty$.

2.3. [1] Prove the translation (2.18), scaling (2.20) and time derivative (2.21) properties of the Fourier transform.

2.4. [1] Let $f_r(t) = \mathrm{Re}[f(t)]$ and $f_i(t) = \mathrm{Ima}[f(t)]$ be the real and imaginary parts of $f(t)$. Prove that $\hat{f}_r(\omega) = [\hat{f}(\omega) + \hat{f}^*(-\omega)]/2$ and $\hat{f}_i(\omega) = [\hat{f}(\omega) - \hat{f}^*(-\omega)]/(2i)$.

2.5. [1] Prove that if $\hat{f}(\omega)$ is differentiable and $\hat{f}(0) = \hat{f}'(0) = 0$ then

$$\int_{-\infty}^{+\infty} f(t) \, dt = \int_{-\infty}^{+\infty} t \, f(t) \, dt = 0.$$

2.6. [1] By using the Fourier transform, verify that

$$\int_{-\infty}^{+\infty} \frac{\sin(\pi t)}{(\pi t)} \, dt = 1 \quad \text{and} \quad \int_{-\infty}^{+\infty} \frac{\sin^2 t}{t^3} \, dt = \frac{3\pi}{4}.$$

2.7. [2] Show that the Fourier transform of $f(t) = \exp(-(a - ib)t^2)$ is

$$\hat{f}(\omega) = \sqrt{\frac{\pi}{a - ib}} \, \exp\left(-\frac{a + ib}{4(a^2 + b^2)} \, \omega^2\right).$$

Hint: write a differential equation similar to (2.33).

2.8. [3] *Riemann-Lebesgue* Prove that if $f \in \mathbf{L}^1(\mathbb{R})$ then $\lim_{\omega \to \infty} \hat{f}(\omega) = 0$.

Hint: Prove it first for $\mathbf{C}^1$ functions with a compact support and use a density argument.

2.9. [2] *Stability of passive circuits*

(a) Let p be a complex number with $\mathrm{Re}[p] < 0$. Compute the Fourier transforms of $f(t) = \exp(pt) \, \mathbf{1}_{[0,+\infty)}(t)$ and of $f(t) = t^n \, \exp(pt) \, \mathbf{1}_{[0,+\infty)}(t)$.

(b) A passive circuit relates the input voltage f to the output voltage g by a differential equation with constant coefficients:

$$\sum_{k=0}^{K} a_k \, f^{(k)}(t) = \sum_{k=0}^{M} b_k \, g^{(k)}(t).$$

Prove that this system is stable and causal if and only if the roots of the equation $\sum_{k=0}^{M} b_k \, z^k = 0$ have a strictly negative real part.

(c) A Butterworth filter satisfies

$$|\hat{h}(\omega)|^2 = \frac{1}{1 + (\omega/\omega_0)^{2N}}.$$

For $N = 3$, compute $\hat{h}(\omega)$ and $h(t)$ so that this filter can be implemented by a stable electronic circuit.

2.10. [1] For any $A > 0$, construct f such that the time and frequency spread measured respectively by σ_t and σ_ω in (2.46) and (2.45) satisfy $\sigma_t > A$ and $\sigma_\omega > A$.

2.11. [3] Suppose that $f(t) \geqslant 0$ and that its support is in $[-T, T]$. Verify that $|\hat{f}(\omega)| \leqslant \hat{f}(0)$. Let ω_c be the half-power point defined by $|\hat{f}(\omega_c)|^2 = |\hat{f}(0)|^2/2$ and $|\hat{f}(\omega)|^2 < |\hat{f}(0)|^2/2$ for $\omega < \omega_c$. Prove that $\omega_c T \geqslant \pi/2$.

2.12. [2] *Rectification* A rectifier computes $g(t) = |f(t)|$, for recovering the envelope of modulated signals [53].

(a) Show that if $f(t) = a(t) \sin \omega_0 t$ with $a(t) \geqslant 0$ then $g(t) = |f(t)|$ satisfies

$$\hat{g}(\omega) = -\frac{2}{\pi} \sum_{n=-\infty}^{+\infty} \frac{\hat{a}(\omega - 2n\omega_0)}{4n^2 - 1}.$$

(b) Suppose that $\hat{a}(\omega) = 0$ for $|\omega| > \omega_0$. Find h such that $a(t) = h \star g(t)$.

2.13. **²** *Amplitude modulation* For $0 \leqslant n < N$, we suppose that $f_n(t)$ is real and that $\hat{f}_n(\omega) = 0$ for $|\omega| > \omega_0$. An amplitude modulated multiplexed signal is defined by

$$g(t) = \sum_{n=0}^{N} f_n(t) \cos(2 \, n \, \omega_0 \, t).$$

Compute $\hat{g}(\omega)$ and verify that the width of its support is $4N\omega_0$. Find a demodulation algorithm that recovers each f_n from g.

2.14. **¹** Show that $\|\phi\|_V = +\infty$ if $\phi(t) = \sin(\pi t)/(\pi t)$. Show that $\|\phi\|_V = 2\,\lambda$ if $\phi(t) = \lambda\,\mathbf{1}_{[a,b]}(t)$.

2.15. **³** Let $f_\xi = f \star h_\xi$ with $\hat{h}_\xi = \mathbf{1}_{[-\xi,\xi]}$. Suppose that f has a bounded variation $\|f\|_V < +\infty$ and that it is continuous in a neighborhood of t_0. Prove that in a neighborhood of t_0, $f_\xi(t)$ converges uniformly to $f(t)$ when ξ goes to $+\infty$.

2.16. **¹** Compute the two-dimensional Fourier transforms of $f(x) = \mathbf{1}_{[0,1]^2}(x_1, x_2)$ and of $f(x) = e^{-(x_1^2 + x_2^2)}$.

2.17. **¹** Compute Radon transform of the indicator function of the unit circle: $f(x_1, x_2) = \mathbf{1}_{x^1 + x^2 \leqslant 1}$.

2.18. **²** Let $f(x_1, x_2)$ be an image which has a discontinuity of amplitude A along a straight line having an angle θ in the plane (x_1, x_2). Compute the amplitude of the Gibbs oscillations of $f \star h_\xi(x_1, x_2)$ as a function of ξ, θ and A, for $\hat{h}_\xi(\omega_1, \omega_2) = \mathbf{1}_{[-\xi,\xi]}(\omega_1)\,\mathbf{1}_{[-\xi,\xi]}(\omega_2)$.

III

Discrete Revolution

Digital signal processing has taken over. First used in the 1950's at the service of analog signal processing to simulate analog transforms, digital algorithms have invaded most traditional fortresses, including phones, music recording, cameras, televisions, and all information processing. Analog computations performed with electronic circuits are faster than digital algorithms implemented with microprocessors, but are less precise and less flexible. Thus analog circuits are often replaced by digital chips once the computational performance of microprocessors is sufficient to operate in real time for a given application.

Whether sound recordings or images, most discrete signals are obtained by sampling an analog signal. An analog to digital conversion is a linear approximation which introduces an error which depends upon the sampling rate. Once more, the Fourier transform is unavoidable because the eigenvectors of discrete time-invariant operators are sinusoidal waves. The Fourier transform is discretized for signals of finite size and implemented with a Fast Fourier Transform algorithm.

3.1 Sampling Analog Signals

The simplest way to discretize an analog signal f is to record its sample values $\{f(ns)\}_{n\in\mathbb{Z}}$ at intervals s. An approximation of $f(t)$ at any $t \in \mathbb{R}$ may be recovered by interpolating these samples. The Shannon-Whittaker sampling theorem gives a sufficient condition on the support of the Fourier transform $\hat{f}$ to recover $f(t)$ exactly. Aliasing and approximation errors are studied when this condition is not satisfied.

Digital acquisition devices often do not satisfy the restrictive hypothesis of the Shannon-Whittaker sampling theorem. General linear analog to discrete conversion is introduced in Section 3.1.3, showing that a stable uniform discretization is a linear approximation. A digital conversion also approximates discrete coefficients with a given precision, to store them with a limited number of bits. This quantization aspect is studied in Chapter 10.

3.1.1 Shannon-Whittaker Sampling Theorem

Sampling is first studied from the more classical Shannon-Whittaker point of view, which tries to recover $f(t)$ from its samples $\{f(ns)\}_{n\in\mathbb{Z}}$. A discrete signal may be represented as a sum of Diracs. We associate to any sample $f(ns)$ a Dirac $f(ns)\delta(t - ns)$ located at $t = ns$. A uniform sampling of f thus corresponds to the weighted Dirac sum

$$f_d(t) = \sum_{n=-\infty}^{+\infty} f(ns)\,\delta(t - ns). \tag{3.1}$$

The Fourier transform of $\delta(t - ns)$ is $\mathrm{e}^{-ins\omega}$ so the Fourier transform of f_d is a Fourier series:

$$\hat{f}_d(\omega) = \sum_{n=-\infty}^{+\infty} f(ns)\,\mathrm{e}^{-ins\omega}. \tag{3.2}$$

To understand how to compute $f(t)$ from the sample values $f(ns)$ and hence f from f_d, we relate their Fourier transforms $\hat{f}$ and $\hat{f}_d$.

Theorem 3.1. *The Fourier transform of the discrete signal obtained by sampling f at intervals s is*

$$\hat{f}_d(\omega) = \frac{1}{s} \sum_{k=-\infty}^{+\infty} \hat{f}\left(\omega - \frac{2k\pi}{s}\right). \tag{3.3}$$

Proof. Since $\delta(t - ns)$ is zero outside $t = ns$,

$$f(ns)\,\delta(t - ns) = f(t)\,\delta(t - ns),$$

so we can rewrite (3.1) as multiplication with a Dirac comb:

$$f_d(t) = f(t) \sum_{n=-\infty}^{+\infty} \delta(t - ns) = f(t)\,c(t). \tag{3.4}$$

Computing the Fourier transform yields

$$\hat{f}_d(\omega) = \frac{1}{2\pi}\,\hat{f} \star \hat{c}(\omega). \tag{3.5}$$

The Poisson formula (2.4) proves that

$$\hat{c}(\omega) = \frac{2\pi}{s} \sum_{k=-\infty}^{+\infty} \delta\left(\omega - \frac{2\pi k}{s}\right). \tag{3.6}$$

Since $\hat{f} \star \delta(\omega - \xi) = \hat{f}(\omega - \xi)$, inserting (3.6) in (3.5) proves (3.3). ■   ■

Theorem 3.1 proves that sampling f at intervals s is equivalent to making its Fourier transform $2\pi/s$ periodic by summing all its translations $\hat{f}(\omega - 2k\pi/s)$. The resulting sampling theorem was first proved by Whittaker [482] in 1935 in a book on interpolation theory. Shannon rediscovered it in 1949 for applications to communication theory [428].

Theorem 3.2 (Shannon, Whittaker)**.** *If the support of $\hat{f}$ is included in $[-\pi/s, \pi/s]$ then*

$$f(t) = \sum_{n=-\infty}^{+\infty} f(ns)\,\phi_s(t - ns), \tag{3.7}$$

with

$$\phi_s(t) = \frac{\sin(\pi t/s)}{\pi t/s}. \tag{3.8}$$

Proof. If $n \neq 0$, the support of $\hat{f}(\omega - n\pi/s)$ does not intersect the support of $\hat{f}(\omega)$ because $\hat{f}(\omega) = 0$ for $|\omega| > \pi/s$. So (3.3) implies

$$\hat{f}_d(\omega) = \frac{\hat{f}(\omega)}{s} \quad \text{if} \ \ |\omega| \leqslant \frac{\pi}{s}. \tag{3.9}$$

The Fourier transform of ϕ_s is $\hat{\phi}_s = s\,\mathbf{1}_{[-\pi/s,\pi/s]}$. Since the support of $\hat{f}$ is in $[-\pi/s, \pi/s]$ it results from (3.9) that $\hat{f}(\omega) = \hat{\phi}_s(\omega)\,\hat{f}_d(\omega)$. The inverse Fourier transform of this equality gives

$$
\begin{aligned}
f(t) = \phi_s \star f_d(t) &= \phi_s \star \sum_{n=-\infty}^{+\infty} f(ns)\,\delta(t - ns) \\
&= \sum_{n=-\infty}^{+\infty} f(ns)\,\phi_s(t - ns).
\end{aligned}
$$

■   ■

The sampling theorem imposes that the support of $\hat{f}$ is included in $[-\pi/s, \pi/s]$, which guarantees that f has no brutal variations between consecutive samples, and can thus be recovered with a smooth interpolation. Section 3.1.3 shows that one can impose other smoothness conditions to recover f from its samples. Figure 3.1 illustrates the different steps of a sampling and reconstruction from samples, in both the time and Fourier domains.

3.1.2 Aliasing

The sampling interval s is often imposed by computation or storage constraints and the support of $\hat{f}$ is generally not included in $[-\pi/s, \pi/s]$. In this case the interpolation formula (3.7) does not recover f. We analyze the resulting error and a filtering procedure to reduce it.
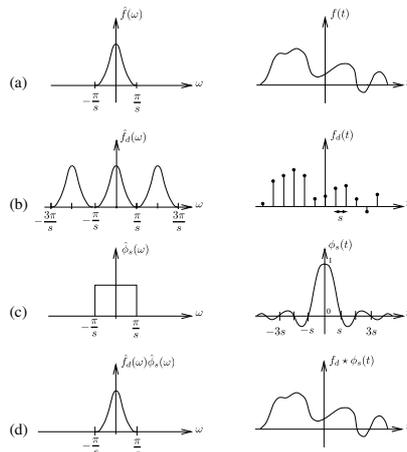
Figure 3.1: (a): Signal f and its Fourier transform $\hat{f}$. (b): A uniform sampling of f makes its Fourier transform periodic. (c): Ideal low-pass filter. (d): The filtering of (b) with (c) recovers f.

Theorem 3.1 proves that

$$\hat{f}_d(\omega) = \frac{1}{s} \sum_{k=-\infty}^{+\infty} \hat{f}\left(\omega - \frac{2k\pi}{s}\right). \tag{3.10}$$

Suppose that the support of $\hat{f}$ goes beyond $[-\pi/s, \pi/s]$. In general the support of $\hat{f}(\omega - 2k\pi/s)$ intersects $[-\pi/s, \pi/s]$ for several $k \neq 0$, as shown in Figure 3.2. This folding of high frequency components over a low frequency interval is called *aliasing*. In the presence of aliasing, the interpolated signal

$$\phi_s \star f_d(t) = \sum_{n=-\infty}^{+\infty} f(ns)\, \phi_s(t - ns)$$

has a Fourier transform

$$\hat{f}_d(\omega)\, \hat{\phi}_s(\omega) = s\, \hat{f}_d(\omega)\, \mathbf{1}_{[-\pi/s, \pi/s]}(\omega) = \mathbf{1}_{[-\pi/s, \pi/s]}(\omega) \sum_{k=-\infty}^{+\infty} \hat{f}\left(\omega - \frac{2k\pi}{s}\right) \tag{3.11}$$

which may be completely different from $\hat{f}(\omega)$ over $[-\pi/s, \pi/s]$. The signal $\phi_s \star f_d$ may not even be a good approximation of f, as shown by Figure 3.2.

Example 3.1. *Let us consider a high frequency oscillation*

$$f(t) = \cos(\omega_0 t) = \frac{e^{i\omega_0 t} + e^{-i\omega_0 t}}{2}.$$

Its Fourier transform is

$$\hat{f}(\omega) = \pi\Big(\delta(\omega - \omega_0) + \delta(\omega + \omega_0)\Big).$$

If $2\pi/s > \omega_0 > \pi/s$ then (3.11) yields

$$\hat{f}_d(\omega)\, \hat{\phi}_s(\omega)$$

$$= \pi\, \mathbf{1}_{[-\pi/s, \pi/s]}(\omega) \sum_{k=-\infty}^{+\infty} \left(\delta\Big(\omega - \omega_0 - \frac{2k\pi}{s}\Big) + \delta\Big(\omega + \omega_0 - \frac{2k\pi}{s}\Big)\right)$$

$$= \pi\Big(\delta(\omega - \frac{2\pi}{s} + \omega_0) + \delta(\omega + \frac{2\pi}{s} - \omega_0)\Big),$$
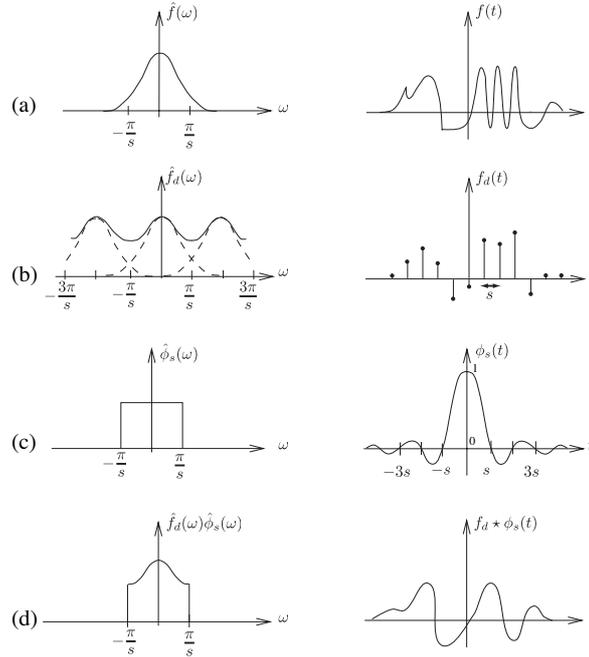
Figure 3.2: (a): Signal f and its Fourier transform $\hat{f}$. (b): Aliasing produced by an overlapping of $\hat{f}(\omega - 2k\pi/s)$ for different k, shown in dashed lines. (c): Ideal low-pass filter. (d): The filtering of (b) with (c) creates a low-frequency signal that is different from f.

so

$$f_d \star \phi_s(t) = \cos\left[\left(\frac{2\pi}{s} - \omega_0\right)t\right].$$

The aliasing reduces the high frequency ω_0 to a lower frequency $2\pi/s - \omega_0 \in [-\pi/s, \pi/s]$. The same frequency folding is observed in a film that samples a fast moving object without enough images per second. A wheel turning rapidly appears as turning much more slowly in the film.

Removal of Aliasing   To apply the sampling theorem, f is approximated by the closest signal $\tilde{f}$ whose Fourier transform has a support in $[-\pi/s, \pi/s]$. The Plancherel formula (2.26) proves that

$$
\begin{aligned}
\|f - \tilde{f}\|^2 &= \frac{1}{2\pi} \int_{-\infty}^{+\infty} |\hat{f}(\omega) - \widehat{\tilde{f}}(\omega)|^2 \, d\omega \\
&= \frac{1}{2\pi} \int_{|\omega|>\pi/s} |\hat{f}(\omega)|^2 \, d\omega + \frac{1}{2\pi} \int_{|\omega|\leqslant\pi/s} |\hat{f}(\omega) - \widehat{\tilde{f}}(\omega)|^2 \, d\omega.
\end{aligned}
$$

This distance is minimum when the second integral is zero and hence

$$\widehat{\tilde{f}}(\omega) = \hat{f}(\omega)\, \mathbf{1}_{[-\pi/s,\pi/s]}(\omega) = \frac{1}{s}\, \hat{\phi}_s(\omega)\, \hat{f}(\omega). \tag{3.12}$$

It corresponds to $\tilde{f} = \frac{1}{s} f \star \phi_s$. The filtering of f by ϕ_s avoids the aliasing by removing any frequency larger than π/s. Since $\widehat{\tilde{f}}$ has a support in $[-\pi/s, \pi/s]$, the sampling theorem proves that $\tilde{f}(t)$ can be recovered from the samples $\tilde{f}(ns)$. An analog to digital converter is therefore composed of a filter that limits the frequency band to $[-\pi/s, \pi/s]$, followed by a uniform sampling at intervals s.

3.1.3   General Sampling and Linear Analog Conversions

The Shannon-Whittaker theorem is a particular example of linear discrete to analog conversion, which does not apply to all digital acquisition devices. This section describes general analog to

discrete conversion and reverse discrete to analog conversion, with a general linear filtering and uniform sampling. Analog signals are approximated by linear projections on approximation spaces.

Sampling Theorems   We want to recover a stable approximation of $f \in \mathbf{L^2}(\mathbb{R})$ from a filtering and uniform sampling that outputs $\{f \star \bar{\phi}_s(ns)\}_{n \in \mathbb{Z}}$, for some real filter $\bar{\phi}_s(t)$. These samples can be written as inner products in $\mathbf{L^2}(\mathbb{R})$:

$$f \star \phi_s(ns) = \int_{-\infty}^{+\infty} f(t)\,\bar{\phi}_s(ns - t)\,dt = \langle f(t), \phi_s(t - ns) \rangle \tag{3.13}$$

with $\phi_s(t) = \bar{\phi}_s(-t)$. Let $\mathbf{U}_s$ be the approximation space generated by linear combinations of the $\{\phi_s(t - ns)\}_{n \in \mathbb{Z}}$. The approximation $\tilde{f} \in \mathbf{U}_s$ which minimizes the maximum possible error $\|f - \tilde{f}\|$ is the orthognal projection of f on $\mathbf{U}_s$ (Exercice 3.5). The calculation of this orthogonal projection is stable if $\{\phi_s(t - ns)\}_{n \in \mathbb{Z}}$ is a Riesz basis of $\mathbf{U}_s$, as defined in Section 5.1.1. Following Definition 5.1, a Riesz basis is a family of linearly independent functions that yields inner product satisfying an energy equivalence. There exists $B \geqslant A > 0$ such that for any $f \in \mathbf{U}_s$

$$A \|f\|^2 \leqslant \sum_{n=-\infty}^{+\infty} |\langle f(t), \phi_s(t - ns) \rangle|^2 \leqslant B \|f\|^2. \tag{3.14}$$

The basis is orthogonal if and only if $A = B$. The following generalized sampling theorem computes the orthogonal projection on the approximation space $\mathbf{U}_s$ [468].

Theorem 3.3 (Linear sampling). *Let $\{\phi_s(t - ns)\}_{n \in \mathbb{Z}}$ be a Riesz basis of $\mathbf{U}_s$ and $\bar{\phi}_s(t) = \phi_s(-t)$. There exists a biorthogonal basis $\{\tilde{\phi}_s(t - ns)\}_{n \in \mathbb{Z}}$ of $\mathbf{U}_s$ such that*

$$\forall f \in \mathbf{L^2}(\mathbb{R}) \quad , \quad P_{\mathbf{U}_s} f(t) = \sum_{n=-\infty}^{+\infty} f \star \bar{\phi}_s(ns)\,\tilde{\phi}_s(t - ns) \; . \tag{3.15}$$

Proof. For any Riesz basis, Section 5.1.2 proves that there exists a biorthogonal basis $\{\tilde{\phi}_{s,n}(t)\}_{n \in \mathbb{Z}}$ that satisfies the biorthogonality relations

$$\forall (n, m) \in \mathbb{Z}^2 \quad , \quad \langle \phi_s(t - ns), \tilde{\phi}_{s,m}(t - ms) \rangle = \delta[n - m] \; . \tag{3.16}$$

Since $\langle \phi_s(t - (n-m)s), \tilde{\phi}_{s,0}(t) \rangle = \langle \phi_s(t - ns), \tilde{\phi}_{s,0}(t - ms) \rangle = 0$ and since the dual basis is unique, necessarily $\tilde{\phi}_{s,m}(t) = \tilde{\phi}_{s,0}(t - ms)$. Section 5.1.2 proves in (5.20) that the orthogonal projection in $\mathbf{U}_s$ can be written

$$P_{\mathbf{U}_s} f(t) = \sum_{n=-\infty}^{+\infty} \langle f(t), \phi_s(t - ns) \rangle\,\tilde{\phi}_s(t - ns)$$

which proves (3.15).   ∎

The orthogonal projection (3.15) can be rewritten as an analog filtering of the discrete signal $f_d(t) = \sum_{n=-\infty}^{+\infty} f \star \bar{\phi}_s(ns)\,\delta(t - ns)$:

$$P_{\mathbf{U}_s} f(t) = f_d \star \tilde{\phi}_s(t) \; . \tag{3.17}$$

If $f \in \mathbf{U}_s$ then $P_{\mathbf{U}_s} f = f$ so it is exactly reconstructed by filtering the uniformly sampled discrete signal $\{f \star \bar{\phi}_s(ns)\}_{n \in \mathbb{Z}}$ with the analog filter $\tilde{\phi}(t)$. If $f \notin \mathbf{U}_s$ then (3.17) recovers the best linear approximation approximation of f in $\mathbf{U}_s$. Section 9.1 shows that the linear approximation error $\|f - P_{\mathbf{U}_s} f\|$ depends essentially on the uniform regularity of f. Given some prior information on f, optimizing the analog discretization filter ϕ_s amounts to optimize the approximation space $\mathbf{U}_s$ in order to minimize this error. The following theorem characterizes filters ϕ_s that generate a Riesz basis and computes the dual filter.

Theorem 3.4. *A filter ϕ_s generates a Riesz basis $\{\phi_s(t-ns)\}_{n\in\mathbb{Z}}$ of a space $\mathbf{U}_s$ if and only if there exists $B \geqslant A > 0$ such that*

$$\forall \omega \in [0, 2\pi/s] \quad , \quad A \leqslant \frac{1}{s} \sum_{k=-\infty}^{+\infty} |\hat{\phi}_s(\omega - \frac{2k\pi}{s})|^2 \leqslant B \ . \tag{3.18}$$

The biorthogonal basis $\{\tilde{\phi}_s(t-ns)\}_{n\in\mathbb{Z}}$ is defined by the dual filter $\tilde{\phi}_s$ which satisfies:

$$\widehat{\tilde{\phi}_s}(\omega) = \frac{s\,\hat{\phi}_s^*(\omega)}{\sum_{k=-\infty}^{+\infty} |\hat{\phi}_s(\omega - 2k\pi/s)|^2} \ . \tag{3.19}$$

Proof. Theorem 5.1 proves that $\{\phi_s(t-ns)\}_{n\in\mathbb{Z}}$ is a Riesz basis of $\mathbf{U}_s$ with Riesz bounds $B \geqslant A > 0$ if and only if

$$\forall a \in \boldsymbol{\ell}^2(\mathbb{Z}) \quad , \quad A\|a\|^2 \leqslant \|\sum_{n\in\mathbb{Z}} a[ns]\,\phi_s(t-ns)\|^2 \leqslant B\|a\|^2 \ , \tag{3.20}$$

with $\|a\|^2 = \sum_{n\in\mathbb{Z}} |a[ns]|^2$.

Let us first write these conditions in the Fourier domain. The Fourier transform of $f(t) = \sum_{n=-\infty}^{+\infty} a[ns]\,\phi_s(t-ns)$ is

$$\hat{f}(\omega) = \sum_{n=-\infty}^{+\infty} a[ns]\,\mathrm{e}^{-ins\omega}\,\hat{\phi}_s(\omega) = \hat{a}(\omega)\,\hat{\phi}_s(\omega) \tag{3.21}$$

where $\hat{a}(\omega)$ is the Fourier series $\hat{a}(\omega) = \sum_{n=-\infty}^{+\infty} a[ns]\,\mathrm{e}^{-ins\omega}$. Let us relate the norm of f and $\hat{a}$. Since $\hat{a}(\omega)$ is $2\pi/s$ periodic, inserting (3.21) in the Plancherel formula (2.26) gives

$$\begin{aligned}
\|f\|^2 &= \frac{1}{2\pi} \int_{-\infty}^{+\infty} |\hat{f}(\omega)|^2\,d\omega = \frac{1}{2\pi} \int_0^{2\pi/s} \sum_{k=-\infty}^{+\infty} |\hat{a}(\omega + 2k\pi/s)|^2\,|\hat{\phi}_s(\omega + 2k\pi/s)|^2\,d\omega \\
&= \frac{1}{2\pi} \int_0^{2\pi/s} |\hat{a}(\omega)|^2 \sum_{k=-\infty}^{+\infty} |\hat{\phi}_s(\omega + 2k\pi/s)|^2\,d\omega \ .
\end{aligned} \tag{3.22}$$

Section 3.2.2 on Fourier series proves that

$$\|a\|^2 = \sum_{n=-\infty}^{+\infty} |a[ns]|^2 = \frac{s}{2\pi} \int_0^{2\pi/s} |\hat{a}(\omega)|^2\,d\omega \ . \tag{3.23}$$

As a consequence of (3.22) and (3.23), the Riesz bound inequalities (3.20) are equivalent to

$$\forall \hat{a} \in \mathbf{L}^2[0, 2\pi/s] \quad , \quad \frac{1}{2\pi} \int_0^{2\pi/s} |\hat{a}(\omega)|^2 \sum_{k=-\infty}^{+\infty} |\hat{\phi}_s(\omega + 2k\pi/s)|^2\,d\omega \leqslant \frac{B\,s}{2\pi} \int_0^{2\pi/s} |\hat{a}(\omega)|^2\,d\omega \tag{3.24}$$

and

$$\forall \hat{a} \in \mathbf{L}^2[0, 2\pi/s] \quad , \quad \frac{1}{2\pi} \int_0^{2\pi/s} |\hat{a}(\omega)|^2 \sum_{k=-\infty}^{+\infty} |\hat{\phi}_s(\omega + 2k\pi/s)|^2\,d\omega \geqslant \frac{A\,s}{2\pi} \int_0^{2\pi/s} |\hat{a}(\omega)|^2\,d\omega \ . \tag{3.25}$$

If $\hat{\phi}_s$ satisfies (3.18) then clearly (3.24) and (3.25) are valid, which proves (3.22).

Conversely, if $\{\phi_s(ns - t)\}_{n\in\mathbb{Z}}$ is a Riesz basis. Suppose that either the upper or the lower bound of (3.18) is not satisfied for ω in a set of non-zero measure. Let $\hat{a}$ be the indicator function of this set. Then either (3.24) or (3.25) are not valid for this $\hat{a}$. This implies that the Riesz bounds (3.20) are not valid for a and hence that it is not a Riesz basis, which contradicts our hypothesis. So (3.18) is indeed valid for almost all $\omega \in [0, 2\pi/s]$.

To compute the biorthogonal basis, we are looking for $\tilde{\phi}_s \in \mathbf{U}_s$ such that $\{\tilde{\phi}_s(t-ns)\}_{n\in\mathbb{Z}}$ satisfies the biorthogonal relations (3.16). Since $\tilde{\phi}_s \in \mathbf{U}_s$ we saw in (3.21) that its Fourier transform can be written $\widehat{\tilde{\phi}_s}(\omega) = \hat{a}(\omega)\hat{\phi}_s(\omega)$ where $\hat{a}(\omega)$ is $2\pi/s$ periodic. Let us define $g(t) = \bar{\phi}_s \star \tilde{\phi}_s(t)$. Its Fourier transform is

$$\hat{g}(\omega) = \hat{\phi}_s^*(\omega)\,\widehat{\tilde{\phi}_s}(\omega) = \hat{a}(\omega)|\hat{\phi}_s(\omega)|^2 \ .$$

The biorthogonal relations (3.16) are satisfied if and only if $g(ns) = 0$ if $n \neq 0$ and $g(0) = 1$. It results that $g_d(t) = \sum_{n=-\infty}^{+\infty} g(ns)\,\delta(t-ns) = \delta(t)$. Theorem 3.1 derives in (3.3) that

$$\hat{g}_d(\omega) = \frac{1}{s} \sum_{k=-\infty}^{+\infty} \hat{g}\,(\omega - 2k\pi/s) = \frac{\hat{a}(\omega)}{s} \sum_{k=-\infty}^{+\infty} |\hat{\phi}_s(\omega - 2k\pi/s)|^2 = 1.$$

It results that

$$\hat{a}(\omega) = s \left(\sum_{k=-\infty}^{+\infty} |\hat{\phi}_s(\omega - 2k\pi/s)|^2 \right)^{-1}$$

which proves (3.19). ∎ ∎

This theorem gives a necessary and sufficient condition on the low-pass filter $\bar{\phi}_s(t) = \phi_s(-t)$ to recover a stable signal approximation from a uniform sampling at intervals s. For various sampling intervals s, the low-pass filter can be obtained by dilating a single filter $\phi_s(t) = s^{-1/2}\phi(t/s)$ and hence $\hat{\phi}_s(\omega) = s^{1/2}\hat{\phi}(s\omega)$. The necessary and sufficient Riesz basis condition (3.18) is then satisfied if and only if

$$\forall \omega \in [-\pi, \pi] \quad , \quad A \leqslant \sum_{k=-\infty}^{+\infty} |\hat{\phi}(\omega - 2k\pi)|^2 \leqslant B. \tag{3.26}$$

It results from (3.19) that the dual filter satisfies $\widehat{\tilde{\phi}_s}(\omega) = s^{1/2}\widehat{\tilde{\phi}}(s\omega)$ and hence $\tilde{\phi}_s(t) = s^{-1/2}\tilde{\phi}(t/s)$. When $A = B = 1$ the Riesz basis is an orthonormal basis, which proves the following corollary.

Corollary 3.1. *The family $\{\phi_s(t-ns)\}_{n\in\mathbb{Z}}$ is an orthonormal basis of the space $\mathbf{U}_s$ it generates, with $\phi_s(t) = s^{-1/2}\phi(t/s)$, if and only if*

$$\forall \omega \in [0, 2\pi] \quad , \quad \sum_{k=-\infty}^{+\infty} |\hat{\phi}(\omega - 2k\pi)|^2 = 1 \ , \tag{3.27}$$

and the dual filter is $\tilde{\phi}_s = \phi_s$.

Shannon-Whittaker revisited The Shannon-Whittaker theorem 3.2 is defined with a sine-cardinal perfect low-pass filter ϕ_s, which we renormalize here to have a unit norm. The following theorem proves that it samples functions in an orthonormal basis.

Theorem 3.5. *If $\phi_s(t) = s^{1/2}\sin(\pi s^{-1}t)/(\pi t)$ then $\{\phi_s(t-ns)\}_{n\in\mathbb{Z}}$ is an orthonormal basis of the space $\mathbf{U}_s$ of functions whose Fourier transforms have a support included in $[-\pi/s, \pi/s]$. If $f \in \mathbf{U}_s$ then*

$$f(nT) = s^{-1/2}\,f \star \phi_s(ns) \ . \tag{3.28}$$

Proof. The filter satisfies $\phi_s(t) = s^{-1/2}\phi(t/s)$ with $\phi(t) = \sin(\pi t)/(\pi t)$. The Fourier transform $\hat{\phi}(\omega) = \mathbf{1}_{[-\pi,\pi]}(\omega)$ satisfies the condition (3.27) of Corollary 3.1, which proves that $\{\phi_s(t-ns)\}_{n\in\mathbb{Z}}$ is an orthonormal basis of a space $\mathbf{U}_s$.

Any $f(t) = \sum_{n=-\infty}^{+\infty} a[ns]\,\phi_s(t-ns) \in \mathbf{U}_s$ has a Fourier transform which can be written

$$\hat{f}(\omega) = \sum_{n=-\infty}^{+\infty} a[ns]\,\mathrm{e}^{-ins\omega}\,\hat{\phi}_s(\omega) = \hat{a}(\omega)\,s^{1/2}\,\mathbf{1}_{[-\pi/s,\pi/s]} \ , \tag{3.29}$$

which implies that $f \in \mathbf{U}_s$ if and only if f has a Fourier transform supported in $[-\pi/s, \pi/s]$.

If $f \in \mathbf{U}_s$ then decomposing it the orthonormal basis $\{\phi_s(t-ns)\}_{n\in\mathbb{Z}}$ gives

$$f(t) = P_{\mathbf{U}_s} f(t) = \sum_{n\in\mathbb{Z}} \langle f(u), \phi_s(u-ns) \rangle \, \phi_s(t-ns) \ .$$

Since $\phi_s(ps) = s^{-1/2}\delta[ps]$ and $\phi_s(-t) = \phi_s(t)$, it results that

$$f(ns) = s^{-1/2}\langle f(u), \phi_s(u-ns) \rangle = s^{-1/2}\,f \star \phi_s(ns) \ .$$

∎

This theorem proves that in the particular case of the Shannon-Whittaker sampling theorem, if $f \in \mathbf{U}_s$ then the sampled low-pass filtered values $f \star \phi_s(ns)$ are proportional to the signal samples $f(ns)$. This comes from the fact that the sine-cardinal $\phi(t) = \sin(\pi t/s)/(\pi t/s)$ satisfies the interpolation property $\phi(ns) = \delta[ns]$. A generalization of such multiscale interpolations is studied in Section 7.6.

The Shannon-Whittaker sampling approximates signals by restricting their Fourier transform to a low frequency interval. It is particularly effective for smooth signals whose Fourier transform have an energy concentrated at low frequencies. It is also well adapted to sound recordings, which are well approximated by lower frequency harmonics.

For discontinuous signals such as images, a low-frequency restriction produces Gibbs oscillations studied in Section 2.3.3. The image visual quality is degraded by these oscillations, which have a total variation (2.65) that is infinite. A piecewise constant approximation has the advantage of creating no such spurious oscillations.

Block Sampler   A block sampler approximates signals with piecewise constant functions. The approximation space $\mathbf{U}_s$ is the set of all functions that are constant on intervals $[ns, (n+1)s]$, for any $n \in \mathbb{Z}$. Let $\phi_s(t) = s^{-1/2} \mathbf{1}_{[0,s]}(t)$. The family $\{\phi_s(t - ns]\}_{n \in \mathbb{Z}}$ is an orthonormal basis of $\mathbf{U}_s$ (Exercise 3.1). If $f \notin \mathbf{U}_s$ then its orthogonal projection on $\mathbf{U}_s$ is calculated with a partial decomposition in the block orthonormal basis of $\mathbf{U}_s$

$$P_{\mathbf{U}_s} f(t) = \sum_{n=-\infty}^{+\infty} \langle f(u), \phi_s(u - ns) \rangle \, \phi_s(t - ns), \tag{3.30}$$

and each coefficient is proportional to the signal average on $[ns, (n+1)s]$

$$\langle f(u), \phi_s(u - ns) \rangle = f \star \phi_s(ns) = s^{-1/2} \int_{ns}^{(n+1)s} f(u) \, du \ .$$

This block analog to digital conversion is particularly simple to implement in analog electronic, where the integration is performed by a capacity.

In domains where f is a regular function, a piecewise constant approximation $P_{\mathbf{U}_s} f$ is not very precise and can be significantly improved. More precise approximations are obtained with approximation spaces $\mathbf{U}_s$ of higher order polynomial splines. The resulting approximations can introduce small Gibbs oscillations, but these oscillations have a finite total variation.

Spline Sampling   Block samplers are generalized by a spline sampling with a space $\mathbf{U}_s$ of splines functions that are $m - 1$ times continuously differentiable and equal to a polynomial of degree m on any interval $[ns, (n+1)s)$, for $n \in \mathbb{Z}$. When $m = 1$, functions in $\mathbf{U}_s$ are piecewise linear and continuous.

A Riesz basis of polynomial splines is constructed with *box splines*. A box spline ϕ of degree m is computed by convolving the box window $\mathbf{1}_{[0,1]}$ with itself $m + 1$ times and centering it at 0 or $1/2$. Its Fourier transform is

$$\hat{\phi}(\omega) = \left(\frac{\sin(\omega/2)}{\omega/2} \right)^{m+1} \exp\left(\frac{-i \varepsilon \omega}{2} \right). \tag{3.31}$$

If m is even then $\varepsilon = 1$ and ϕ has a support centered at $t = 1/2$. If m is odd then $\varepsilon = 0$ and $\phi(t)$ is symmetric about $t = 0$. One can verify that $\hat{\phi}(\omega)$ satisfies the sampling condition (3.26) using a closed form expression (7.19) of the resulting series. It results that for any $s > 0$, a box splines family $\{\phi_s(t - ns)\}_{n \in \mathbb{Z}}$ defines a Riesz basis of $\mathbf{U}_s$, and thus a stable sampling.

3.2   Discrete Time-Invariant Filters

3.2.1   Impulse Response and Transfer Function

Classical discrete signal processing algorithms are mostly based on time-invariant linear operators [50, 54]. The time-invariance is limited to translations on the sampling grid. To simplify notation,

the sampling interval is normalized $s = 1$, and we denote $f[n]$ the sample values. A linear discrete operator L is time-invariant if an input $f[n]$ delayed by $p \in \mathbb{Z}$, $f_p[n] = f[n-p]$, produces an output also delayed by p:

$$Lf_p[n] = Lf[n-p].$$

Impulse Response   We denote by $\delta[n]$ the discrete Dirac

$$\delta[n] = \left\{ \begin{array}{ll} 1 & \text{if } n = 0 \\ 0 & \text{if } n \neq 0 \end{array} \right. . \tag{3.32}$$

Any signal $f[n]$ can be decomposed as a sum of shifted Diracs

$$f[n] = \sum_{p=-\infty}^{+\infty} f[p]\,\delta[n-p].$$

Let $L\delta[n] = h[n]$ be the discrete *impulse response*. The linearity and time-invariance implies that

$$Lf[n] = \sum_{p=-\infty}^{+\infty} f[p]\,h[n-p] = f \star h[n]. \tag{3.33}$$

A discrete linear time-invariant operator is thus computed with a discrete convolution. If $h[n]$ has a finite support the sum (3.33) is calculated with a finite number of operations. These are called *Finite Impulse Response* (FIR) filters. Convolutions with infinite impulse response filters may also be calculated with a finite number of operations if they can be rewritten with a recursive equation (3.45).

Causality and Stability   A discrete filter L is *causal* if $Lf[p]$ depends only on the values of $f[n]$ for $n \leqslant p$. The convolution formula (3.33) implies that $h[n] = 0$ if $n < 0$.

   The filter is *stable* if any bounded input signal $f[n]$ produces a bounded output signal $Lf[n]$. Since

$$|Lf[n]| \leqslant \sup_{n \in \mathbb{Z}} |f[n]| \sum_{k=-\infty}^{+\infty} |h[k]|,$$

it is sufficient that $\sum_{n=-\infty}^{+\infty} |h[n]| < +\infty$, which means that $h \in \ell^1(\mathbb{Z})$. One can verify that this sufficient condition is also necessary. The filter is thus stable if and only if $h \in \ell^1(\mathbb{Z})$ (Exercise 3.6)

Transfer Function   The Fourier transform plays a fundamental role in analyzing discrete time-invariant operators, because discrete sinusoidal waves $e_\omega[n] = \mathrm{e}^{i\omega n}$ are eigenvectors:

$$Le_\omega[n] = \sum_{p=-\infty}^{+\infty} \mathrm{e}^{i\omega(n-p)}\,h[p] = \mathrm{e}^{i\omega n} \sum_{p=-\infty}^{+\infty} h[p]\,\mathrm{e}^{-i\omega p}. \tag{3.34}$$

The eigenvalue is a Fourier series

$$\hat{h}(\omega) = \sum_{p=-\infty}^{+\infty} h[p]\,\mathrm{e}^{-i\omega p}. \tag{3.35}$$

It is the filter *transfer function*.

Example 3.2. *The uniform discrete average*

$$Lf[n] = \frac{1}{2N+1} \sum_{p=n-N}^{n+N} f[p]$$

is a time-invariant discrete filter whose impulse response is $h = (2N+1)^{-1}\mathbf{1}_{[-N,N]}$. *Its transfer function is*

$$\hat{h}(\omega) = \frac{1}{2N+1} \sum_{n=-N}^{+N} \mathrm{e}^{-in\omega} = \frac{1}{2N+1} \frac{\sin(N+1/2)\omega}{\sin \omega/2}. \tag{3.36}$$

3.2.2 Fourier Series

The properties of Fourier series are essentially the same as the properties of the Fourier transform since Fourier series are particular instances of Fourier transforms for Dirac sums. If $f(t) = \sum_{n=-\infty}^{+\infty} f[n]\, \delta(t-n)$ then $\hat{f}(\omega) = \sum_{n=-\infty}^{+\infty} f[n]\, \mathrm{e}^{-i\omega n}$.

For any $n \in \mathbb{Z}$, $\mathrm{e}^{-i\omega n}$ has period 2π, so Fourier series have period 2π. An important issue is to understand whether all functions with period 2π can be written as Fourier series. Such functions are characterized by their restriction to $[-\pi, \pi]$. We therefore consider functions $\hat{a} \in \mathbf{L}^2[-\pi, \pi]$ that are square integrable over $[-\pi, \pi]$. The space $\mathbf{L}^2[-\pi, \pi]$ is a Hilbert space with the inner product

$$\langle \hat{a}, \hat{b} \rangle = \frac{1}{2\pi} \int_{-\pi}^{\pi} \hat{a}(\omega)\, \hat{b}^*(\omega)\, d\omega \tag{3.37}$$

and the resulting norm

$$\|\hat{a}\|^2 = \frac{1}{2\pi} \int_{-\pi}^{\pi} |\hat{a}(\omega)|^2\, d\omega.$$

The following theorem proves that any function in $\mathbf{L}^2[-\pi, \pi]$ can be written as a Fourier series.

Theorem 3.6. *The family of functions $\{\mathrm{e}^{-ik\omega}\}_{k \in \mathbb{Z}}$ is an orthonormal basis of $\mathbf{L}^2[-\pi, \pi]$.*

Proof. The orthogonality with respect to the inner product (3.37) is established with a direct integration. To prove that $\{\exp(-ik\omega)\}_{k \in \mathbb{Z}}$ is a basis, we must show that linear expansions of these vectors are dense in $\mathbf{L}^2[-\pi, \pi]$.

We first prove that any continuously differentiable function $\hat{\phi}$ with a support included in $[-\pi, \pi]$ satisfies

$$\hat{\phi}(\omega) = \sum_{k=-\infty}^{+\infty} \langle \hat{\phi}(\xi), \exp(-ik\xi) \rangle \, \exp(-ik\omega) \, , \tag{3.38}$$

with a pointwise convergence for any $\omega \in [-\pi, \pi]$. Let us compute the partial sum

$$
\begin{aligned}
S_N(\omega) &= \sum_{k=-N}^{N} \langle \hat{\phi}(\xi), \exp(-ik\xi) \rangle \, \exp(-ik\omega) \\
&= \sum_{k=-N}^{N} \frac{1}{2\pi} \int_{-\pi}^{\pi} \hat{\phi}(\xi) \exp(ik\xi)\, d\xi \, \exp(-ik\omega) \\
&= \frac{1}{2\pi} \int_{-\pi}^{\pi} \hat{\phi}(\xi) \sum_{k=-N}^{N} \exp[ik(\xi - \omega)]\, d\xi.
\end{aligned}
$$

The Poisson formula (2.37) proves the distribution equality

$$\lim_{N \to +\infty} \sum_{k=-N}^{N} \exp[ik(\xi - \omega)] = 2\pi \sum_{k=-\infty}^{+\infty} \delta(\xi - \omega - 2\pi k),$$

and since the support of $\hat{\phi}$ is in $[-\pi, \pi]$ we get

$$\lim_{N \to +\infty} S_N(\omega) = \hat{\phi}(\omega).$$

Since $\hat{\phi}$ is continuously differentiable, following the steps (2.38-2.40) in the proof of the Poisson formula shows that $S_N(\omega)$ converges uniformly to $\hat{\phi}(\omega)$ on $[-\pi, \pi]$.

To prove that linear expansions of sinusoidal waves $\{\exp(-ik\omega)\}_{k \in \mathbb{Z}}$ are dense in $\mathbf{L}^2[-\pi, \pi]$, let us verify that the distance between $\hat{a} \in \mathbf{L}^2[-\pi, \pi]$ and such a linear expansion is less than ε, for any $\varepsilon > 0$. Continuously differentiable functions with a support included in $[-\pi, \pi]$ are dense in $\mathbf{L}^2[-\pi, \pi]$, hence there exists $\hat{\phi}$ such that $\|\hat{a} - \hat{\phi}\| \leqslant \varepsilon/2$. The uniform pointwise convergence proves that there exists N for which

$$\sup_{\omega \in [-\pi, \pi]} |S_N(\omega) - \hat{\phi}(\omega)| \leqslant \frac{\varepsilon}{2},$$

which implies that

$$\|S_N - \hat{\phi}\|^2 = \frac{1}{2\pi} \int_{-\pi}^{\pi} |S_N(\omega) - \hat{\phi}(\omega)|^2\, d\omega \leqslant \frac{\varepsilon^2}{4} \, .$$

It follows that $\hat{a}$ is approximated by the Fourier series S_N with an error

$$\|\hat{a} - S_N\| \leqslant \|\hat{a} - \hat{\phi}\| + \|\hat{\phi} - S_N\| \leqslant \varepsilon .$$

■   ■

Theorem 3.6 proves that if $f \in \boldsymbol{\ell}^2(\mathbb{Z})$, the Fourier series

$$\hat{f}(\omega) = \sum_{n=-\infty}^{+\infty} f[n] \, \mathrm{e}^{-i\omega n} \tag{3.39}$$

can be interpreted as the decomposition of $\hat{f}$ in an orthonormal basis of $\mathbf{L}^2[-\pi, \pi]$. The Fourier series coefficients can thus be written as inner products in $\mathbf{L}^2[-\pi, \pi]$:

$$f[n] = \langle \hat{f}(\omega), \mathrm{e}^{-i\omega n} \rangle = \frac{1}{2\pi} \int_{-\pi}^{\pi} \hat{f}(\omega) \, \mathrm{e}^{i\omega n} \, d\omega. \tag{3.40}$$

The energy conservation of orthonormal bases (A.10) yields a Plancherel identity:

$$\sum_{n=-\infty}^{+\infty} |f[n]|^2 = \|\hat{f}\|^2 = \frac{1}{2\pi} \int_{-\pi}^{\pi} |\hat{f}(\omega)|^2 \, d\omega. \tag{3.41}$$

Pointwise Convergence   The equality (3.39) is meant in the sense of mean-square convergence

$$\lim_{N \to +\infty} \left\| \hat{f}(\omega) - \sum_{k=-N}^{N} f[k] \, \mathrm{e}^{-i\omega k} \right\| = 0.$$

It does not imply a pointwise convergence at all $\omega \in \mathbb{R}$. In 1873, Dubois-Reymond constructed a periodic function $\hat{f}(\omega)$ that is continuous and whose Fourier series diverges at some points. On the other hand, if $\hat{f}(\omega)$ is continuously differentiable, then the proof of Theorem 3.6 shows that its Fourier series converges uniformly to $\hat{f}(\omega)$ on $[-\pi, \pi]$. It was only in 1966 that Carleson [148] was able to prove that if $\hat{f} \in \mathbf{L}^2[-\pi, \pi]$ then its Fourier series converges almost everywhere. The proof is very technical.

Convolutions   Since $\{\mathrm{e}^{-i\omega k}\}_{k \in \mathbb{Z}}$ are eigenvectors of discrete convolution operators, we also have a discrete convolution theorem.

Theorem 3.7. *If $f \in \boldsymbol{\ell}^1(\mathbb{Z})$ and $h \in \boldsymbol{\ell}^1(\mathbb{Z})$ then $g = f \star h \in \boldsymbol{\ell}^1(\mathbb{Z})$ and*

$$\hat{g}(\omega) = \hat{f}(\omega) \, \hat{h}(\omega). \tag{3.42}$$

The proof is identical to the proof of the convolution Theorem 2.2, if we replace integrals by discrete sums. The reconstruction formula (3.40) shows that a filtered signal can be written

$$f \star h[n] = \frac{1}{2\pi} \int_{-\pi}^{\pi} \hat{h}(\omega) \hat{f}(\omega) \, \mathrm{e}^{i\omega n} \, d\omega. \tag{3.43}$$

The transfer function $\hat{h}(\omega)$ amplifies or attenuates the frequency components $\hat{f}(\omega)$ of $f[n]$.

Example 3.3. *An ideal discrete low-pass filter has a 2π periodic transfer function defined by $\hat{h}(\omega) = \mathbf{1}_{[-\xi,\xi]}(\omega)$, for $\omega \in [-\pi, \pi]$ and $0 < \xi < \pi$. Its impulse response is computed with (3.40):*

$$h[n] = \frac{1}{2\pi} \int_{-\xi}^{\xi} \mathrm{e}^{i\omega n} \, d\omega = \frac{\sin \xi n}{\pi n}. \tag{3.44}$$

It is a uniform sampling of the ideal analog low-pass filter (2.29).

Example 3.4. *A recursive filter computes $g = Lf$ which is solution of a recursive equation*

$$\sum_{k=0}^{K} a_k \, f[n-k] = \sum_{k=0}^{M} b_k \, g[n-k], \tag{3.45}$$

with $b_0 \neq 0$. If $g[n] = 0$ and $f[n] = 0$ for $n < 0$ then g has a linear and time-invariant dependency upon f, and can thus be written $g = f \star h$. The transfer function is obtained by computing the Fourier transform of (3.45). The Fourier transform of $f_k[n] = f[n-k]$ is $\hat{f}_k(\omega) = \hat{f}(\omega) \, \mathrm{e}^{-ik\omega}$ so

$$\hat{h}(\omega) = \frac{\hat{g}(\omega)}{\hat{f}(\omega)} = \frac{\sum_{k=0}^{K} a_k \, \mathrm{e}^{-ik\omega}}{\sum_{k=0}^{M} b_k \, \mathrm{e}^{-ik\omega}}.$$

It is a rational function of $\mathrm{e}^{-i\omega}$. If $b_k \neq 0$ for some $k > 0$ then one can verify that the impulse response h has an infinite support. The stability of such filters is studied in Exercise 3.18. A direct calculation of the convolution sum $g[n] = f \star h[n]$ would require an infinite number of operations whereas (3.45) computes $g[n]$ with $K + M$ additions and multiplications from its past values.

Window Multiplication   An infinite impulse response filter h such as the ideal low-pass filter (3.44) may be approximated by a finite response filter $\tilde{h}$ by multiplying h with a window g of finite support:

$$\tilde{h}[n] = g[n] \, h[n].$$

One can verify (Exercise 3.7) that a multiplication in time is equivalent to a convolution in the frequency domain:

$$\widehat{\tilde{h}}(\omega) = \frac{1}{2\pi} \int_{-\pi}^{\pi} \hat{h}(\xi) \, \hat{g}(\omega - \xi) \, d\xi = \frac{1}{2\pi} \hat{h} \star \hat{g}(\omega). \tag{3.46}$$

Clearly $\widehat{\tilde{h}} = \hat{h}$ only if $\hat{g} = 2\pi\delta$, which would imply that g has an infinite support and $g[n] = 1$. The approximation $\widehat{\tilde{h}}$ is close to $\hat{h}$ only if $\hat{g}$ approximates a Dirac, which means that all its energy is concentrated at low frequencies. In time, g should therefore have smooth variations.

The rectangular window $g = \mathbf{1}_{[-N,N]}$ has a Fourier transform $\hat{g}$ computed in (3.36). It has important side lobes far away from $\omega = 0$. The resulting $\widehat{\tilde{h}}$ is a poor approximation of $\hat{h}$. The Hanning window

$$g[n] = \cos^2\left(\frac{\pi n}{2N}\right) \mathbf{1}_{[-N,N]}[n]$$

is smoother and thus has a Fourier transform better concentrated at low frequencies. The spectral properties of other windows are studied in Section 4.2.2.

3.3   Finite Signals

Up to now, we have considered discrete signals $f[n]$ defined for all $n \in \mathbb{Z}$. In practice, $f[n]$ is known over a finite domain, say $0 \leqslant n < N$. Convolutions must therefore be modified to take into account the border effects at $n = 0$ and $n = N - 1$. The Fourier transform must also be redefined over finite sequences for numerical computations. The fast Fourier transform algorithm is explained as well as its application to fast convolutions.

3.3.1   Circular Convolutions

Let $\tilde{f}$ and $\tilde{h}$ be signals of N samples. To compute the convolution product

$$\tilde{f} \star \tilde{h}[n] = \sum_{p=-\infty}^{+\infty} \tilde{f}[p] \, \tilde{h}[n-p] \ \text{ for } 0 \leqslant n < N,$$

we must know $\tilde{f}[n]$ and $\tilde{h}[n]$ beyond $0 \leqslant n < N$. One approach is to extend $\tilde{f}$ and $\tilde{h}$ with a periodization over N samples, and define

$$f[n] = \tilde{f}[n \bmod N] \ \ , \ \ h[n] = \tilde{h}[n \bmod N].$$

The *circular convolution* of two such signals, both with period N, is defined as a sum over their period:

$$f \circledast h[n] = \sum_{p=0}^{N-1} f[p] \, h[n-p] = \sum_{p=0}^{N-1} f[n-p] \, h[p].$$

It is also a signal of period N.

The eigenvectors of a circular convolution operator

$$Lf[n] = f \circledast h[n]$$

are the discrete complex exponentials $e_k[n] = \exp(i2\pi kn/N)$. Indeed

$$Le_k[n] = \exp\left(\frac{i2\pi kn}{N}\right) \sum_{p=0}^{N-1} h[p] \, \exp\left(\frac{-i2\pi kp}{N}\right),$$

and the eigenvalue is the discrete Fourier transform of h:

$$\hat{h}[k] = \sum_{p=0}^{N-1} h[p] \, \exp\left(\frac{-i2\pi kp}{N}\right).$$

3.3.2 Discrete Fourier Transform

The space of signals of period N is an Euclidean space of dimension N and the inner product of two such signals f and g is

$$\langle f, g \rangle = \sum_{n=0}^{N-1} f[n] \, g^*[n]. \tag{3.47}$$

The following theorem proves that any signal with period N can be decomposed as a sum of discrete sinusoidal waves.

Theorem 3.8. *The family*

$$\left\{ e_k[n] = \exp\left(\frac{i2\pi kn}{N}\right) \right\}_{0 \leqslant k < N}$$

is an orthogonal basis of the space of signals of period N.

Since the space is of dimension N, any orthogonal family of N vectors is an orthogonal basis. To prove this theorem it is therefore sufficient to verify that $\{e_k\}_{0 \leqslant k < N}$ is orthogonal with respect to the inner product (3.47) (Exercise 3.8). Any signal f of period N can be decomposed in this basis:

$$f = \sum_{k=0}^{N-1} \frac{\langle f, e_k \rangle}{\|e_k\|^2} \, e_k. \tag{3.48}$$

By definition, the *discrete Fourier transform* (DFT) of f is

$$\hat{f}[k] = \langle f, e_k \rangle = \sum_{n=0}^{N-1} f[n] \, \exp\left(\frac{-i2\pi kn}{N}\right). \tag{3.49}$$

Since $\|e_k\|^2 = N$, (3.48) gives an inverse discrete Fourier formula:

$$f[n] = \frac{1}{N} \sum_{k=0}^{N-1} \hat{f}[k] \, \exp\left(\frac{i2\pi kn}{N}\right). \tag{3.50}$$

The orthogonality of the basis also implies a Plancherel formula

$$\|f\|^2 = \sum_{n=0}^{N-1} |f[n]|^2 = \frac{1}{N} \sum_{k=0}^{N-1} |\hat{f}[k]|^2. \tag{3.51}$$

The discrete Fourier transform of a signal f of period N is computed from its values for $0 \leqslant n < N$. Then why is it important to consider it a periodic signal with period N rather than a finite signal of N samples? The answer lies in the interpretation of the Fourier coefficients. The discrete Fourier sum (3.50) defines a signal of period N for which the samples $f[0]$ and $f[N-1]$ are side by side. If $f[0]$ and $f[N-1]$ are very different, this produces a brutal transition in the periodic signal, creating relatively high amplitude Fourier coefficients at high frequencies. For example, Figure 3.3 shows that the "smooth" ramp $f[n] = n$ for $0 \leqslant n < N$ has sharp transitions at $n = 0$ and $n = N$ once made periodic.

Circular Convolutions   Since $\{\exp(i2\pi kn/N)\}_{0 \leqslant k < N}$ are eigenvectors of circular convolutions, we derive a convolution theorem.
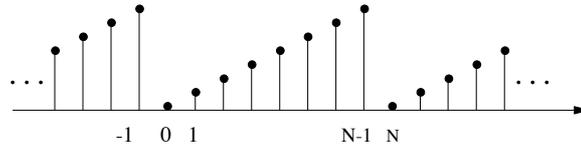
Figure 3.3: Signal $f[n] = n$ for $0 \leqslant n < N$ made periodic over N samples.

Theorem 3.9. *If f and h have period N then the discrete Fourier transform of $g = f \circledast h$ is*

$$\hat{g}[k] = \hat{f}[k]\,\hat{h}[k]. \tag{3.52}$$

The proof is similar to the proof of the two previous convolution Theorems 2.2 and 3.7. This theorem shows that a circular convolution can be interpreted as a discrete frequency filtering. It also opens the door to fast computations of convolutions using the fast Fourier transform.

3.3.3   Fast Fourier Transform

For a signal f of N points, a direct calculation of the N discrete Fourier sums

$$\hat{f}[k] = \sum_{n=0}^{N-1} f[n] \exp\left(\frac{-i2\pi kn}{N}\right) \quad , \quad \text{for } 0 \leqslant k < N, \tag{3.53}$$

requires N^2 complex multiplications and additions. The fast Fourier transform (FFT) algorithm reduces the numerical complexity to $O(N \log_2 N)$ by reorganizing the calculations.

When the frequency index is even, we group the terms n and $n + N/2$:

$$\hat{f}[2k] = \sum_{n=0}^{N/2-1} \left(f[n] + f[n+N/2]\right) \exp\left(\frac{-i2\pi kn}{N/2}\right). \tag{3.54}$$

When the frequency index is odd, the same grouping becomes

$$\hat{f}[2k+1] = \sum_{n=0}^{N/2-1} \exp\left(\frac{-i2\pi n}{N}\right) \left(f[n] - f[n+N/2]\right) \exp\left(\frac{-i2\pi kn}{N/2}\right). \tag{3.55}$$

Equation (3.54) proves that even frequencies are obtained by calculating the discrete Fourier transform of the $N/2$ periodic signal

$$f_e[n] = f[n] + f[n+N/2].$$

Odd frequencies are derived from (3.55) by computing the Fourier transform of the $N/2$ periodic signal

$$f_o[n] = \exp\left(\frac{-i2\pi n}{N}\right) \left(f[n] - f[n+N/2]\right).$$

A discrete Fourier transform of size N may thus be calculated with two discrete Fourier transforms of size $N/2$ plus $O(N)$ operations.

The inverse fast Fourier transform of $\hat{f}$ is derived from the forward fast Fourier transform of its complex conjugate $\hat{f}^*$ by observing that

$$f^*[n] = \frac{1}{N} \sum_{k=0}^{N-1} \hat{f}^*[k] \, \exp\left(\frac{-i2\pi kn}{N}\right). \tag{3.56}$$

Complexity   Let $C(N)$ be the number of elementary operations needed to compute a discrete Fourier transform with the FFT. Since f is complex, the calculation of f_e and f_o requires N complex additions and $N/2$ complex multiplications. Let KN be the corresponding number of elementary operations. We have

$$C(N) = 2\,C(N/2) + K\,N. \tag{3.57}$$

Since the Fourier transform of a single point is equal to itself, $C(1) = 0$. With the change of variable $l = \log_2 N$ and the change of function $T(l) = \frac{C(N)}{N}$, we derive from (3.57) that

$$T(l) = T(l-1) + K.$$

Since $T(0) = 0$ we get $T(l) = K\,l$ and hence

$$C(N) = K\,N\,\log_2(N).$$

There exist several variations of this fast algorithm [236, 48]. The goal is to minimize the constant K. The most efficient fast discrete Fourier transform to this date is the split-radix FFT algorithm, which is slightly more complicated than the procedure just described, but which requires only $N \log_2 N$ real multiplications and $3N \log_2 N$ additions. When the input signal f is real, there are half as many parameters to compute, since $\hat{f}[-k] = \hat{f}^*[k]$. The number of multiplications and additions is thus reduced by 2.

3.3.4   Fast Convolutions

The low computational complexity of a fast Fourier transform makes it efficient to compute finite discrete convolutions by using the circular convolution Theorem 3.9. Let f and h be two signals whose samples are non-zero only for $0 \leqslant n < M$. The causal signal

$$g[n] = f \star h[n] = \sum_{k=-\infty}^{+\infty} f[k]\,h[n-k] \tag{3.58}$$

is non-zero only for $0 \leqslant n < 2M$. If h and f have M non-zero samples, calculating this convolution product with the sum (3.58) requires $M(M+1)$ multiplications and additions. When $M \geqslant 32$, the number of computations is reduced by using the fast Fourier transform [10, 48].

To use the fast Fourier transform with the circular convolution Theorem 3.9, the non-circular convolution (3.58) is written as a circular convolution. We define two signals of period $2M$:

$$a[n] = \begin{cases} f[n] & \text{if } 0 \leqslant n < M \\ 0 & \text{if } M \leqslant n < 2M \end{cases} \tag{3.59}$$

$$b[n] = \begin{cases} h[n] & \text{if } 0 \leqslant n < M \\ 0 & \text{if } M \leqslant n < 2M \end{cases}. \tag{3.60}$$

Let $c = a \circledast b$, one can verify that $c[n] = g[n]$ for $0 \leqslant n < 2M$. The $2M$ non-zero coefficients of g are thus obtained by computing $\hat{a}$ and $\hat{b}$ from a and b and then calculating the inverse discrete Fourier transform of $\hat{c} = \hat{a}\,\hat{b}$. With the fast Fourier transform algorithm, this requires a total of $O(M \log_2 M)$ additions and multiplications instead of $M(M+1)$. A single FFT or inverse FFT of a real signal of size N is calculated with $2^{-1}N \log_2 N$ multiplications, using a split-radix algorithm. The FFT convolution is thus performed with a total of $3M \log_2 M + 11M$ real multiplications. For $M \geqslant 32$ the FFT algorithm is faster than the direct convolution approach. For $M \leqslant 16$, it is faster to use a direct convolution sum.

Fast Overlap-Add Convolutions   The convolution of a signal f of L non-zero samples with a smaller causal signal h of M samples is calculated with an overlap-add procedure that is faster than the previous algorithm. The signal f is decomposed with a sum of L/M blocks f_r having M non-zero samples:

$$f[n] = \sum_{r=0}^{L/M-1} f_r[n - rM] \ \ \text{with} \ \ f_r[n] = f[n + rM]\, \mathbf{1}_{[0,M-1]}[n]. \tag{3.61}$$

For each $0 \leqslant r < L/M$, the $2M$ non-zero samples of $g_r = f_r \star h$ are computed with the FFT based convolution algorithm, which requires $O(M \log_2 M)$ operations. These L/M convolutions are thus obtained with $O(L \log_2 M)$ operations. The block decomposition (3.61) implies that

$$f \star h[n] = \sum_{r=0}^{L/M-1} g_r[n - rM]. \tag{3.62}$$

The addition of these L/M translated signals of size $2M$ is done with $2L$ additions. The overall convolution is thus performed with $O(L \log_2 M)$ operations.

3.4   Discrete Image Processing

Two-dimensional signal processing poses many specific geometrical and topological problems that do not exist in one dimension [20, 32]. For example, a simple concept such as causality is not well defined in two dimensions. We can avoid the complexity introduced by the second dimension by extending one-dimensional algorithms with a separable approach. This not only simplifies the mathematics but also leads to fast numerical algorithms along the rows and columns of images. Appendix A.5 reviews the properties of tensor products for separable calculations.

3.4.1   Two-Dimensional Sampling Theorems

The light intensity measured by a camera is generally sampled over a rectangular array of picture elements, called *pixels*. One-dimensional sampling theorems are extended to this two-dimensional sampling array. Other two-dimensional sampling grids such as hexagonal grids are also possible, but non-rectangular sampling arrays are not often used.

Let s_1 and s_2 be the sampling intervals along the x_1 and x_2 axes of an infinite rectangular sampling grid. In the following we renormalize the axes so that $s_1 = s_2 = s$. A discrete image obtained by sampling $f(x)$ with $x = (x_1, x_2)$ can be represented as a sum of Diracs located at the grid points:

$$f_d(x) = \sum_{n \in \mathbb{Z}^2} f(sn)\, \delta(x - ns) \ .$$

The two-dimensional Fourier transform of $\delta(x - sn)$ is $\mathrm{e}^{-isn\cdot\omega}$ with $\omega = (\omega_1, \omega_2)$ and $n \cdot \omega = n_1\omega_1 + n_2\omega_2$. The Fourier transform of f_d is thus a two-dimensional Fourier series

$$\hat{f}_d(\omega) = \sum_{n \in \mathbb{Z}^2} f(sn)\, \mathrm{e}^{-isn\cdot\omega} \ . \tag{3.63}$$

It is $2\pi/s$ periodic along ω_1 and along ω_2. An extension of Theorem 3.1 relates $\hat{f}_d$ to the two-dimensional Fourier transform $\hat{f}$ of f.

Theorem 3.10. *The Fourier transform of the discrete image $f_d(x)$ is*

$$\hat{f}_d(\omega) = \frac{1}{s^2} \sum_{k \in \mathbb{Z}^2} \hat{f}(\omega - 2k\pi/s) \ \ \textit{with} \ \ k = (k_1, k_2) \ . \tag{3.64}$$

We derive the following two-dimensional sampling theorem, which is analogous to Theorem 3.2.

Theorem 3.11. *If $\hat{f}$ has a support included in $[-\pi/s, \pi/s]^2$ then*

$$f(x) = s \sum_{n \in \mathbb{Z}^2} f(ns)\, \phi_s(x - ns) \; , \tag{3.65}$$

where

$$\phi_s(x_1, x_2) = \frac{1}{s} \frac{\sin(\pi x_1/s)}{\pi x_1/s} \frac{\sin(\pi x_2/s)}{\pi x_2/s} \; . \tag{3.66}$$

If the support of $\hat{f}$ is not included in the low-frequency rectangle $[-\pi/s, \pi/s]^2$, the interpolation formula (3.65) introduces aliasing errors. This aliasing is eliminated by prefiltering f with the ideal low-pass separable filter $\phi_s(x)$ whose Fourier transform is the indicator function of $[-\pi/s, \pi/s]^2$.

General Sampling Theorems   As explained in Section 3.1.3, the Shannon-Whittaker sampling theorem is a particular case of more general linear sampling theorems with low-pass filters. The following theorem is a two-dimensional extension of Theorems 3.3 and 3.4, which characterizes these filters to obtain a stable reconstruction.

Theorem 3.12. *If there exists $B \geqslant A > 0$ such that the Fourier transform of $\phi_s \in \mathbf{L}^2(\mathbb{R}^2)$ satisfies*

$$\forall \omega \in [0, 2\pi/s]^2 \quad A \leqslant \hat{h}(\omega) = \sum_{k \in \mathbb{Z}^2} |\hat{\phi}_s(\omega - 2k\pi/s)|^2 \leqslant B \; ,$$

then $\{\phi_s(x - ns)\}_{n \in \mathbb{Z}^2}$ is a Riesz basis of a space $\mathbf{U}_s$. The Fourier transform of the dual filter $\tilde{\phi}_s$ is $\widehat{\tilde{\phi}_s}(\omega) = \hat{\phi}_s^*(\omega)/h(\omega)$, and the orthogonal projection of $f \in \mathbf{L}^2(\mathbb{R}^2)$ in $\mathbf{U}_s$ is

$$P_{\mathbf{U}_s} f(x) = \sum_{n \in \mathbb{Z}^2} f \star \bar{\phi}_s(ns)\, \tilde{\phi}_s(x - ns) \quad \text{with} \quad \bar{\phi}_s(x) = \phi_s(-x) \; . \tag{3.67}$$

This theorem gives a necessary and sufficient condition to obtain a stable linear reconstruction from samples computed with a linear filter. The proof is a direct extension of the proofs of Theorems 3.3 and 3.4. It recovers a signal approximation as an orthogonal projection, by filtering the discrete signal $f_d(x) = \sum_{n \in \mathbb{Z}^2} f \star \bar{\phi}_s(ns)\, \delta(x - ns)$

$$P_{\mathbf{U}_s} f(x) = f_d \star \tilde{\phi}_s(x) \; .$$

Like in one-dimension, the filter ϕ_s can be obtained by scaling a single filter $\phi_s(x) = s^{-1}\phi(s^{-1}x)$. The two-dimensional Shannon-Whittaker theorem is a particular example, where $\hat{\phi}_s = s\mathbf{1}_{[-\pi/s, \pi/s]^2}$, which defines an orthonormal basis of the space $\mathbf{U}_s$ of functions having a Fourier transform supported in $[-\pi/s, \pi/s]^2$.

3.4.2   Discrete Image Filtering

The properties of two-dimensional space-invariant operators are essentially the same as in one dimension. The sampling interval s is normalized to 1. A pixel value located at $n = (n_1, n_2)$ is written $f[n]$. A linear operator L is space-invariant if $Lf_p[n] = Lf[n - p]$ for any $f_p[n] = f[n - p]$, with $p = (p_1, p_2) \in \mathbb{Z}^2$. A discrete Dirac is defined by $\delta[n] = 1$ if $n = (0, 0)$ and $\delta[n] = 0$ if $n \neq (0, 0)$.

Impulse Response   Since $f[n] = \sum_{p \in \mathbb{Z}^2} f[p]\, \delta[n - p]$, the linearity and time invariance implies

$$Lf[n] = \sum_{p \in \mathbb{Z}^2} f[p]\, h[n - p] = f \star h[n], \tag{3.68}$$

where $h[n]$ is the response of the impulse $h[n] = L\delta[n]$. If the impulse response is separable:

$$h[n_1, n_2] = h_1[n_1]\, h_2[n_2], \tag{3.69}$$

then the two-dimensional convolution (3.68) is computed as one-dimensional convolutions along the columns of the image followed by one-dimensional convolutions along the rows (or vice-versa):

$$f \star h[n_1, n_2] = \sum_{p_1=-\infty}^{+\infty} h_1[n_1 - p_1] \sum_{p_2=-\infty}^{+\infty} h_2[n_2 - p_2] \, f[p_1, p_2]. \tag{3.70}$$

This factorization reduces the number of operations. If h_1 and h_2 are finite impulse response filters respectively of size M_1 and M_2 then the separable calculation (3.70) requires $M_1 + M_2$ additions and multiplications per point (n_1, n_2) as opposed to $M_1 M_2$ in a non-separable computation (3.68).

Transfer Function   The Fourier transform of a discrete image f is defined by the Fourier series

$$\hat{f}(\omega) = \sum_{n \in \mathbb{Z}^2} f[n] \, \mathrm{e}^{-i\omega \cdot n} \quad \text{with} \quad \omega \cdot n = n_1 \omega_1 + n_2 \omega_2 \; . \tag{3.71}$$

The two-dimensional extension of the convolution Theorem 3.7 proves that if $g[n] = Lf[n] = f \star h[n]$ then its Fourier transform is $\hat{g}(\omega) = \hat{f}(\omega) \, \hat{h}(\omega)$, and $\hat{h}(\omega)$ is the transfer function of the filter. When a filter is separable $h[n_1, n_2] = h_1[n_1] \, h_2[n_2]$, its transfer function is also separable:

$$\hat{h}(\omega_1, \omega_2) = \hat{h}_1(\omega_1) \, \hat{h}_2(\omega_2). \tag{3.72}$$

3.4.3   Circular Convolutions and Fourier Basis

The discrete convolution of a finite image $\tilde{f}$ raises border problems. As in one dimension, these border issues are solved by extending the image, making it periodic along its rows and columns:

$$f[n_1, n_2] = \tilde{f}[n_1 \bmod N_1 \, , n_2 \bmod N_2] \; ,$$

where $N = N_1 N_2$ is the image size. The resulting periodic image $f[n_1, n_2]$ is defined for all $(n_1, n_2) \in \mathbb{Z}^2$, and each of its rows and columns are periodic one-dimensional signals.

A discrete convolution is replaced by a circular convolution over the image period. If f and h have a periodicity N_1 and N_2 along (n_1, n_2) then

$$f \circledast h[n_1, n_2] = \sum_{p_1=0}^{N_1-1} \sum_{p_2=0}^{N_2-1} f[p_1, p_2] \, h[n_1 - p_1, n_2 - p_2] \; . \tag{3.73}$$

Discrete Fourier Transform   The eigenvectors of circular convolutions are two-dimensional discrete sinusoidal waves:

$$e_k[n] = \mathrm{e}^{\omega_k \cdot n} \quad \text{with} \quad \omega_k = (2\pi k_1/N_1, 2\pi k_2/N_2) \quad \text{for} \quad 0 \leqslant k_1 < N_1 \; , \; 0 \leqslant k_2 < N_2 \; .$$

This family of $N = N_1 N_2$ discrete vectors is the separable product of two one-dimensional discrete Fourier bases $\{\mathrm{e}^{i2\pi k_1 n/N_1}\}_{0 \leqslant k_1 < N_1}$ and $\{\mathrm{e}^{i2\pi k_1 n/N_2}\}_{0 \leqslant k_2 < N_2}$. Theorem A.6 thus proves that the family $\{e_k[n]\}_{0 \leqslant k_1 < N_1, 0 \leqslant k_2 < N_2}$ is an orthogonal basis of $\mathbb{C}^N = \mathbb{C}^{N_1} \otimes \mathbb{C}^{N_2}$ (Exercise 3.23). Any image $f \in \mathbb{C}^N$ can be decomposed in this orthogonal basis:

$$f[n] = \frac{1}{N} \sum_{k_1=0}^{N_1-1} \sum_{k_2=0}^{N_2-1} \hat{f}[k] \, \mathrm{e}^{i\omega_k \cdot n} \; , \tag{3.74}$$

where $\hat{f}$ is the two-dimensional discrete Fourier transform of f

$$\hat{f}[k] = \langle f, e_k \rangle = \sum_{n_1=0}^{N_1-1} \sum_{n_2=0}^{N_2-1} f[n] \, \mathrm{e}^{-i\omega_k \cdot n} \; . \tag{3.75}$$

Fast Convolutions   Since $e^{i\omega_k \cdot n}$ are eigenvectors of two-dimensional circular convolutions, the discrete Fourier transform of $g = f \circledast h$ is

$$\hat{g}[k] = \hat{f}[k]\,\hat{h}[k]. \tag{3.76}$$

A direct computation of $f \circledast h$ with the summation (3.73) requires $O(N^2)$ multiplications. With the two-dimensional FFT described next, $\hat{f}[k]$ and $\hat{h}[k]$ as well as the inverse DFT of their product (3.76) are calculated with $O(N \log N)$ operations. Non-circular convolutions are computed with a fast algorithm by reducing them to circular convolutions, with the same approach as in Section 3.3.4.

Separable Basis Decomposition   Let $\mathcal{B}_1 = \{e^1_{k_1}\}_{0 \leqslant k_1 < N_1}$ and $\mathcal{B}_2 = \{e^2_{k_2}\}_{0 \leqslant k_2 < N_2}$ be two orthogonal bases of $\mathbb{C}^{N_1}$ and $\mathbb{C}^{N_2}$. Suppose the calculation of decomposition coefficients of $f_1 \in \mathbb{C}^{N_1}$ in the basis $\mathcal{B}_1$ requires $C_1(N_1)$ operations and of $f_2 \in \mathbb{C}^{N_1}$ in the basis $\mathcal{B}_2$ requires $C_2(N_2)$ operations. One can verify (Exercise 3.23) that the family $\mathcal{B} = \{e_k[n] = e^1_{k_1}[n_1]\,e^2_{k_2}[n_2]\}_{0 \leqslant k_1 < N_1, 0 \leqslant k_2 < N_2}$ is an orthogonal basis of the space $\mathbb{C}^N = \mathbb{C}^{N_1} \otimes \mathbb{C}^{N_2}$ of images $f[n_1, n_2]$ of $N = N_1 N_2$ pixels. We describe a fast separable algorithm that computes the decomposition coefficients of an image f in $\mathcal{B}$ with $N_2\,C_1(N_1) + N_1\,C_2(N_2)$ operations as opposed to N^2. A fast two-dimensional FFT is derived.

Two-dimensional inner products are calculated with

$$\begin{aligned}
\langle f, e^1_{k_1} e^2_{k_2}\rangle &= \sum_{n_1=0}^{N_1-1} \sum_{n_2=0}^{N_2-1} f[n_1, n_2]\,e^{1*}_{k_1}[n_1]\,e^{2*}_{k_2}[n_2] \\
&= \sum_{n_1=0}^{N_1-1} e^{1*}_{k_1}[n_1] \sum_{n_2=0}^{N_2-1} f[n_1, n_2]\,e^{2*}_{k_2}[n_2].
\end{aligned} \tag{3.77}$$

For $0 \leqslant n_1 < N_1$, we must compute

$$U f[n_1, k_2] = \sum_{n_2=0}^{N_2-1} f[n_1, n_2]\,e^{2*}_{k_2}[n_2],$$

which are the decomposition coefficients of the N_1 image rows of size N_2 in the basis $\mathcal{B}_2$. The coefficients $\{\langle f, e^1_{k_1} e^2_{k_2}\rangle\}_{0 \leqslant k_1 < N_1, 0 \leqslant k_2 < N_2}$ are calculated in (3.77) as the inner products of the columns of the transformed image $U f[n_1, k_2]$ in the basis $\mathcal{B}_1$. The overall algorithm thus requires to perform N_1 one-dimensional transforms in the basis $\mathcal{B}_2$ plus N_2 one-dimensional transforms in the basis $\mathcal{B}_1$ and hence requires $N_2\,C_1(N_1) + N_1\,C_2(N_2)$ operations.

The fast Fourier transform algorithm of Section 3.3.3 decomposes signals of size N_1 and N_2 in the discrete Fourier bases $\mathcal{B}_1 = \{e^1_{k_1}[n_1] = e^{i2\pi k_1 n_1/N_1}\}_{0 \leqslant k_1 < N_1}$ and $\mathcal{B}_2 = \{e^2_{k_2}[n_2] = e^{i2\pi k_2 n_2/N_2}\}_{0 \leqslant k_2 < N_2}$, with $C_1(N_1) = K N_1 \log_2 N_1$ and $C_2(N_2) = K N_2 \log_2 N_2$ operations. A separable implementation of a two-dimensional FFT thus requires $N_2\,C_1(N_1) + N_1\,C_2(N_2) = K N \log_2 N$ operations, with $N = N_1 N_2$. A split-radix FFT corresponds to $K = 3$.

3.5 Exercises

3.1. [1] Show that if $\phi_s(t) = s^{-1/2}\,\mathbf{1}_{[0,s)}(t)$ then $\{\phi_s(t - ns)\}_{n \in \mathbb{Z}}$ is an orthonormal basis of the space of piecewise constant function on intervals $[ns, (n+1)s)$, for any $n \in \mathbb{Z}$.

3.2. [2] Prove that if f has a Fourier transform included in $[-\pi/s, \pi/s]$, then

$$\forall u \in \mathbb{R} \quad, \quad f(u) = \frac{1}{s}\,\langle f(t), \phi_s(t-u)\rangle \quad \text{with} \quad \phi_s(t) = \frac{\sin(\pi t/s)}{\pi t/s}.$$

3.3. [2] An interpolation function $f(t)$ satisfies $f(n) = \delta[n]$ for any $n \in \mathbb{Z}$.
   (a) Prove that $\sum_{k=-\infty}^{+\infty} \hat{f}(\omega + 2k\pi) = 1$ if and only if f is an interpolation function.
   (b) Suppose that $f(t) = \sum_{n=-\infty}^{+\infty} h[n]\,\theta(t-n)$ with $\theta \in \mathbf{L}^2(\mathbb{R})$. Find $\hat{h}(\omega)$ as a function of $\hat{\theta}(\omega)$ so that $f(t)$ is an interpolation function. Relate $\hat{f}(\omega)$ to $\hat{\theta}(\omega)$, and give a sufficient condition on $\hat{\theta}$ to guarantee that $f \in \mathbf{L}^2(\mathbb{R})$.

3.4. [2] Prove that if $f \in \mathbf{L}^2(\mathbb{R})$ and $\sum_{n=-\infty}^{+\infty} f(t-n) \in \mathbf{L}^2[0,1]$ then

$$\sum_{n=-\infty}^{+\infty} f(t-n) = \sum_{k=-\infty}^{+\infty} \hat{f}(2k\pi)\, e^{i2\pi kt} .$$

3.5. [1] We want to approximate f by a signal $\tilde{f}$ in an approximation space $\mathbf{U}_s$. Prove that the approximation $\tilde{f}$ which minimizes $\|\tilde{f} - f\|$, is the orthogonal projection of f in $\mathbf{U}_s$.

3.6. [2] Prove that the discrete filter $Lf[n] = f \star h[n]$ is stable if and only if $h \in \ell^1(\mathbb{Z})$.

3.7. [2] If $\hat{h}(\omega)$ and $\hat{g}(\omega)$ are the Fourier transforms of $h[n]$ and $g[n]$, we write

$$\hat{h} \star \hat{g}(\omega) = \int_{-\pi}^{+\pi} \hat{h}(\xi)\, \hat{g}(\omega - \xi)\, d\xi.$$

Prove that if $f[n] = g[n]\, h[n]$ then $\hat{f}(\omega) = (2\pi)^{-1} \hat{h} \star \hat{g}(\omega)$.

3.8. [1] Prove that $\{e^{i2\pi kn/N}\}_{0 \leqslant k < N}$ is an orthogonal family and hence an orthogonal basis of $\mathbb{C}^N$. What renormalization factor is needed to obtain an orthonormal basis ?

3.9. [2] Suppose that $\hat{f}$ has a support in $[-(n+1)\pi/s, -n\pi/s] \cup [n\pi/s, (n+1)\pi/s]$ and that $f(t)$ is real. Find an interpolation formula that recovers $f(t)$ from $\{f(ns)\}_{n \in \mathbb{Z}}$.

3.10. [3] Suppose that $\hat{f}$ has a support in $[-\pi/s, \pi/s]$.

    (a) Give the filter $\phi_s(t)$ such that for any f

$$\forall n \in \mathbb{Z} \ , \quad \tilde{f}(ns) = \int_{(n-1/2)s}^{(n+1/2)s} f(t)\, dt = f \star \phi_s(ns).$$

    (b) Show that $\tilde{f}(t) = f \star \phi_s(t)$ can be recovered from $\{\tilde{f}(ns)\}_{n \in \mathbb{Z}}$ with an interpolation formula.
    (c) Reconstruct f from $\tilde{f}$ by inverting ϕ_s.
    (d) Prove that the reconstruction of $f(t)$ from $\{\tilde{f}(ns)\}_{n \in \mathbb{Z}}$ is stable.

3.11. [2] The linear box spline $\phi(t)$ is defined in (3.31) for $m = 1$.

    (a) Give an analytical formula for $\phi(t)$ and specify its support.
    (b) Prove with (7.19) that $\{\phi(t-n)\}_{n \in \mathbb{Z}}$ is a Riesz basis of the space of finite energy functions that are continuous and linear on intervals $[ns, (n+1)]$ for $n \in \mathbb{Z}$.
    (c) Does the dual filter $\tilde{\phi}(t)$ have a compact support ? Compute its graph numerically.

3.12. [1] If $f[n]$ is defined for $0 \leqslant n < N$ prove that $|\hat{f}[k]| \leqslant \sum_{n=0}^{N-1} |f[n]|$ for any $0 \leqslant< N$.

3.13. [2] The discrete and periodic total variation is

$$\|f\|_V = \sum_{n=0}^{N-1} |f[n] - f[n-1]| + |f[N-1] - f[0]|.$$

    (a) Prove that $\|f\|_V = \sum_{n=0}^{N-1} |f \circledast h[n]|$ where $h[n]$ is a filter and specify $\hat{h}[k]$.
    (b) Derive an upper bound of $|\hat{f}[k]|$ as a function of k^{-1}.

3.14. [1] Let $g[n] = (-1)^n\, h[n]$. Relate $\hat{g}(\omega)$ to $\hat{h}(\omega)$. If h is a low-pass filter, what kind of filter is g ?

3.15. [2] Prove the convolution Theorem 3.7.

3.16. [2] Let h^{-1} be the inverse of h defined by $h \star h^{-1}[n] = \delta[n]$.

    (a) Compute $\hat{h}^{-1}(\omega)$ as a function of $\hat{h}(\omega)$.
    (b) Prove that if h has a finite support then h^{-1} has a finite support if and only if $h[n] = \delta[n-p]$ for some $p \in \mathbb{Z}$.

3.17. [1] *All pass filters*

    (a) Verify that

$$\hat{h}(\omega) = \prod_{k=1}^{K} \frac{a_k^* - e^{-i\omega}}{1 + a_k\, e^{i\omega}}$$

        is an all-pass filter, i.e. $|\hat{h}(\omega)| = 1$.
    (b) Prove that $\{h[n-m]\}_{m \in \mathbb{Z}}$ is an orthonormal basis of $\ell^2(\mathbb{Z})$.

3.18. [2] *Recursive filters*

    (a) Compute the Fourier transform of $h[n] = a^n\, \mathbf{1}_{[0,+\infty)}[n]$ for $|a| < 1$. Compute the inverse Fourier transform of $\hat{h}(\omega) = (1 - a\, e^{-i\omega})^{-p}$.

(b) Suppose that $g = f \star h$ is calculated by a recursive equation with real coefficients

$$\sum_{k=0}^{K} a_k \, f[n-k] = \sum_{k=0}^{M} b_k \, g[n-k],$$

write $\hat{h}(\omega)$ as a function of the parameters a_k and b_k.

(c) Show that h is a stable filter if and only if the equation $\sum_{k=0}^{M} b_k \, z^{-k} = 0$ has roots with a modulus strictly smaller than 1.

3.19. [1] *Discrete interpolation* Let $\hat{f}[k]$ be the DFT of a signal $f[n]$ of size N. We define a signal $\tilde{f}[n]$ of size $2N$ by $\widehat{\tilde{f}}[N/2] = \widehat{\tilde{f}}[3N/2] = \hat{f}[N/2]$ and

$$\widehat{\tilde{f}}[k] = \begin{cases} 2\hat{f}[k] & \text{if } 0 \leqslant k < N/2 \\ 0 & \text{if } N/2 < k < 3N/2 \\ 2\hat{f}[k-N] & \text{if } 3N/2 < k < 2N \end{cases}.$$

Prove that $\tilde{f}$ is an interpolation of f which satisfies $\tilde{f}[2n] = f[n]$.

3.20. [2] *Decimation* Let $x[n] = y[Mn]$ with $M > 1$.

(a) Show that $\hat{x}(\omega) = M^{-1} \sum_{k=0}^{M-1} \hat{y}(M^{-1}(\omega - 2k\pi))$.

(b) Give a sufficient condition on $\hat{y}(\omega)$ to recover y from x and give the interpolation formula which recovers $y[n]$ from x.

3.21. [3] We want to compute numerically the Fourier transform of $f(t)$. Let $f_d[n] = f(ns)$, and $f_p[n] = \sum_{p=-\infty}^{+\infty} f_d[n - pN]$.

(a) Prove that the DFT of $f_p[n]$ is related to the Fourier series of $f_d[n]$ and to the Fourier transform of $f(t)$ by

$$\hat{f}_p[k] = \hat{f}_d\left(\frac{2\pi k}{N}\right) = \frac{1}{s} \sum_{l=-\infty}^{+\infty} \hat{f}\left(\frac{2k\pi}{Ns} - \frac{2l\pi}{s}\right).$$

(b) Suppose that $|f(t)|$ and $|\hat{f}(\omega)|$ are negligible when $t \notin [-t_0, t_0]$ and $\omega \notin [-\omega_0, \omega_0]$. Relate N and s to t_0 and ω_0 so that one can compute an approximate value of $\hat{f}(\omega)$ for all $\omega \in \mathbb{R}$ by interpolating the samples $\hat{f}_p[k]$. Is it possible to compute exactly $\hat{f}(\omega)$ with such an interpolation formula?

(c) Let $f(t) = \left(\sin(\pi t)/(\pi t)\right)^4$. What is the support of $\hat{f}$? Sample f appropriately and compute $\hat{f}$ numerically with an FFT algorithm.

3.22. [2] The analytic part $f_a[n]$ of a real discrete signal $f[n]$ of size N is defined by

$$\hat{f}_a[k] = \begin{cases} \hat{f}[k] & \text{if } k = 0, \, N/2 \\ 2\,\hat{f}[k] & \text{if } 0 < k < N/2 \\ 0 & \text{if } N/2 < k < N \end{cases}.$$

(a) Compute $f_a[n]$ for $f[n] = \cos(2\pi kn/N)$ with $0 < k < N/2$.

(b) Prove that the real part $g[n] = \mathrm{Re}(f[n])$ satisfies $\hat{g}[k] = (\hat{f}[k] + \hat{f}^*[-k])/2$.

(c) Prove that $\mathrm{Re}(f_a) = f$.

3.23. [1] Prove that if $\{e_{k_1}[n_1]\}_{0 \leqslant k_1 < N_1}$ is an orthonormal basis of $\mathbf{C}^{N_1}$ and $\{e_{k_2}[n_2]\}_{0 \leqslant k_1 < N_1}$ is an orthonormal basis of $\mathbf{C}^{N_2}$ then $\{e_{k_1}[n_1] \, e_{k_2}[n_2]\}_{0 \leqslant k_1 < N_1, 0 \leqslant k_2 < N_2}$ is orthogonal basis of the space $\mathbb{C}^N = \mathbb{C}^{N_1 N_2}$ of images $f[n_1, n_2]$ of $N = N_1 N_2$ pixels.

3.24. [2] Let $h[n_1, n_2]$ be a non-separable filter which is non-zero for $0 \leqslant n_1, n_2 < M$. Let $f[n_1, n_2]$ be a square image defined for $0 \leqslant n_1, n_2 \leqslant LM$ of $N = (LM)^2$ pixels. Describe an overlap-add algorithm to compute $g[n_1, n_2] = f \star h[n_1, n_2]$. By using an FFT that requires $KP \log P$ operators to compute the Fourier transform of an image of P pixels, how many operations does your algorithm require? If $K = 6$, for what range of M is it better to compute the convolution with a direct summation?

3.25. [2] Let $f[n_1, n_2, n_3]$ be a 3-dimensional signal of size $N = N_1 N_2 N_3$. The discrete Fourier transform is defined as a decomposition in a separable discrete Fourier basis. Give a separable algorithm which decomposes f in this basis with $KN \log N$ operations, by using a one-dimensional FFT algorithm that requires $KP \log P$ operations for a one-dimensional signal of size P.

IV

Time Meets Frequency

When we listen to music, we clearly "hear" the time variation of the sound "frequencies." These localized frequency events are not pure tones but packets of close frequencies. The properties of sounds are revealed by transforms that decompose signals over elementary functions that are well concentrated in time and frequency. Windowed Fourier transforms and wavelet transforms are two important classes of local time-frequency decompositions. Measuring the time variations of "instantaneous" frequencies illustrates the limitations imposed by the Heisenberg uncertainty. These instantaneous frequencies are detected as local maxima in windowed Fourier and wavelet dictionaries, and define a signal approximation support. Audio processing algorithms are implemented by modifying the geometry of this approximation support.

There is no unique definition of time-frequency energy density. All quadratic time-frequency distributions are related through the averaging of a single quadratic form called the Wigner-Ville distribution. This framework gives another perspective on windowed Fourier and wavelet transforms.

4.1   Time-Frequency Atoms

A linear time-frequency transform correlates the signal with a dictionary of waveforms that are well concentrated in time and in frequency. These waveforms are called *time-frequency atoms*. Let us consider a general dictionary of time-frequency atoms $\mathcal{D} = \{\phi_\gamma\}_{\gamma \in \Gamma}$, where γ might be a multi-index parameter. We suppose that $\phi_\gamma \in \mathbf{L}^2(\mathbb{R})$ and that $\|\phi_\gamma\| = 1$. The corresponding linear time-frequency transform of $f \in \mathbf{L}^2(\mathbb{R})$ is defined by

$$\Phi f(\gamma) = \int_{-\infty}^{+\infty} f(t)\, \phi_\gamma^*(t)\, dt = \langle f, \phi_\gamma \rangle.$$

The Parseval formula (2.25) proves that

$$\Phi f(\gamma) = \int_{-\infty}^{+\infty} f(t)\, \phi_\gamma^*(t)\, dt = \frac{1}{2\pi} \int_{-\infty}^{+\infty} \hat{f}(\omega)\, \hat{\phi}_\gamma^*(\omega)\, d\omega. \tag{4.1}$$

If $\phi_\gamma(t)$ is nearly zero when t is outside a neighborhood of an abscissa u, then $\langle f, \phi_\gamma \rangle$ depends only on the values of f in this neighborhood. Similarly, if $\hat{\phi}_\gamma(\omega)$ is negligible for ω far from ξ, then the right integral of (4.1) proves that $\langle f, \phi_\gamma \rangle$ reveals the properties of $\hat{f}$ in the neighborhood of ξ.

Heisenberg Boxes   The slice of information provided by $\langle f, \phi_\gamma \rangle$ is represented in a time-frequency plane (t, ω) by a region whose location and width depends on the time-frequency spread of ϕ_γ. Since

$$\|\phi_\gamma\|^2 = \int_{-\infty}^{+\infty} |\phi_\gamma(t)|^2\, dt = 1,$$

we interpret $|\phi_\gamma(t)|^2$ as a probability distribution centered at

$$u_\gamma = \int_{-\infty}^{+\infty} t\,|\phi_\gamma(t)|^2\,dt. \tag{4.2}$$

The spread around u_γ is measured by the variance

$$\sigma_t^2(\gamma) = \int_{-\infty}^{+\infty} (t - u_\gamma)^2\,|\phi_\gamma(t)|^2\,dt. \tag{4.3}$$

The Plancherel formula (2.26) proves that $\int_{-\infty}^{+\infty} |\hat{\phi}_\gamma(\omega)|^2\,d\omega = 2\pi\|\phi_\gamma\|^2$. The center frequency of $\hat{\phi}_\gamma$ is therefore defined by

$$\xi_\gamma = \frac{1}{2\pi} \int_{-\infty}^{+\infty} \omega\,|\hat{\phi}_\gamma(\omega)|^2\,d\omega, \tag{4.4}$$

and its spread around ξ_γ is

$$\sigma_\omega^2(\gamma) = \frac{1}{2\pi} \int_{-\infty}^{+\infty} (\omega - \xi_\gamma)^2\,|\hat{\phi}_\gamma(\omega)|^2\,d\omega. \tag{4.5}$$

The time-frequency resolution of ϕ_γ is represented in the time-frequency plane (t, ω) by a Heisenberg box centered at (u_γ, ξ_γ), whose width along time is $\sigma_t(\gamma)$ and whose width along frequency is $\sigma_\omega(\gamma)$. This is illustrated by Figure 4.1. The Heisenberg uncertainty Theorem 2.6 proves that the area of the rectangle is at least $1/2$:

$$\sigma_t\,\sigma_\omega \geqslant \frac{1}{2}. \tag{4.6}$$

This limits the joint resolution of ϕ_γ in time and frequency. The time-frequency plane must be manipulated carefully because a point (t_0, ω_0) is ill-defined. There is no function that is perfectly well concentrated at a point t_0 and a frequency ω_0. Only rectangles with area at least $1/2$ may correspond to time-frequency atoms.
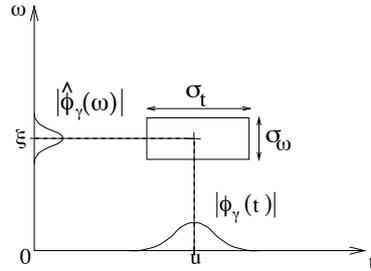
Figure 4.1: Heisenberg box representing an atom ϕ_γ.

Translation Invariant Dictionaries   For pattern recognition, it can be important to construct signal representations that are translation invariant. When a pattern is translated, its numerical descriptors are then translated but not modified. Observe that for any $\phi_\gamma \in \mathcal{D}$ and any shift u

$$\langle f(t-u), \phi_\gamma(t) \rangle = \langle f(t), \phi_\gamma(t+u) \rangle\,.$$

A translation invariant representation is thus obtained if $\phi_\gamma(t+u)$ is in $\mathcal{D}$ up to a multiplicative constant. Such a dictionary is said to be translation invariant.

A translation invariant dictionary is obtained by translating a family of generators $\{\phi_\gamma\}_{\gamma \in \Gamma}$, and can be written $\mathcal{D} = \{\phi_{u,\gamma}\}_{\gamma \in \Gamma, u \in \mathbb{R}}$, with $\phi_{u,\gamma}(t) = \lambda_{u,\gamma}\,\phi_\gamma(t-u)$. The resulting time-frequency transform of f can then be written as a convolution:

$$\Phi f(u, \gamma) = \langle f, \phi_{u,\gamma} \rangle = \int_{-\infty}^{+\infty} f(t)\,\lambda_{u,\gamma}\phi_\gamma^*(t-u)\,dt = \lambda_{u,\gamma}\,f \star \tilde{\phi}_\gamma(u)$$

with $\tilde{\phi}_\gamma(t) = \phi_\gamma^*(-t)$.

Energy Density Let us suppose that $\phi_\gamma(t)$ is centered at $t = 0$ so that $\phi_{u,\gamma}(t)$ is centered at u. Let ξ_γ be the center frequency of $\hat{\phi}_\gamma(\omega)$ defined in (4.4). The time-frequency box of $\phi_{u,\gamma}$ specifies a neighborhood of (u, ξ_γ) where the energy of f is measured by

$$P_\Phi f(u, \xi_\gamma) = |\langle f, \phi_{u,\gamma}\rangle|^2 = \left| \int_{-\infty}^{+\infty} f(t)\, \phi_{u,\gamma}^*(t)\, dt \right|^2. \tag{4.7}$$

Section 4.5.1 proves that any such energy density is an averaging of the Wigner-Ville distribution, with a kernel that depends on the atoms $\phi_{u,\gamma}$.

Example 4.1. *A windowed Fourier atom is constructed with a window g modulated by the frequency ξ and translated by u:*

$$\phi_{u,\gamma}(t) = g_{u,\xi}(t) = \mathrm{e}^{i\xi t}\, g(t - u). \tag{4.8}$$

The resulting window Fourier dictionary $\mathcal{D} = \{g_{u,\xi}(t)\}_{u,\xi \in \mathbb{R}^2}$ is translation invariant since $g_{u,\xi} = \mathrm{e}^{i\xi u}\, g_{0,\xi}(t - u)$. A windowed Fourier dictionary is also frequency shit invariant because

$$\mathrm{e}^{i\omega t}\, g_{u,\xi}(t) = g_{u,\xi+\omega}(t) \in \mathcal{D} .$$

This dictionary is thus particularly useful to analyze patterns that are translated in time and frequency.

A wavelet atom is a dilation by s and a translation by u of a mother wavelet ψ:

$$\phi_{u,\gamma}(t) = \psi_{u,s}(t) = \frac{1}{\sqrt{s}}\, \psi\left(\frac{t - u}{s}\right). \tag{4.9}$$

A wavelet dictionary $\mathcal{D} = \{\psi_{u,s}(t)\}_{u \in \mathbb{R}, s \in \mathbb{R}^+}$ translation invariant but also scale invariant, because scaling any wavelet produces a dilated wavelet that remains in the dictionary. A wavelet dictionary is well adapted to analyze patterns that are translated and scaled by arbitrary factors.

Wavelets and windowed Fourier atoms have their energy well localized in time, while their Fourier transform is mostly concentrated in a limited frequency band. The properties of the resulting transforms are studied in Sections 4.2 and 4.3.

4.2 Windowed Fourier Transform

In 1946, Gabor [266] introduced windowed Fourier atoms to measure the "frequency variations" of sounds. A real and symmetric window $g(t) = g(-t)$ is translated by u and modulated by the frequency ξ:

$$g_{u,\xi}(t) = \mathrm{e}^{i\xi t} g(t - u). \tag{4.10}$$

It is normalized $\|g\| = 1$ so that $\|g_{u,\xi}\| = 1$ for any $(u, \xi) \in \mathbb{R}^2$. The resulting windowed Fourier transform of $f \in \mathbf{L}^2(\mathbb{R})$ is

$$Sf(u, \xi) = \langle f, g_{u,\xi}\rangle = \int_{-\infty}^{+\infty} f(t)\, g(t - u)\, \mathrm{e}^{-i\xi t}\, dt. \tag{4.11}$$

This transform is also called the *short time Fourier transform* because the multiplication by $g(t-u)$ localizes the Fourier integral in the neighborhood of $t = u$.

As in (4.7), one can define an energy density called a *spectrogram*, denoted P_S:

$$P_S f(u, \xi) = |Sf(u, \xi)|^2 = \left| \int_{-\infty}^{+\infty} f(t)\, g(t - u)\, \mathrm{e}^{-i\xi t}\, dt \right|^2. \tag{4.12}$$

The spectrogram measures the energy of f in a time-frequency neighborhood of (u, ξ) specified by the Heisenberg box of $g_{u,\xi}$.

Heisenberg Boxes   Since g is even, $g_{u,\xi}(t) = \mathrm{e}^{i\xi t} g(t - u)$ is centered at u. The time spread around u is independent of u and ξ:

$$\sigma_t^2 = \int_{-\infty}^{+\infty} (t - u)^2 \, |g_{u,\xi}(t)|^2 \, dt = \int_{-\infty}^{+\infty} t^2 \, |g(t)|^2 \, dt. \tag{4.13}$$

The Fourier transform $\hat{g}$ of g is real and symmetric because g is real and symmetric. The Fourier transform of $g_{u,\xi}$ is

$$\hat{g}_{u,\xi}(\omega) = \hat{g}(\omega - \xi) \, \exp[-iu(\omega - \xi)]. \tag{4.14}$$

It is a translation by ξ of the frequency window $\hat{g}$, so its center frequency is ξ. The frequency spread around ξ is

$$\sigma_\omega^2 = \frac{1}{2\pi} \int_{-\infty}^{+\infty} (\omega - \xi)^2 \, |\hat{g}_{u,\xi}(\omega)|^2 \, d\omega = \frac{1}{2\pi} \int_{-\infty}^{+\infty} \omega^2 \, |\hat{g}(\omega)|^2 \, d\omega. \tag{4.15}$$

It is independent of u and ξ. Hence $g_{u,\xi}$ corresponds to a Heisenberg box of area $\sigma_t \, \sigma_\omega$ centered at (u, ξ), as illustrated by Figure 4.2. The size of this box is independent of (u, ξ), which means that a windowed Fourier transform has the same resolution across the time-frequency plane.
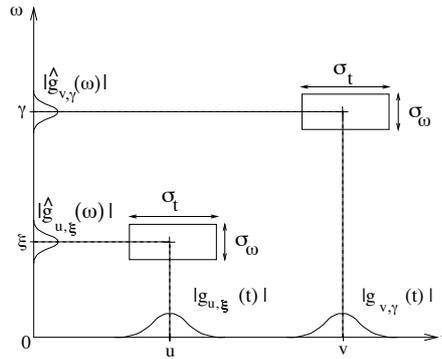
Figure 4.2: Heisenberg boxes of two windowed Fourier atoms $g_{u,\xi}$ and $g_{\nu,\gamma}$.

Example 4.2. *A sinusoidal wave $f(t) = \exp(i\xi_0 t)$ whose Fourier transform is a Dirac $\hat{f}(\omega) = 2\pi\delta(\omega - \xi_0)$ has a windowed Fourier transform*

$$Sf(u, \xi) = \hat{g}(\xi - \xi_0) \, \exp[-iu(\xi - \xi_0)].$$

Its energy is spread over the frequency interval $[\xi_0 - \sigma_\omega/2, \xi_0 + \sigma_\omega/2]$.

Example 4.3. *The windowed Fourier transform of a Dirac $f(t) = \delta(t - u_0)$ is*

$$Sf(u, \xi) = g(u_0 - u) \, \exp(-i\xi u_0).$$

Its energy is spread in the time interval $[u_0 - \sigma_t/2, u_0 + \sigma_t/2]$.

Example 4.4. *A linear chirp $f(t) = \exp(iat^2)$ has an "instantaneous frequency" that increases linearly in time. For a Gaussian window $g(t) = (\pi\sigma^2)^{-1/4} \exp[-t^2/(2\sigma^2)]$, the windowed Fourier transform of f is calculated using the Fourier transform (2.34) of Gaussian chirps. One can verify that its spectrogram is*

$$P_S f(u, \xi) = |Sf(u, \xi)|^2 = \left(\frac{4\pi\sigma^2}{1 + 4a^2\sigma^4} \right)^{1/2} \exp\left(-\frac{\sigma^2(\xi - 2au)^2}{1 + 4a^2\sigma^4} \right).$$

$$\tag{4.16}$$

For a fixed time u, $P_S f(u, \xi)$ is a Gaussian that reaches its maximum at the frequency $\xi(u) = 2au$. Observe that if we write $f(t) = \exp[i\phi(t)]$, then $\xi(u)$ is equal to the "instantaneous frequency," defined as the derivative of the phase: $\omega(u) = \phi'(u) = 2au$. Section 4.4.1 explains this result.
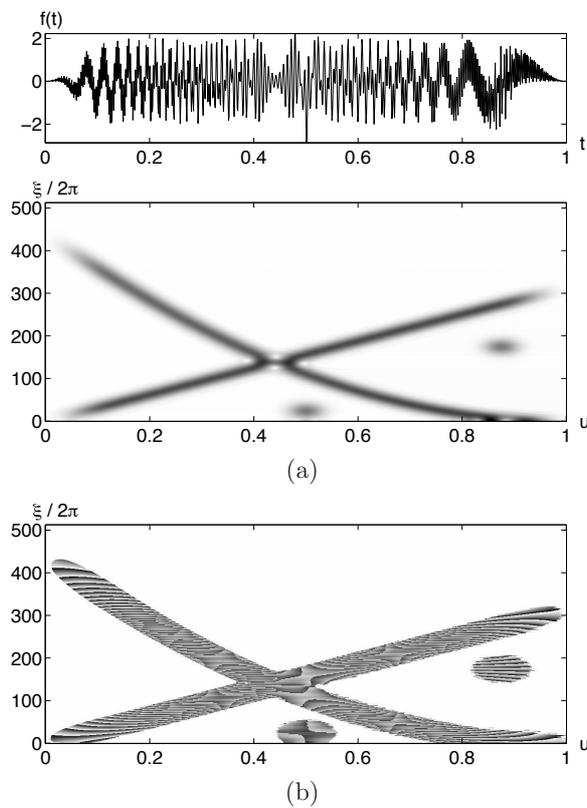
Figure 4.3: The signal includes a linear chirp whose frequency increases, a quadratic chirp whose frequency decreases, and two modulated Gaussian functions located at $t = 0.5$ and $t = 0.87$. (a) Spectrogram $P_S f(u, \xi)$. Dark points indicate large amplitude coefficients. (b) Complex phase of $Sf(u, \xi)$ in regions where the modulus $P_S f(u, \xi)$ is non-zero.

Example 4.5. *Figure 4.3 gives the spectrogram of a signal that includes a linear chirp, a quadratic chirp and two modulated Gaussians. The spectrogram is computed with a Gaussian window dilated by $\sigma = 0.05$. As expected from (4.16), the linear chirp yields large amplitude coefficients along the trajectory of its instantaneous frequency, which is a straight line. The quadratic chirp yields large coefficients along a parabola. The two modulated Gaussians produce low and high frequency blobs at $u = 0.5$ and $u = 0.87$.*

4.2.1   Completeness and Stability

When the time-frequency indices (u, ξ) vary across $\mathbb{R}^2$, the Heisenberg boxes of the atoms $g_{u,\xi}$ cover the whole time-frequency plane. One can thus expect that f can be recovered from its windowed Fourier transform $Sf(u, \xi)$. The following theorem gives a reconstruction formula and proves that the energy is conserved.

Theorem 4.1. *If $f \in \mathbf{L}^2(\mathbb{R})$ then*

$$f(t) = \frac{1}{2\pi} \int_{-\infty}^{+\infty} \int_{-\infty}^{+\infty} Sf(u, \xi)\, g(t - u)\, e^{i\xi t}\, d\xi\, du \tag{4.17}$$

and

$$\int_{-\infty}^{+\infty} |f(t)|^2\, dt = \frac{1}{2\pi} \int_{-\infty}^{+\infty} \int_{-\infty}^{+\infty} |Sf(u, \xi)|^2\, d\xi\, du. \tag{4.18}$$

Proof. The reconstruction formula (4.17) is proved first. Let us apply the Fourier Parseval formula (2.25) to the integral (4.17) with respect to the integration in u. The Fourier transform of $f_\xi(u) = Sf(u, \xi)$ with respect to u is computed by observing that

$$Sf(u, \xi) = \exp(-iu\xi) \int_{-\infty}^{+\infty} f(t)\, g(t - u)\, \exp[i\xi(u - t)]\, dt = \exp(-iu\xi)\, f \star g_\xi(u),$$

where $g_\xi(t) = g(t) \exp(i\xi t)$, because $g(t) = g(-t)$. Its Fourier transform is therefore

$$\hat{f}_\xi(\omega) = \hat{f}(\omega + \xi)\, \hat{g}_\xi(\omega + \xi) = \hat{f}(\omega + \xi)\, \hat{g}(\omega).$$

The Fourier transform of $g(t - u)$ with respect to u is $\hat{g}(\omega) \exp(-it\omega)$. Hence

$$\frac{1}{2\pi} \left(\int_{-\infty}^{+\infty} \int_{-\infty}^{+\infty} Sf(u, \xi)\, g(t - u)\, \exp(i\xi t)\, du \right) d\xi \quad =$$

$$\frac{1}{2\pi} \int_{-\infty}^{+\infty} \left(\frac{1}{2\pi} \int_{-\infty}^{+\infty} \hat{f}(\omega + \xi)\, |\hat{g}(\omega)|^2\, \exp[it(\omega + \xi)]\, d\omega \right) d\xi \quad .$$

If $\hat{f} \in \mathbf{L}^1(\mathbb{R})$, we can apply the Fubini Theorem A.2 to reverse the integration order. The inverse Fourier transform proves that

$$\frac{1}{2\pi} \int_{-\infty}^{+\infty} \hat{f}(\omega + \xi)\, \exp[it(\omega + \xi)]\, d\xi = f(t).$$

Since $\frac{1}{2\pi} \int_{-\infty}^{+\infty} |\hat{g}(\omega)|^2\, d\omega = 1$, we derive (4.17). If $\hat{f} \notin \mathbf{L}^1(\mathbb{R})$, a density argument is used to verify this formula.

Let us now prove the energy conservation (4.18). Since the Fourier transform in u of $Sf(u, \xi)$ is $\hat{f}(\omega + \xi)\, \hat{g}(\omega)$, the Plancherel formula (2.26) applied to the right-hand side of (4.18) gives

$$\frac{1}{2\pi} \int_{-\infty}^{+\infty} \int_{-\infty}^{+\infty} |Sf(u, \xi)|^2\, du\, d\xi = \frac{1}{2\pi} \int_{-\infty}^{+\infty} \frac{1}{2\pi} \int_{-\infty}^{+\infty} |\hat{f}(\omega + \xi)\, \hat{g}(\omega)|^2\, d\omega\, d\xi.$$

The Fubini theorem applies and the Plancherel formula proves that

$$\frac{1}{2\pi} \int_{-\infty}^{+\infty} |\hat{f}(\omega + \xi)|^2\, d\xi = \|f\|^2,$$

which implies (4.18).   ■   ■

The reconstruction formula (4.17) can be rewritten

$$f(t) = \frac{1}{2\pi} \int_{-\infty}^{+\infty} \int_{-\infty}^{+\infty} \langle f, g_{u,\xi} \rangle \, g_{u,\xi}(t) \, d\xi \, du. \tag{4.19}$$

It resembles the decomposition of a signal in an orthonormal basis but it is not, since the functions $\{g_{u,\xi}\}_{u,\xi \in \mathbb{R}^2}$ are very redundant in $\mathbf{L}^2(\mathbb{R})$. The second equality (4.18) justifies the interpretation of the spectrogram $P_S f(u,\xi) = |Sf(u,\xi)|^2$ as an energy density, since its time-frequency sum equals the signal energy.

Reproducing Kernel   A windowed Fourier transform represents a one-dimensional signal $f(t)$ by a two-dimensional function $Sf(u,\xi)$. The energy conservation proves that $Sf \in \mathbf{L}^2(\mathbb{R}^2)$. Because $Sf(u,\xi)$ is redundant, it is not true that any $\Phi \in \mathbf{L}^2(\mathbb{R}^2)$ is the windowed Fourier transform of some $f \in \mathbf{L}^2(\mathbb{R})$. The next theorem gives a necessary and sufficient condition for such a function to be a windowed Fourier transform.

Theorem 4.2. *Let $\Phi \in \mathbf{L}^2(\mathbb{R}^2)$. There exists $f \in \mathbf{L}^2(\mathbb{R})$ such that $\Phi(u,\xi) = Sf(u,\xi)$ if and only if*

$$\Phi(u_0, \xi_0) = \frac{1}{2\pi} \int_{-\infty}^{+\infty} \int_{-\infty}^{+\infty} \Phi(u,\xi) \, K(u_0, u, \xi_0, \xi) \, du \, d\xi, \tag{4.20}$$

with

$$K(u_0, u, \xi_0, \xi) = \langle g_{u,\xi}, g_{u_0,\xi_0} \rangle \, . \tag{4.21}$$

Proof. Suppose that there exists f such that $\Phi(u,\xi) = Sf(u,\xi)$. Let us replace f with its reconstruction integral (4.17) in the windowed Fourier transform definition:

$$Sf(u_0, \xi_0) = \int_{-\infty}^{+\infty} \left(\frac{1}{2\pi} \int_{-\infty}^{+\infty} \int_{-\infty}^{+\infty} Sf(u,\xi) \, g_{u,\xi}(t) \, du \, d\xi \right) g_{u_0,\xi_0}^*(t) \, dt. \tag{4.22}$$

Inverting the integral on t with the integrals on u and ξ yields (4.20). To prove that the condition (4.20) is sufficient, we define f as in the reconstruction formula (4.17):

$$f(t) = \frac{1}{2\pi} \int_{-\infty}^{+\infty} \int_{-\infty}^{+\infty} \Phi(u,\xi) \, g(t-u) \, \exp(i\xi t) \, d\xi \, du$$

and show that (4.20) implies that $\Phi(u,\xi) = Sf(u,\xi)$.   ■   ■

Ambiguity Function   The reproducing kernel $K(u_0, u, \xi_0, \xi)$ measures the time-frequency overlap of the two atoms $g_{u,\xi}$ and g_{u_0,ξ_0}. The amplitude of $K(u_0, u, \xi_0, \xi)$ decays with $u_0 - u$ and $\xi_0 - \xi$ at a rate that depends on the energy concentration of g and $\hat{g}$. Replacing $g_{u,\xi}$ and g_{u_0,ξ_0} by their expression and making the change of variable $v = t - (u + u_0)/2$ in the inner product integral (4.21) yields

$$K(u_0, u, \xi_0, \xi) = \exp\left(-\frac{i}{2}(\xi_0 - \xi)(u + u_0) \right) Ag(u_0 - u, \xi_0 - \xi) \tag{4.23}$$

where

$$Ag(\tau, \gamma) = \int_{-\infty}^{+\infty} g\left(v + \frac{\tau}{2} \right) g\left(v - \frac{\tau}{2} \right) e^{-i\gamma v} \, dv \tag{4.24}$$

is called the *ambiguity function* of g. Using the Parseval formula to replace this time integral with a Fourier integral gives

$$Ag(\tau, \gamma) = \frac{1}{2\pi} \int_{-\infty}^{+\infty} \hat{g}\left(\omega + \frac{\gamma}{2} \right) \hat{g}\left(\omega - \frac{\gamma}{2} \right) e^{i\tau\omega} \, d\omega. \tag{4.25}$$

The decay of the ambiguity function measures the spread of g in time and of $\hat{g}$ in frequency. For example, if g has a support included in an interval of size T, then $Ag(\tau, \omega) = 0$ for $|\tau| \geqslant T/2$. The integral (4.25) shows that the same result applies to the support of $\hat{g}$.

4.2.2 Choice of Window

The resolution in time and frequency of the windowed Fourier transform depends on the spread of the window in time and frequency. This can be measured from the decay of the ambiguity function (4.24) or more simply from the area $\sigma_t \, \sigma_\omega$ of the Heisenberg box. The uncertainty Theorem 2.6 proves that this area reaches the minimum value $1/2$ if and only if g is a Gaussian. The ambiguity function $Ag(\tau, \gamma)$ is then a two-dimensional Gaussian.

Window Scale The time-frequency localization of g can be modified with a scaling. Suppose that g has a Heisenberg time and frequency width respectively equal to σ_t and σ_ω. Let $g_s(t) = s^{-1/2} g(t/s)$ be its dilation by s. A change of variables in the integrals (4.13) and (4.15) shows that the Heisenberg time and frequency width of g_s are respectively $s\sigma_t$ and σ_ω/s. The area of the Heisenberg box is not modified but it is dilated by s in time and compressed by s in frequency. Similarly, a change of variable in the ambiguity integral (4.24) shows that the ambiguity function is dilated in time and frequency respectively by s and $1/s$

$$Ag_s(\tau, \gamma) = Ag\left(\frac{\tau}{s}, s\gamma\right).$$

The choice of a particular scale s depends on the desired resolution trade-off between time and frequency.

Finite Support In numerical applications, g must have a compact support. Theorem 2.7 proves that its Fourier transform $\hat{g}$ necessarily has an infinite support. It is a symmetric function with a main lobe centered at $\omega = 0$, which decays to zero with oscillations. Figure 4.4 illustrates its behavior. To maximize the frequency resolution of the transform, we must concentrate the energy of $\hat{g}$ near $\omega = 0$. Three important parameters evaluate the spread of $\hat{g}$:

- The root mean-square bandwidth $\Delta\omega$, which is defined by

$$\frac{|\hat{g}(\Delta\omega/2)|^2}{|\hat{g}(0)|^2} = \frac{1}{2}.$$

- The maximum amplitude A of the first side-lobes located at $\omega = \pm\omega_0$ in Figure 4.4. It is measured in decibels:

$$A = 10\log_{10}\frac{|\hat{g}(\omega_0)|^2}{|\hat{g}(0)|^2}.$$

- The polynomial exponent p, which gives the asymptotic decay of $|\hat{g}(\omega)|$ for large frequencies:

$$|\hat{g}(\omega)| = O(\omega^{-p-1}). \tag{4.26}$$

Table 4.1 gives the values of these three parameters for several windows g whose supports are restricted to $[-1/2, 1/2]$ [292]. Figure 4.5 shows the graph of these windows.
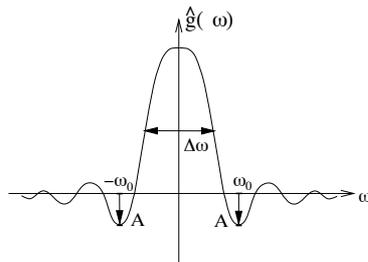
Figure 4.4: The energy spread of $\hat{g}$ is measured by its bandwidth $\Delta\omega$ and the maximum amplitude A of the first side-lobes, located at $\omega = \pm\omega_0$.

Name	$g(t)$	$\Delta\omega$	A	p
Rectangle	1	0.89	-13db	0
Hamming	$0.54 + 0.46\cos(2\pi t)$	1.36	-43db	0
Gaussian	$\exp(-18t^2)$	1.55	-55db	0
Hanning	$\cos^2(\pi t)$	1.44	-32db	2
Blackman	$0.42 + 0.5\cos(2\pi t)$ $+0.08\cos(4\pi t)$	1.68	-58db	2

Table 4.1: Frequency parameters of five windows g whose supports are restricted to $[-1/2, 1/2]$. These windows are normalized so that $g(0) = 1$ but $\|g\| \neq 1$.

To interpret these three frequency parameters, let us consider the spectrogram of a frequency tone $f(t) = \exp(i\xi_0 t)$. If $\Delta\omega$ is small, then $|Sf(u,\xi)|^2 = |\hat{g}(\xi - \xi_0)|^2$ has an energy concentrated near $\xi = \xi_0$. The side-lobes of $\hat{g}$ create "shadows" at $\xi = \xi_0 \pm \omega_0$, which can be neglected if A is also small.

If the frequency tone is embedded in a signal that has other components of much higher energy at different frequencies, the tone can still be detected if $\hat{g}(\omega - \xi)$ attenuates these components rapidly when $|\omega - \xi|$ increases. This means that $|\hat{g}(\omega)|$ has a rapid decay, and Theorem 2.5 proves that this decay depends on the regularity of g. Property (4.26) is typically satisfied by windows that are p times differentiable.
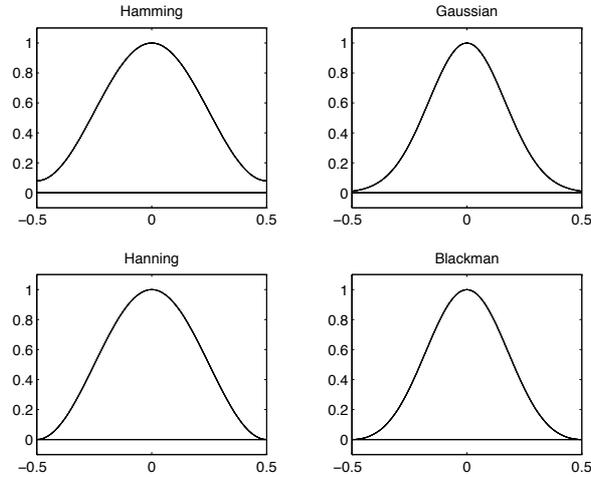
Figure 4.5: Graphs of four windows g whose support are $[-1/2, 1/2]$.

4.2.3 Discrete Windowed Fourier Transform

The discretization and fast computation of the windowed Fourier transform follow the same ideas as the discretization of the Fourier transform described in Section 3.3. We consider discrete signals of period N. The window $g[n]$ is chosen to be a symmetric discrete signal of period N with unit norm $\|g\| = 1$. Discrete windowed Fourier atoms are defined by

$$g_{m,l}[n] = g[n - m] \exp\left(\frac{i2\pi ln}{N}\right).$$

The discrete Fourier transform of $g_{m,l}$ is

$$\hat{g}_{m,l}[k] = \hat{g}[k - l] \exp\left(\frac{-i2\pi m(k - l)}{N}\right).$$

The discrete windowed Fourier transform of a signal f of period N is

$$Sf[m,l] = \langle f, g_{m,l} \rangle = \sum_{n=0}^{N-1} f[n]\, g[n-m]\, \exp\left(\frac{-i2\pi ln}{N}\right),\tag{4.27}$$

For each $0 \leqslant m < N$, $Sf[m,l]$ is calculated for $0 \leqslant l < N$ with a discrete Fourier transform of $f[n]g[n-m]$. This is performed with N FFT procedures of size N, and thus requires a total of $O(N^2 \log_2 N)$ operations. Figure 4.3 is computed with this algorithm.

Inverse Transform   The following theorem discretizes the reconstruction formula and the energy conservation of Theorem 4.1.

Theorem 4.3. *If f is a signal of period N then*

$$f[n] = \frac{1}{N} \sum_{m=0}^{N-1} \sum_{l=0}^{N-1} Sf[m,l]\, g[n-m]\, \exp\left(\frac{i2\pi ln}{N}\right)\tag{4.28}$$

and

$$\sum_{n=0}^{N-1} |f[n]|^2 = \frac{1}{N} \sum_{l=0}^{N-1} \sum_{m=0}^{N-1} |Sf[m,l]|^2.\tag{4.29}$$

This theorem is proved by applying the Parseval and Plancherel formulas of the discrete Fourier transform, exactly as in the proof of Theorem 4.1. The energy conservation (4.29) proves that this windowed Fourier transform defines a tight frame, as explained in Chapter 5. The reconstruction formula (4.28) is rewritten

$$f[n] = \frac{1}{N} \sum_{m=0}^{N-1} g[n-m] \sum_{l=0}^{N-1} Sf[m,l]\, \exp\left(\frac{i2\pi ln}{N}\right).$$

The second sum computes for each $0 \leqslant m < N$ the inverse discrete Fourier transform of $Sf[m,l]$ with respect to l. This is calculated with N FFT procedures, requiring a total of $O(N^2 \log_2 N)$ operations.

A discrete windowed Fourier transform is an N^2 image $Sf[l,m]$ that is very redundant, since it is entirely specified by a signal f of size N. The redundancy is characterized by a discrete reproducing kernel equation, which is the discrete equivalent of (4.20).

4.3   Wavelet Transforms

To analyze signal structures of very different sizes, it is necessary to use time-frequency atoms with different time supports. The wavelet transform decomposes signals over dilated and translated wavelets. A wavelet is a function $\psi \in \mathbf{L^2}(\mathbb{R})$ with a zero average:

$$\int_{-\infty}^{+\infty} \psi(t)\, dt = 0.\tag{4.30}$$

It is normalized $\|\psi\| = 1$, and centered in the neighborhood of $t = 0$. A dictionary of time-frequency atoms is obtained by scaling ψ by s and translating it by u:

$$\mathcal{D} = \left\{ \psi_{u,s}(t) = \frac{1}{\sqrt{s}}\, \psi\left(\frac{t-u}{s}\right) \right\}_{u \in \mathbb{R}, s \in \mathbb{R}^+}.$$

These atoms remain normalized: $\|\psi_{u,s}\| = 1$. The wavelet transform of $f \in \mathbf{L^2}(\mathbb{R})$ at time u and scale s is

$$Wf(u,s) = \langle f, \psi_{u,s} \rangle = \int_{-\infty}^{+\infty} f(t)\, \frac{1}{\sqrt{s}}\, \psi^*\left(\frac{t-u}{s}\right) dt.\tag{4.31}$$

Linear Filtering   The wavelet transform can be rewritten as a convolution product:

$$W f(u, s) = \int_{-\infty}^{+\infty} f(t) \frac{1}{\sqrt{s}} \psi^* \left(\frac{t - u}{s} \right) dt = f \star \bar{\psi}_s(u) \tag{4.32}$$

with

$$\bar{\psi}_s(t) = \frac{1}{\sqrt{s}} \psi^* \left(\frac{-t}{s} \right).$$

The Fourier transform of $\bar{\psi}_s(t)$ is

$$\widehat{\bar{\psi}_s}(\omega) = \sqrt{s} \, \hat{\psi}^*(s\omega). \tag{4.33}$$

Since $\hat{\psi}(0) = \int_{-\infty}^{+\infty} \psi(t) \, dt = 0$, it appears that $\hat{\psi}$ is the transfer function of a band-pass filter. The convolution (4.32) computes the wavelet transform with dilated band-pass filters.

Analytic Versus Real Wavelets   Like a windowed Fourier transform, a wavelet transform can measure the time evolution of frequency transients. This requires using a complex analytic wavelet, which can separate amplitude and phase components. The properties of this analytic wavelet transform are described in Section 4.3.2, and its application to the measurement of instantaneous frequencies is explained in Section 4.4.2. In contrast, real wavelets are often used to detect sharp signal transitions. Section 4.3.1 introduces elementary properties of real wavelets, which are developed in Chapter 6.

4.3.1   Real Wavelets

Suppose that ψ is a real wavelet. Since it has a zero average, the wavelet integral

$$W f(u, s) = \int_{-\infty}^{+\infty} f(t) \frac{1}{\sqrt{s}} \psi^* \left(\frac{t - u}{s} \right) dt$$

measures the variation of f in a neighborhood of u, whose size is proportional to s. Section 6.1.3 proves that when the scale s goes to zero, the decay of the wavelet coefficients characterizes the regularity of f in the neighborhood of u. This has important applications for detecting transients and analyzing fractals. This section concentrates on the completeness and redundancy properties of real wavelet transforms.

Example 4.6. *Wavelets equal to the second derivative of a Gaussian are called Mexican hats. They were first used in computer vision to detect multiscale edges [487]. The normalized Mexican hat wavelet is*

$$\psi(t) = \frac{2}{\pi^{1/4} \sqrt{3\sigma}} \left(\frac{t^2}{\sigma^2} - 1 \right) \exp \left(\frac{-t^2}{2\sigma^2} \right). \tag{4.34}$$

For $\sigma = 1$, Figure 4.6 plots $-\psi$ and its Fourier transform

$$\hat{\psi}(\omega) = \frac{-\sqrt{8} \, \sigma^{5/2} \, \pi^{1/4}}{\sqrt{3}} \, \omega^2 \, \exp \left(\frac{-\sigma^2 \omega^2}{2} \right). \tag{4.35}$$

Figure 4.7 shows the wavelet transform of a signal that is piecewise regular on the left and almost everywhere singular on the right. The maximum scale is smaller than 1 because the support of f is normalized to $[0, 1]$. The minimum scale is limited by the sampling interval of the discretized signal used in numerical calculations. When the scale decreases, the wavelet transform has a rapid decay to zero in the regions where the signal is regular. The isolated singularities on the left create cones of large amplitude wavelet coefficients that converge to the locations of the singularities. This is further explained in Chapter 6.

A real wavelet transform is complete and maintains an energy conservation, as long as the wavelet satisfies a weak admissibility condition, specified by the following theorem. This theorem was first proved in 1964 by the mathematician Calderón [131], from a different point of view. Wavelets did not appear as such, but Calderón defines a wavelet transform as a convolution operator that decomposes the identity. Grossmann and Morlet [287] were not aware of Calderón's work when they proved the same formula for signal processing.
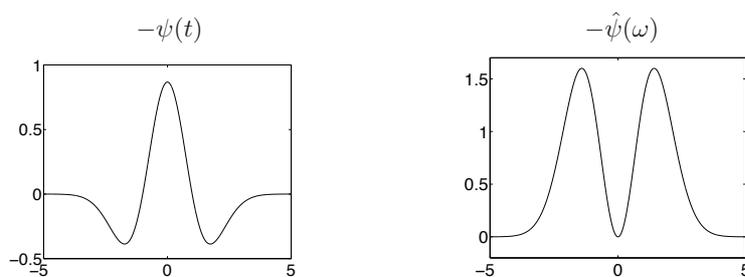
Figure 4.6: Mexican hat wavelet (4.34) for $\sigma = 1$ and its Fourier transform.
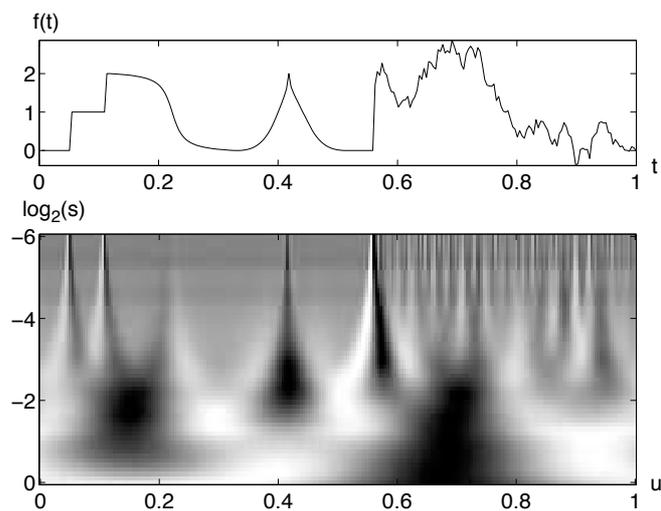

Figure 4.7: Real wavelet transform $Wf(u, s)$ computed with a Mexican hat wavelet (4.34). The vertical axis represents $\log_2 s$. Black, grey and white points correspond respectively to positive, zero and negative wavelet coefficients.

Theorem 4.4 (Calderón, Grossmann, Morlet). *Let $\psi \in \mathbf{L}^2(\mathbb{R})$ be a real function such that*

$$C_\psi = \int_0^{+\infty} \frac{|\hat{\psi}(\omega)|^2}{\omega} \, d\omega < +\infty. \tag{4.36}$$

Any $f \in \mathbf{L}^2(\mathbb{R})$ satisfies

$$f(t) = \frac{1}{C_\psi} \int_0^{+\infty} \int_{-\infty}^{+\infty} Wf(u,s) \frac{1}{\sqrt{s}} \, \psi\left(\frac{t-u}{s}\right) du \, \frac{ds}{s^2}, \tag{4.37}$$

and

$$\int_{-\infty}^{+\infty} |f(t)|^2 dt = \frac{1}{C_\psi} \int_0^{+\infty} \int_{-\infty}^{+\infty} |Wf(u,s)|^2 \, du \, \frac{ds}{s^2}. \tag{4.38}$$

Proof. The proof of (4.38) is almost identical to the proof of (4.18). Let us concentrate on the proof of (4.37). The right integral $b(t)$ of (4.37) can be rewritten as a sum of convolutions. Inserting $Wf(u,s) = f \star \bar{\psi}_s(u)$ with $\psi_s(t) = s^{-1/2} \, \psi(t/s)$ yields

$$
\begin{aligned}
b(t) &= \frac{1}{C_\psi} \int_0^{+\infty} Wf(.,s) \star \psi_s(t) \, \frac{ds}{s^2} \\
&= \frac{1}{C_\psi} \int_0^{+\infty} f \star \bar{\psi}_s \star \psi_s(t) \, \frac{ds}{s^2}.
\end{aligned} \tag{4.39}
$$

The "." indicates the variable over which the convolution is calculated. We prove that $b = f$ by showing that their Fourier transforms are equal. The Fourier transform of b is

$$\hat{b}(\omega) = \frac{1}{C_\psi} \int_0^{+\infty} \hat{f}(\omega) \sqrt{s} \, \hat{\psi}^*(s\omega) \sqrt{s} \, \hat{\psi}(s\omega) \, \frac{ds}{s^2} = \frac{\hat{f}(\omega)}{C_\psi} \int_0^{+\infty} |\hat{\psi}(s\,\omega)|^2 \, \frac{ds}{s}.$$

Since ψ is real we know that $|\hat{\psi}(-\omega)|^2 = |\hat{\psi}(\omega)|^2$. The change of variable $\xi = s\omega$ thus proves that

$$\hat{b}(\omega) = \frac{1}{C_\psi} \hat{f}(\omega) \int_0^{+\infty} \frac{|\hat{\psi}(\xi)|^2}{\xi} \, d\xi = \hat{f}(\omega).$$

■   ■

The theorem hypothesis

$$C_\psi = \int_0^{+\infty} \frac{|\hat{\psi}(\omega)|^2}{\omega} \, d\omega < +\infty$$

is called the wavelet *admissibility condition*. To guarantee that this integral is finite we must ensure that $\hat{\psi}(0) = 0$, which explains why we imposed that wavelets must have a zero average. This condition is nearly sufficient. If $\hat{\psi}(0) = 0$ and $\hat{\psi}(\omega)$ is continuously differentiable then the admissibility condition is satisfied. One can verify that $\hat{\psi}(\omega)$ is continuously differentiable if ψ has a sufficient time decay

$$\int_{-\infty}^{+\infty} (1 + |t|) \, |\psi(t)| \, dt < +\infty.$$

Reproducing Kernel   Like a windowed Fourier transform, a wavelet transform is a redundant representation, whose redundancy is characterized by a reproducing kernel equation. Inserting the reconstruction formula (4.37) into the definition of the wavelet transform yields

$$Wf(u_0, s_0) = \int_{-\infty}^{+\infty} \left(\frac{1}{C_\psi} \int_0^{+\infty} \int_{-\infty}^{+\infty} Wf(u,s) \, \psi_{u,s}(t) \, du \frac{ds}{s^2} \right) \psi_{u_0,s_0}^*(t) \, dt.$$

Interchanging these integrals gives

$$Wf(u_0, s_0) = \frac{1}{C_\psi} \int_{-\infty}^{+\infty} K(u, u_0, s, s_0) \, Wf(u,s) \, du \, \frac{ds}{s^2}, \tag{4.40}$$

with

$$K(u_0, u, s_0, s) = \langle \psi_{u,s}, \psi_{u_0,s_0} \rangle \, . \tag{4.41}$$

The reproducing kernel $K(u_0, u, s_0, s)$ measures the correlation of two wavelets $\psi_{u,s}$ and ψ_{u_0,s_0}. The reader can verify that any function $\Phi(u,s)$ is the wavelet transform of some $f \in \mathbf{L}^2(\mathbb{R})$ if and only if it satisfies the reproducing kernel equation (4.40).

Scaling Function   When $Wf(u,s)$ is known only for $s < s_0$, to recover f we need a complement of information corresponding to $Wf(u,s)$ for $s > s_0$. This is obtained by introducing a *scaling function* ϕ that is an aggregation of wavelets at scales larger than 1. The modulus of its Fourier transform is defined by

$$|\hat{\phi}(\omega)|^2 = \int_1^{+\infty} |\hat{\psi}(s\omega)|^2 \, \frac{ds}{s} = \int_\omega^{+\infty} \frac{|\hat{\psi}(\xi)|^2}{\xi} \, d\xi, \qquad (4.42)$$

and the complex phase of $\hat{\phi}(\omega)$ can be arbitrarily chosen. One can verify that $\|\phi\| = 1$ and we derive from the admissibility condition (4.36) that

$$\lim_{\omega \to 0} |\hat{\phi}(\omega)|^2 = C_\psi. \qquad (4.43)$$

The scaling function can thus be interpreted as the impulse response of a low-pass filter. Let us denote

$$\phi_s(t) = \frac{1}{\sqrt{s}} \, \phi\left(\frac{t}{s}\right) \quad \text{and} \quad \bar{\phi}_s(t) = \phi_s^*(-t).$$

The low-frequency approximation of f at the scale s is

$$Lf(u,s) = \left\langle f(t), \frac{1}{\sqrt{s}} \phi\left(\frac{t-u}{s}\right) \right\rangle = f \star \bar{\phi}_s(u). \qquad (4.44)$$

With a minor modification of the proof of Theorem 4.4, it can be shown that (Exercise 4.6)

$$f(t) = \frac{1}{C_\psi} \int_0^{s_0} Wf(.,s) \star \psi_s(t) \, \frac{ds}{s^2} + \frac{1}{C_\psi s_0} \, Lf(.,s_0) \star \phi_{s_0}(t). \qquad (4.45)$$

Example 4.7. *If ψ is the second order derivative of a Gaussian whose Fourier transform is given by (4.35), then the integration (4.42) yields*

$$\hat{\phi}(\omega) = \frac{2\sigma^{3/2}\pi^{1/4}}{\sqrt{3}} \sqrt{\omega^2 + \frac{1}{\sigma^2}} \, \exp\left(-\frac{\sigma^2\omega^2}{2}\right). \qquad (4.46)$$

Figure 4.8 displays ϕ and $\hat{\phi}$ for $\sigma = 1$.
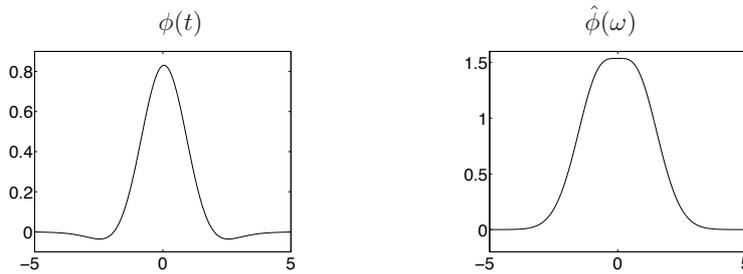

Figure 4.8: Scaling function associated to a Mexican hat wavelet and its Fourier transform calculated with (4.46).

4.3.2   Analytic Wavelets

To analyze the time evolution of frequency tones, it is necessary to use an analytic wavelet to separate the phase and amplitude information of signals. The properties of the resulting analytic wavelet transform are studied.

Analytic Signal A function $f_a \in \mathbf{L}^2(\mathbb{R})$ is said to be *analytic* if its Fourier transform is zero for negative frequencies:

$$\hat{f}_a(\omega) = 0 \;\; \text{if } \omega < 0.$$

An analytic function is necessarily complex but is entirely characterized by its real part. Indeed, the Fourier transform of its real part $f = \mathrm{Re}[f_a]$ is

$$\hat{f}(\omega) = \frac{\hat{f}_a(\omega) + \hat{f}_a^*(-\omega)}{2},$$

and this relation can be inverted:

$$\hat{f}_a(\omega) = \begin{cases} 2\,\hat{f}(\omega) & \text{if } \omega \geqslant 0 \\ 0 & \text{if } \omega < 0 \end{cases}. \tag{4.47}$$

The analytic part $f_a(t)$ of a signal $f(t)$ is the inverse Fourier transform of $\hat{f}_a(\omega)$ defined by (4.47).

Discrete Analytic Part The analytic part $f_a[n]$ of a discrete signal $f[n]$ of size N is also computed by setting to zero the negative frequency components of its discrete Fourier transform. The Fourier transform values at $k = 0$ and $k = N/2$ must be carefully adjusted so that $\mathrm{Re}[f_a] = f$ (Exercise 3.22):

$$\hat{f}_a[k] = \begin{cases} \hat{f}[k] & \text{if } k = 0,\, N/2 \\ 2\,\hat{f}[k] & \text{if } 0 < k < N/2 \\ 0 & \text{if } N/2 < k < N \end{cases}. \tag{4.48}$$

We obtain $f_a[n]$ by computing the inverse discrete Fourier transform.

Example 4.8. *The Fourier transform of*

$$f(t) = a\cos(\omega_0 t + \phi) = \frac{a}{2}\Big(\exp[i(\omega_0 t + \phi)] + \exp[-i(\omega_0 t + \phi)]\Big)$$

is

$$\hat{f}(\omega) = \pi a\Big(\exp(i\phi)\,\delta(\omega - \omega_0) + \exp(-i\phi)\,\delta(\omega + \omega_0)\Big).$$

The Fourier transform of the analytic part computed with (4.47) is $\hat{f}_a(\omega) = 2\pi a\,\exp(i\phi)\,\delta(\omega - \omega_0)$ and hence

$$f_a(t) = a\,\exp[i(\omega_0 t + \phi)]. \tag{4.49}$$

Time-Frequency Resolution An analytic wavelet transform is calculated with an analytic wavelet ψ:

$$Wf(u, s) = \langle f, \psi_{u,s} \rangle = \int_{-\infty}^{+\infty} f(t)\,\frac{1}{\sqrt{s}}\,\psi^*\left(\frac{t - u}{s}\right)\,dt. \tag{4.50}$$

Its time-frequency resolution depends on the time-frequency spread of the wavelet atoms $\psi_{u,s}$. We suppose that ψ is centered at 0, which implies that $\psi_{u,s}$ is centered at $t = u$. With the change of variable $v = \frac{t-u}{s}$, we verify that

$$\int_{-\infty}^{+\infty} (t - u)^2\,|\psi_{u,s}(t)|^2\,dt = s^2\,\sigma_t^2, \tag{4.51}$$

with $\sigma_t^2 = \int_{-\infty}^{+\infty} t^2\,|\psi(t)|^2\,dt$. Since $\hat{\psi}(\omega)$ is zero at negative frequencies, the center frequency η of $\hat{\psi}$ is

$$\eta = \frac{1}{2\pi}\int_0^{+\infty} \omega\,|\hat{\psi}(\omega)|^2\,d\omega. \tag{4.52}$$

The Fourier transform of $\psi_{u,s}$ is a dilation of $\hat{\psi}$ by $1/s$:

$$\hat{\psi}_{u,s}(\omega) = \sqrt{s}\,\hat{\psi}(s\omega)\,\exp(-i\omega u). \tag{4.53}$$

Its center frequency is therefore η/s. The energy spread of $\hat{\psi}_{u,s}$ around η/s is

$$\frac{1}{2\pi} \int_0^{+\infty} \left(\omega - \frac{\eta}{s}\right)^2 \left|\hat{\psi}_{u,s}(\omega)\right|^2 d\omega = \frac{\sigma_\omega^2}{s^2}, \tag{4.54}$$

with

$$\sigma_\omega^2 = \frac{1}{2\pi} \int_0^{+\infty} (\omega - \eta)^2 \, |\hat{\psi}(\omega)|^2 \, d\omega.$$

The energy spread of a wavelet time-frequency atom $\psi_{u,s}$ thus corresponds to a Heisenberg box centered at $(u, \eta/s)$, of size $s\sigma_t$ along time and σ_ω/s along frequency. The area of the rectangle remains equal to $\sigma_t\,\sigma_\omega$ at all scales but the resolution in time and frequency depends on s, as illustrated in Figure 4.9.

An analytic wavelet transform defines a local time-frequency energy density $P_W f$, which measures the energy of f in the Heisenberg box of each wavelet $\psi_{u,s}$ centered at $(u, \xi = \eta/s)$:

$$P_W f(u, \xi) = |Wf(u, s)|^2 = \left|Wf\left(u, \frac{\eta}{\xi}\right)\right|^2. \tag{4.55}$$

This energy density is called a *scalogram*.
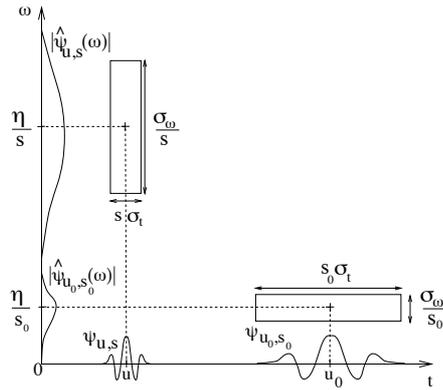
Figure 4.9: Heisenberg boxes of two wavelets. Smaller scales decrease the time spread but increase the frequency support, which is shifted towards higher frequencies.

Completeness   An analytic wavelet transform of f depends only on its analytic part f_a. The following theorem derives a reconstruction formula and proves that energy is conserved for real signals.

Theorem 4.5. *For any* $f \in \mathbf{L^2}(\mathbb{R})$

$$Wf(u, s) = \frac{1}{2}\,Wf_a(u, s). \tag{4.56}$$

If $C_\psi = \int_0^{+\infty} \omega^{-1}\,|\hat{\psi}(\omega)|^2\,d\omega < +\infty$ *and* f *is real then*

$$f(t) = \frac{2}{C_\psi}\mathrm{Re}\left[\int_0^{+\infty}\int_{-\infty}^{+\infty} Wf(u, s)\,\psi_s(t - u)\,du\,\frac{ds}{s^2}\right], \tag{4.57}$$

and

$$\|f\|^2 = \frac{2}{C_\psi}\int_0^{+\infty}\int_{-\infty}^{+\infty} |Wf(u, s)|^2\,du\,\frac{ds}{s^2}. \tag{4.58}$$

Proof. Let us first prove (4.56). The Fourier transform with respect to u of

$$f_s(u) = Wf(u, s) = f \star \bar{\psi}_s(u)$$

is

$$\hat{f}_s(\omega) = \hat{f}(\omega)\sqrt{s}\,\hat{\psi}^*(s\omega).$$

Since $\hat{\psi}(\omega) = 0$ at negative frequencies, and $\hat{f}_a(\omega) = 2\hat{f}(\omega)$ for $\omega \geqslant 0$, we derive that

$$\hat{f}_s(\omega) = \frac{1}{2}\,\hat{f}_a(\omega)\,\sqrt{s}\,\hat{\psi}^*(s\omega),$$

which is the Fourier transform of (4.56).

With the same derivations as in the proof of (4.37) one can verify that the inverse wavelet formula reconstructs the analytic part of f:

$$f_a(t) = \frac{1}{C_\psi} \int_0^{+\infty} \int_{-\infty}^{+\infty} Wf_a(u,s)\,\psi_s(t-u)\,\frac{ds}{s^2}\,du. \tag{4.59}$$

Since $f = \mathrm{Re}[f_a]$, inserting (4.56) proves (4.57).

An energy conservation for the analytic part f_a is proved as in (4.38) by applying the Plancherel formula:

$$\int_{-\infty}^{+\infty} |f_a(t)|^2\,dt = \frac{1}{C_\psi} \int_0^{+\infty} \int_{-\infty}^{+\infty} |W_a f(u,s)|^2\,du\,\frac{ds}{s^2}.$$

Since $Wf_a(u,s) = 2Wf(u,s)$ and $\|f_a\|^2 = 2\|f\|^2$, equation (4.58) follows. ■ ■

If f is real the change of variable $\xi = 1/s$ in the energy conservation (4.58) proves that

$$\|f\|^2 = \frac{2}{C_\psi} \int_0^{+\infty} \int_{-\infty}^{+\infty} P_W f(u,\xi)\,du\,d\xi.$$

It justifies the interpretation of a scalogram as a time-frequency energy density.

Wavelet Modulated Windows An analytic wavelet can be constructed with a frequency modulation of a real and symmetric window g. The Fourier transform of

$$\psi(t) = g(t)\,\exp(i\eta t) \tag{4.60}$$

is $\hat{\psi}(\omega) = \hat{g}(\omega - \eta)$. If $\hat{g}(\omega) = 0$ for $|\omega| > \eta$ then $\hat{\psi}(\omega) = 0$ for $\omega < 0$. Hence ψ is analytic, as shown in Figure 4.10. Since g is real and even, $\hat{g}$ is also real and symmetric. The center frequency of $\hat{\psi}$ is therefore η and

$$|\hat{\psi}(\eta)| = \sup_{\omega \in \mathbb{R}} |\hat{\psi}(\omega)| = \hat{g}(0). \tag{4.61}$$

A *Gabor wavelet* $\psi(t) = g(t)\,e^{i\eta t}$ is obtained with a Gaussian window

$$g(t) = \frac{1}{(\sigma^2 \pi)^{1/4}}\,\exp\left(\frac{-t^2}{2\sigma^2}\right). \tag{4.62}$$

The Fourier transform of this window is $\hat{g}(\omega) = (4\pi\sigma^2)^{1/4}\exp(-\sigma^2\omega^2/2)$. If $\sigma^2\eta^2 \gg 1$ then $\hat{g}(\omega) \approx 0$ for $|\omega| > \eta$. Such Gabor wavelets are thus considered to be approximately analytic.
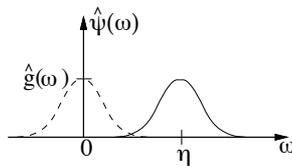
Figure 4.10: Fourier transform $\hat{\psi}(\omega)$ of a wavelet $\psi(t) = g(t)\,\exp(i\eta t)$.

Example 4.9. *The wavelet transform of $f(t) = a \exp(i\omega_0 t)$ is*

$$Wf(u, s) = a\sqrt{s}\,\hat{\psi}^*(s\omega_0)\,\exp(i\omega_0 t) = a\,\sqrt{s}\,\hat{g}(s\omega_0 - \eta)\,\exp(i\omega_0 t).$$

Observe that the normalized scalogram is maximum at $\xi = \omega_0$:

$$\frac{\xi}{\eta} P_W f(u, \xi) = \frac{1}{s}|Wf(u,s)|^2 = a^2 \left| \hat{g}\left(\eta\left(\frac{\omega_0}{\xi} - 1\right)\right)\right|^2.$$

Example 4.10. *The wavelet transform of a linear chirp $f(t) = \exp(iat^2) = \exp[i\phi(t)]$ is computed for a Gabor wavelet whose Gaussian window is (4.62). By using the Fourier transform of Gaussian chirps (2.34) one can verify that*

$$\frac{|Wf(u,s)|^2}{s} = \left(\frac{4\pi\sigma^2}{1 + 4s^2 a^2 \sigma^4} \right)^{1/2} \exp\left(\frac{-\sigma^2}{1 + 4a^2 s^4 \sigma^4}(\eta - 2asu)^2 \right).$$

As long as $4a^2 s^4 \sigma^4 \ll 1$, at a fixed time u the renormalized scalogram $\eta^{-1}\xi P_W f(u, \xi)$ is a Gaussian function of s that reaches its maximum at

$$\xi(u) = \frac{\eta}{s(u)} = \phi'(u) = 2\,a\,u. \tag{4.63}$$

Section 4.4.2 explains why the amplitude is maximum at the instantaneous frequency $\phi'(u)$.

(a)

(b)
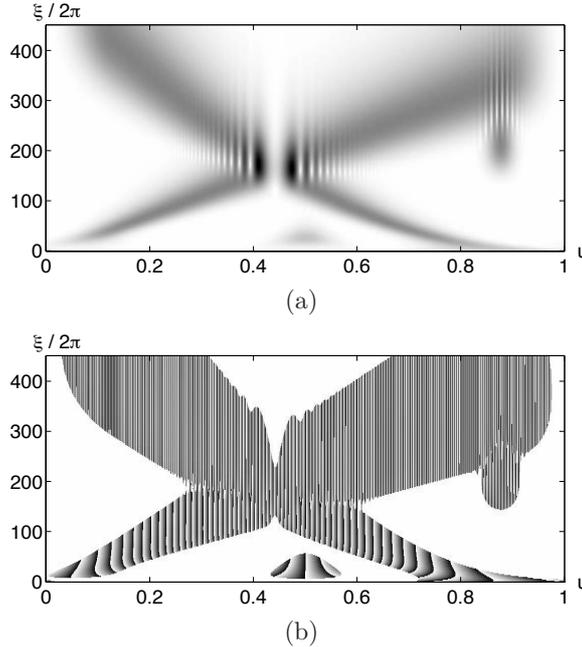
Figure 4.11: (a) Normalized scalogram $\eta^{-1}\xi P_W f(u, \xi)$ computed from the signal in Figure 4.3. Dark points indicate large amplitude coefficients. (b) Complex phase $\Theta_W(u, \xi)$ of $Wf(u, \eta/\xi)$, where the modulus is non-zero.

Example 4.11. *Figure 4.11 displays the normalized scalogram $\eta^{-1}\xi P_W f(u, \xi)$, and the complex phase $\Theta_W(u, \xi)$ of $Wf(u, s)$, for the signal f of Figure 4.3. The frequency bandwidth of wavelet atoms is proportional to $1/s = \xi/\eta$. The frequency resolution of the scalogram is therefore finer than the spectrogram at low frequencies but coarser than the spectrogram at higher frequencies. This explains why the wavelet transform produces interference patterns between the high frequency Gabor function at the abscissa $t = 0.87$ and the quadratic chirp at the same location, whereas the spectrogram in Figure 4.3 separates them well.*

4.3.3 Discrete Wavelets

Let $\bar{f}(t)$ be a continuous time signal defined over $[0, 1]$. Let $f[n]$ be a the discrete signal obtained by a low-pass filtering of $\bar{f}$ and a uniform sampling at intervals N^{-1}. Its discrete wavelet transform can only be calculated at scales $N^{-1} < s < 1$, as shown in Figure 4.7. It is calculated for $s = a^j$, with $a = 2^{1/v}$, which provides v intermediate scales in each octave $[2^j, 2^{j+1})$.

Let $\psi(t)$ be a wavelet whose support is included in $[-K/2, K/2]$. For $1 \leqslant a^j \leqslant N K^{-1}$, a discrete wavelet scaled by a^j is defined by

$$\psi_j[n] = \frac{1}{\sqrt{a^j}} \, \psi\left(\frac{n}{a^j}\right) \, .$$

This discrete wavelet has $K a^j$ non-zero values on $[-N/2, N/2]$. The scale a^j is larger than 1 otherwise the sampling interval may be larger than the wavelet support.

Fast Transform To avoid border problems, we treat $f[n]$ and the wavelets $\psi_j[n]$ as periodic signals of period N. The discrete wavelet transform can then be written as a circular convolution with $\bar{\psi}_j[n] = \psi_j^*[-n]$:

$$Wf[n, a^j] = \sum_{m=0}^{N-1} f[m] \, \psi_j^*[m - n] = f \circledast \bar{\psi}_j[n]. \tag{4.64}$$

This circular convolution is calculated with the fast Fourier transform algorithm, which requires $O(N \log_2 N)$ operations. If $a = 2^{1/v}$, there are $v \log_2(N/(2K))$ scales $a^j \in [2N^{-1}, K^{-1}]$. The total number of operations to compute the wavelet transform over all scales is therefore $O(v N (\log_2 N)^2)$ [407].

To compute the scalogram $P_W[n, \xi] = |Wf[n, \frac{\eta}{\xi}]|^2$ we calculate $Wf[n, s]$ at any scale s with a parabola interpolation. Let j be the closest integer to $\log_2 s / \log_2 a$, and $p(x)$ be the parabola such that

$$p(j - 1) = Wf[n, a^{j-1}] \quad , \quad p(j) = Wf[n, a^j] \quad , \quad p(j + 1) = Wf[n, a^{j+1}].$$

A second order interpolation computes

$$Wf[n, s] = p\left(\frac{\log_2 s}{\log_2 a}\right) \, .$$

Parabolic interpolations are used instead of linear interpolations in order to locate more precisely the ridges defined in Section 4.4.2.

Discrete Scaling Filter A wavelet transform computed up to a scale a^J is not a complete signal representation. It is necessary to add the low frequencies $Lf[n, a^J]$ corresponding to scales larger than a^J. A discrete and periodic scaling filter is computed by sampling the scaling function $\phi(t)$ defined in (4.42):

$$\phi_J[n] = \frac{1}{\sqrt{a^J}} \, \phi\left(\frac{n}{a^J}\right) \quad \text{for } n \in [-N/2, N/2].$$

Let $\bar{\phi}_J[n] = \phi_J^*[-n]$. The low frequencies are carried by

$$Lf[n, a^J] = \sum_{m=0}^{N-1} f[m] \, \phi_J^*[m - n] = f \circledast \bar{\phi}_J[n]. \tag{4.65}$$

Reconstruction An approximate inverse wavelet transform is implemented by discretizing the integral (4.45). Suppose that $a^I = 1$ is the finest scale. Since $ds/s^2 = d\log_e s / s$ and the discrete wavelet transform is computed along an exponential scale sequence $\{a^j\}_j$ with a logarithmic increment $d \log_e s = \log_e a$, we obtain

$$f[n] \approx \frac{\log_e a}{C_\psi} \sum_{j=I}^{J} \frac{1}{a^j} \, Wf[., a^j] \circledast \psi_j[n] + \frac{1}{C_\psi \, a^J} \, Lf[., a^J] \circledast \phi_J[n]. \tag{4.66}$$

The "." indicates the variable over which the convolution is calculated. These circular convolutions are calculated using the FFT, with $O(vN(\log_2 N)^2)$ operations.

Analytic wavelet transforms are often computed over real signals $f[n]$ that have no energy at low frequencies. The scaling filter component is then negligible. Theorem 4.5 shows that

$$f[n] \approx \frac{2\log_e a}{C_\psi} \operatorname{Re}\left(\sum_{j=I}^{J} \frac{1}{a^j} \, Wf[.,a^j] \circledast \psi_j[n]\right). \tag{4.67}$$

The error introduced by the discretization of scales decreases when the number v of voices per octave increases. However, the approximation of continuous time convolutions with discrete convolutions also creates high frequency errors. Perfect reconstructions are obtained with a more careful design of the reconstruction filters (Exercise 4.3). Section 5.2.2 describes an exact inverse wavelet transform computed at dyadic scales $a^j = 2^j$.

4.4   Time-Frequency Geometry of Instantaneous Frequencies

When listening to music, we perceive several frequencies that change with time. In music, it is associated to the geometric perception of "movements". This notion of instantaneous frequency remains to be defined. The time variation of several instantaneous frequencies is measured with local maxima of windowed Fourier transforms and wavelet transforms. They define a geometric time-frequency support from which signal approximations are recovered. Audio processing is implemented by modifying this time-frequency support.

Analytic Instantaneous Frequency   A cosine modulation

$$f(t) = a\cos(\omega_0 t + \theta_0) = a\cos\theta(t)$$

has a frequency ω_0 that is the derivative of the phase $\theta(t) = \omega_0 t + \theta_0$. To generalize this notion, real signals f are written as an amplitude $a(t)$ modulated with a time varying phase $\theta(t)$:

$$f(t) = a(t)\,\cos\theta(t) \quad\text{with}\quad a(t) \geqslant 0 \;. \tag{4.68}$$

The *instantaneous frequency* is defined as a positive derivative of the phase:

$$\omega(t) = \theta'(t) \geqslant 0 \;.$$

The derivative can be chosen to be positive by adapting the sign of $\theta(t)$. For a given $f(t)$ there are many possible choices of $a(t)$ and $\theta(t)$ satisfying (4.68), so $\omega(t)$ is not uniquely defined relative to f.

A particular decomposition (4.68) is obtained from the analytic part f_a of f, whose Fourier transform is defined in (4.47) by

$$\hat{f}_a(\omega) = \begin{cases} 2\,\hat{f}(\omega) & \text{if } \omega \geqslant 0 \\ 0 & \text{if } \omega < 0 \end{cases} \;. \tag{4.69}$$

This complex signal is represented by separating the modulus and the complex phase:

$$f_a(t) = a(t)\,\exp[i\theta(t)] \;. \tag{4.70}$$

Since $f = \operatorname{Re}[f_a]$, it follows that

$$f(t) = a(t)\,\cos\theta(t).$$

We call $a(t)$ the *analytic* amplitude of $f(t)$ and $\theta'(t)$ its *instantaneous frequency*; they are uniquely defined.

Example 4.12. *If $f(t) = a(t) \cos(\omega_0 t + \theta_0)$, then*

$$\hat{f}(\omega) = \frac{1}{2} \Big(\exp(i\theta_0) \, \hat{a}(\omega - \omega_0) + \exp(-i\theta_0) \, \hat{a}(\omega + \omega_0) \Big) \ .$$

If the variations of $a(t)$ are slow compared to the period $2\pi/\omega_0$, which is achieved by requiring that the support of $\hat{a}$ be included in $[-\omega_0, \omega_0]$, then

$$\hat{f}_a(\omega) = \hat{a}(\omega - \omega_0) \, \exp(i\theta_0)$$

so $f_a(t) = a(t) \, \exp[i(\omega_0 t + \theta_0)]$.

If a signal f is the sum of two sinusoidal waves:

$$f(t) = a \cos(\omega_1 t) + a \cos(\omega_2 t),$$

then

$$f_a(t) = a \, \exp(i\omega_1 t) + a \, \exp(i\omega_2 t) = 2a \, \cos\left(\frac{1}{2} \, (\omega_1 - \omega_2) \, t \right) \exp\left(\frac{i}{2} \, (\omega_1 + \omega_2) \, t \right).$$

The instantaneous frequency is $\theta'(t) = (\omega_1 + \omega_2)/2$ and the amplitude is

$$a(t) = 2a \, \left| \cos\left(\frac{1}{2} \, (\omega_1 - \omega_2) \, t \right) \right|.$$

This result is not satisfying because it does not reveal that the signal includes two sinusoidal waves of the same amplitude. It measures an average frequency value. Next sections explain how to measure the instantaneous frequencies of several spectral components by separating them with a windowed Fourier transform or a wavelet transform. We first describe two important applications of instantaneous frequencies.

Frequency Modulation   In signal communications, information can be transmitted through the amplitude $a(t)$ (amplitude modulation) or the instantaneous frequency $\theta'(t)$ (frequency modulation) [59]. Frequency modulation is more robust in the presence of additive Gaussian white noise. In addition, it better resists multi-path interferences, which destroy the amplitude information. A frequency modulation sends a message $m(t)$ through a signal

$$f(t) = a \cos \theta(t) \quad \text{with} \quad \theta'(t) = \omega_0 + k \, m(t).$$

The frequency bandwidth of f is proportional to k. This constant is adjusted depending on the transmission noise and the available bandwidth. At the reception, the message $m(t)$ is restored with a frequency demodulation that computes the instantaneous frequency $\theta'(t)$ [118].

Additive Sound Models   Musical sounds and voiced speech segments can be modeled with sums of sinusoidal *partials*:

$$f(t) = \sum_{k=1}^{K} f_k(t) = \sum_{k=1}^{K} a_k(t) \cos \theta_k(t) \ , \tag{4.71}$$

where a_k and θ'_k are vary slowly [412, 413]. Such decompositions are useful for pattern recognition and for modifying sound properties [338]. Sections 4.4.1 and 4.4.2 explain how to compute a_k and the instantaneous frequency θ'_k and reconstruct signals from these information.

Reducing or increasing the duration of a sound f by a factor α in time is used by radio broadcast to adjust recorded sequences to a precise time schedule. A scaling $f(\alpha t)$ transforms each $\theta_k(t)$ in $\theta_k(\alpha t)$ and hence $\theta'_k(t)$ in $\alpha \theta'_k(t)$. For sound reduction, with $\alpha > 1$ all frequencies are thus increased. To avoid modifying the values of θ'_k and a_k, a new sound is synthesized

$$f_\alpha(t) = \sum_{k=1}^{K} a_k(\alpha \, t) \, \cos\Big(\frac{1}{\alpha} \, \theta_k(\alpha \, t) \Big). \tag{4.72}$$

The partials of f_α at $t = t_0/\alpha$ and the partials of f at $t = t_0$ have the same amplitudes and the same instantaneous frequencies and the properties of these sounds are thus perceived as identical.

A frequency transposition with same duration is calculated by dividing each phase by a constant α in order to shift the sound harmonics:

$$f_\alpha(t) = \sum_{k=1}^{K} b_k(t) \, \cos\Big(\theta_k(t)/\alpha\Big). \tag{4.73}$$

The instantaneous frequency of each partial is now $\theta_k'(t)/\alpha$. To maintain the sound properties, the amplitudes $b_k(t)$ must be adjusted in order not to modify the global frequency envelop $F(t,\omega)$ of the harmonics:

$$a_k(t) = F\Big(t, \theta_k'(t)\Big) \;\; \text{and} \;\; b_k(t) = F\Big(t, \theta_k'(t)/\alpha\Big) \,. \tag{4.74}$$

Many types of sounds, musical instruments or speech, are produced by an excitation that propagates across a wave guide. Locally, $F(t,\omega)$ is the transfer function of the wave guide. In speech processing, it is called a *formant*. This transfer function is often approximated with an autoregressive filter of order M, in which case:

$$F(t,\omega) = \frac{C}{\sum_{m=0}^{M-1} c_m \, \mathrm{e}^{-im\omega}} \,. \tag{4.75}$$

The parameters c_m are identified with (4.74) from the a_k and the b_k are then derived with (4.74) and (4.75).

4.4.1   Windowed Fourier Ridges

The spectrogram $P_S f(u,\xi) = |Sf(u,\xi)|^2$ measures the energy of f in a time-frequency neighborhood of (u,ξ). The ridge algorithm computes the signal instantaneous frequencies and amplitudes from the local maxima of $P_S f(u,\xi)$. These local maxima define a geometric support in the time-frequency plane. Modifications of sound durations or frequency transpositions are computed with time or frequency dilations of the ridge support.

Time-frequency ridges were introduced by Delprat, Escudié, Guillemain, Kronland-Martinet, Tchamitchian and Torrésani [203, 64] to analyze musical sounds. Since then it has found applications for a wide range of signals [288, 64] that have time varying frequency tones.

The windowed Fourier transform is computed with a symmetric window $g(t) = g(-t)$ whose support is equal to $[-1/2, 1/2]$. The Fourier transform $\hat{g}$ is a real symmetric function. We suppose that $|\hat{g}(\omega)| \leqslant \hat{g}(0)$ for all $\omega \in \mathbb{R}$, and that $\hat{g}(0) = \int_{-1/2}^{1/2} g(t)\, dt$ is on the order of 1. Table 4.1 gives several examples of such windows. The window g is normalized so that $\|g\| = 1$. For a fixed scale s, $g_s(t) = s^{-1/2} g(t/s)$ has a support of size s and a unit norm. The corresponding windowed Fourier atoms are

$$g_{s,u,\xi}(t) = g_s(t - u)\, \mathrm{e}^{i\xi t},$$

and the resulting windowed Fourier transform is

$$Sf(u,\xi) \;=\; \langle f, g_{s,u,\xi} \rangle \;=\; \int_{-\infty}^{+\infty} f(t)\, g_s(t-u)\, \mathrm{e}^{-i\xi t}\, dt. \tag{4.76}$$

The following theorem relates $Sf(u,\xi)$ to the instantaneous frequency of f.

Theorem 4.6. *Let* $f(t) = a(t) \cos\theta(t)$. *If* $\xi \geqslant 0$ *then*

$$\langle f, g_{s,u,\xi} \rangle = \frac{\sqrt{s}}{2}\, a(u) \, \exp(i[\theta(u) - \xi u]) \left(\hat{g}(s[\xi - \theta'(u)]) + \varepsilon(u,\xi) \right) . \tag{4.77}$$

The corrective term satisfies

$$|\varepsilon(u,\xi)| \leqslant \varepsilon_{a,1} + \varepsilon_{a,2} + \varepsilon_{\theta,2} + \sup_{|\omega| \geqslant s\theta'(u)} |\hat{g}(\omega)| \tag{4.78}$$

with

$$\varepsilon_{a,1} \leqslant \frac{s\,|a'(u)|}{|a(u)|} \quad , \quad \varepsilon_{a,2} \leqslant \sup_{|t-u|\leqslant s/2} \frac{s^2\,|a''(t)|}{|a(u)|} \, , \tag{4.79}$$

and if $s\,|a'(u)|\,|a(u)|^{-1} \leqslant 1$*, then*

$$\varepsilon_{\theta,2} \leqslant \sup_{|t-u|\leqslant s/2} s^2|\theta''(t)| \, . \tag{4.80}$$

If $\xi = \theta'(u)$ *then*

$$\varepsilon_{a,1} = \frac{s\,|a'(u)|}{|a(u)|} \left| \hat{g}'\Big(2\,s\,\theta'(u)\Big) \right| \, . \tag{4.81}$$

Proof. Observe that

$$\begin{aligned}
\langle f, g_{s,u,\xi} \rangle &= \int_{-\infty}^{+\infty} a(t)\,\cos\theta(t)\,g_s(t-u)\,\exp(-i\xi t)\,dt \\
&= \frac{1}{2} \int_{-\infty}^{+\infty} a(t)\,(\exp[i\theta(t)] + \exp[-i\theta(t)])\,g_s(t-u)\,\exp[-i\xi t]\,dt \\
&= I(\theta) + I(-\theta).
\end{aligned}$$

We first concentrate on

$$\begin{aligned}
I(\theta) &= \frac{1}{2} \int_{-\infty}^{+\infty} a(t)\,\exp[i\theta(t)]\,g_s(t-u)\,\exp(-i\xi t)\,dt \\
&= \frac{1}{2} \int_{-\infty}^{+\infty} a(t+u)\,e^{i\theta(t+u)}\,g_s(t)\,\exp[-i\xi(t+u)]\,dt.
\end{aligned}$$

This integral is computed by using second order Taylor expansions:

$$\begin{aligned}
a(t+u) &= a(u) + t\,a'(u) + \frac{t^2}{2}\,\alpha(t) \quad \text{with} \quad |\alpha(t)| \leqslant \sup_{h\in[u,t+u]} |a''(h)| \\
\theta(t+u) &= \theta(u) + t\,\theta'(u) + \frac{t^2}{2}\,\beta(t) \quad \text{with} \quad |\beta(t)| \leqslant \sup_{h\in[u,t+u]} |\theta''(h)| \, .
\end{aligned}$$

We get

$$\begin{aligned}
2\quad &\exp\Big(-i(\theta(u)-\xi u)\Big)\,I(\theta) = \\
&\int_{-\infty}^{+\infty} a(u)\,g_s(t)\,\exp\Big(-it(\xi-\theta'(u))\Big)\,\exp\Big(i\frac{t^2}{2}\beta(t)\Big)\,dt \\
&+ \int_{-\infty}^{+\infty} a'(u)\,t\,g_s(t)\,\exp\Big(-it(\xi-\theta'(u))\Big)\,\exp\Big(i\frac{t^2}{2}\beta(t)\Big)\,dt \\
&+ \frac{1}{2} \int_{-\infty}^{+\infty} \alpha(t)\,t^2\,g_s(t)\,\exp\Big(-i(t\xi+\theta(u)-\theta(t+u))\Big)\,dt \, .
\end{aligned}$$

A first order Taylor expansion of $\exp(ix)$ gives

$$\exp\Big(i\frac{t^2}{2}\beta(t)\Big) = 1 + \frac{t^2}{2}\beta(t)\,\gamma(t) \quad \text{with} \quad |\gamma(t)| \leqslant 1 \, . \tag{4.82}$$

Since

$$\int_{-\infty}^{+\infty} g_s(t)\,\exp[-it(\xi-\theta'(u))]\,dt = \sqrt{s}\,\hat{g}(s[\xi-\theta'(u)]) \, ,$$

inserting (4.82) in the expression of $I(\theta)$ yields

$$\left| I(\theta) - \frac{\sqrt{s}}{2} a(u)\,\exp[i(\theta(u)-\xi u)]\,\hat{g}(\xi-\theta'(u)) \right| \leqslant \frac{\sqrt{s}\,|a(u)|}{4}(\varepsilon_{a,1}^+ + \varepsilon_{a,2} + \varepsilon_{\theta,2}) \tag{4.83}$$

with

$$\varepsilon_{a,1}^+ = \frac{2|a'(u)|}{|a(u)|} \left| \int_{-\infty}^{+\infty} t \frac{1}{\sqrt{s}} g_s(t) \exp[-it(\xi - \theta'(u))] \, dt \right| , \tag{4.84}$$

$$\varepsilon_{a,2} = \int_{-\infty}^{+\infty} t^2 |\alpha(t)| \frac{1}{\sqrt{s}} |g_s(t)| \, dt , \tag{4.85}$$

$$\varepsilon_{\theta,2} = \int_{-\infty}^{+\infty} t^2 |\beta(t)| \frac{1}{\sqrt{s}} |g_s(t)| \, dt \tag{4.86}$$

$$+ \frac{|a'(u)|}{|a(u)|} \int_{-\infty}^{+\infty} |t^3| \, |\beta(t)| \frac{1}{\sqrt{s}} |g_s(t)| dt.$$

Applying (4.83) to $I(-\theta)$ gives

$$|I(-\theta)| \leqslant \frac{\sqrt{s} \, |a(u)|}{2} |\hat{g}(\xi + \theta'(u))| + \frac{\sqrt{s} \, |a(u)|}{4} (\varepsilon_{a,1}^- + \varepsilon_{a,2} + \varepsilon_{\theta,2}) ,$$

with

$$\varepsilon_{a,1}^- = \frac{2|a'(u)|}{|a(u)|} \left| \int_{-\infty}^{+\infty} t \frac{1}{\sqrt{s}} g_s(t) \exp[-it(\xi + \theta'(u))] \, dt \right| . \tag{4.87}$$

Since $\xi \geqslant 0$ and $\theta'(u) \geqslant 0$, we derive that

$$|\hat{g}(s[\xi + \theta'(u)])| \leqslant \sup_{|\omega| \geqslant s\theta'(u)} |\hat{g}(\omega)| ,$$

and hence

$$I(\theta) + I(-\theta) = \frac{\sqrt{s}}{2} a(u) \exp[i(\theta(u) - \xi u)] \Big(\hat{g}\big(s[\xi - \theta'(u)]\big) + \varepsilon(u, \xi) \Big)$$

with

$$\varepsilon(u, \xi) = \frac{\varepsilon_{a,1}^+ + \varepsilon_{a,1}^-}{2} + \varepsilon_{a,2} + \varepsilon_{\theta,2} + \sup_{|\omega| \geqslant s|\theta'(u)|} |\hat{g}(\omega)| .$$

Let us now verify the upper bound (4.79) for $\varepsilon_{a,1} = (\varepsilon_{a,1}^+ + \varepsilon_{a,1}^-)/2$. Since $g_s(t) = s^{-1/2} g(t/s)$, a simple calculation shows that for $n \geqslant 0$

$$\int_{-\infty}^{+\infty} |t|^n \frac{1}{\sqrt{s}} |g_s(t)| \, dt = s^n \int_{-1/2}^{1/2} |t|^n |g(t)| \, dt \leqslant \frac{s^n}{2^n} \|g\|^2 = \frac{s^n}{2^n}. \tag{4.88}$$

Inserting this for $n = 1$ in (4.84) and (4.87) gives

$$\varepsilon_{a,1} = \frac{\varepsilon_{a,1}^+ + \varepsilon_{a,1}^-}{2} \leqslant \frac{s \, |a'(u)|}{|a(u)|} .$$

The upper bounds (4.79) and (4.80) of the second order terms $\varepsilon_{a,2}$ and $\varepsilon_{\theta,2}$ are obtained by observing that the remainder $\alpha(t)$ and $\beta(t)$ of the Taylor expansion of $a(t+u)$ and $\theta(t+u)$ satisfy

$$\sup_{|t| \leqslant s/2} |\alpha(t)| \leqslant \sup_{|t-u| \leqslant s/2} |a''(t)| \quad , \quad \sup_{|t| \leqslant s/2} |\beta(t)| \leqslant \sup_{|t-u| \leqslant s/2} |\theta''(t)|. \tag{4.89}$$

Inserting this in (4.85) yields

$$\varepsilon_{a,2} \leqslant \sup_{|t-u| \leqslant s/2} \frac{s^2 \, |a''(t)|}{|a(u)|} .$$

When $s \, |a'(u)| |a(u)|^{-1} \leqslant 1$, replacing $|\beta(t)|$ by its upper bound in (4.86) gives

$$\varepsilon_{\theta,2} \leqslant \frac{1}{2} \Big(1 + \frac{s \, |a'(u)|}{|a(u)|}\Big) \sup_{|t-u| \leqslant s/2} s^2 |\theta''(t)| \leqslant \sup_{|t-u| \leqslant s/2} s^2 |\theta''(t)|.$$

Let us finally compute ε_a when $\xi = \theta'(u)$. Since $g(t) = g(-t)$, we derive from (4.84) that

$$\varepsilon_{a,1}^+ = \frac{2|a'(u)|}{|a(u)|} \left| \int_{-\infty}^{+\infty} t \frac{1}{\sqrt{s}} g_s(t) \, dt \right| = 0 .$$

We also derive from (2.22) that the Fourier transform of $t \frac{1}{\sqrt{s}} g_s(t)$ is $i \, s \, \hat{g}'(s\omega)$, so (4.87) gives

$$\varepsilon_a = \frac{1}{2} \varepsilon_{a,1}^- = \frac{s|a'(u)|}{|a(u)|} |\hat{g}'(2s\theta'(u))| .$$

∎   ∎

Delprat et al. [203] give a different proof of a similar result when $g(t)$ is a Gaussian, using a stationary phase approximation. If we can neglect the corrective term $\varepsilon(u, \xi)$ we shall see that (4.77) enables us to measure $a(u)$ and $\theta'(u)$ from $Sf(u, \xi)$. This implies that the decomposition $f(t) = a(t) \cos \theta(t)$ is uniquely defined. By reviewing the proof of Theorem 4.6, one can verify that a and θ' are the analytic amplitude and instantaneous frequencies of f.

The expressions (4.79, 4.80) show that the three corrective terms $\varepsilon_{a,1}$, $\varepsilon_{a,2}$ and $\varepsilon_{\theta,2}$ are small if $a(t)$ and $\theta'(t)$ have small relative variations over the support of the window g_s. Let $\Delta \omega$ be the bandwidth of $\hat{g}$ defined by

$$|\hat{g}(\omega)| \ll 1 \quad \text{for } |\omega| \geqslant \Delta \omega. \tag{4.90}$$

The term $\sup\limits_{|\omega| \geqslant s|\theta'(u)|} |\hat{g}(\omega)|$ of $\varepsilon(u, \xi)$ is negligible if

$$\theta'(u) \geqslant \frac{\Delta \omega}{s} \ .$$

Ridge Points Let us suppose that $a(t)$ and $\theta'(t)$ have small variations over intervals of size s and that $\theta'(t) \geqslant \Delta\omega/s$ so that the corrective term $\varepsilon(u, \xi)$ in (4.77) can be neglected. Since $|\hat{g}(\omega)|$ is maximum at $\omega = 0$, (4.77) shows that for each u the spectrogram $|Sf(u, \xi)|^2 = |\langle f, g_{s,u,\xi} \rangle|^2$ is maximum at $\xi(u) = \theta'(u)$. The corresponding time-frequency points $(u, \xi(u))$ are called *ridges*. At ridge points, (4.77) becomes

$$Sf(u, \xi) = \frac{\sqrt{s}}{2} a(u) \exp(i[\theta(u) - \xi u]) \left(\hat{g}(0) + \varepsilon(u, \xi) \right) . \tag{4.91}$$

Theorem 4.6 proves that the $\varepsilon(u, \xi)$ is smaller at a ridge point because the first order term $\varepsilon_{a,1}$ becomes negligible in (4.81). This is shown by verifying that $|\hat{g}'(2s\theta'(u))|$ is negligible when $s\theta'(u) \geqslant \Delta\omega$. At ridge points, the second order terms $\varepsilon_{a,2}$ and $\varepsilon_{\theta,2}$ are predominant in $\varepsilon(u, \xi)$.

The ridge frequency gives the instantaneous frequency $\xi(u) = \theta'(u)$ and the amplitude is calculated by

$$a(u) = \frac{2 \left| Sf(u, \xi(u)) \right|}{\sqrt{s} \left| \hat{g}(0) \right|}. \tag{4.92}$$

Let $\Theta_S(u, \xi)$ be the complex phase of $Sf(u, \xi)$. If we neglect the corrective term, then (4.91) proves that ridges are also points of stationary phase:

$$\frac{\partial \Theta_S(u, \xi)}{\partial u} = \theta'(u) - \xi = 0.$$

Testing the stationarity of the phase locates the ridges more precisely.

Multiple Frequencies When the signal contains several spectral lines whose frequencies are sufficiently apart, the windowed Fourier transform separates each of these components and the ridges detect the evolution in time of each spectral component. Let us consider

$$f(t) = a_1(t) \cos \theta_1(t) + a_2(t) \cos \theta_2(t),$$

where $a_k(t)$ and $\theta'_k(t)$ have small variations over intervals of size s and $s\theta'_k(t) \geqslant \Delta\omega$. Since the windowed Fourier transform is linear, we apply (4.77) to each spectral component and neglect the corrective terms:

$$\begin{aligned} Sf(u, \xi) &= \frac{\sqrt{s}}{2} a_1(u) \hat{g}(s[\xi - \theta'_1(u)]) \exp(i[\theta_1(u) - \xi u]) \\ &+ \frac{\sqrt{s}}{2} a_2(u) \hat{g}(s[\xi - \theta'_2(u)]) \exp(i[\theta_2(u) - \xi u]) . \end{aligned} \tag{4.93}$$

The two spectral components are discriminated if for all u

$$\hat{g}(s|\theta'_1(u) - \theta'_2(u)|) \ll 1, \tag{4.94}$$

which means that the frequency difference is larger than the bandwidth of $\hat{g}(s\omega)$:

$$|\theta'_1(u) - \theta'_2(u)| \geqslant \frac{\Delta\omega}{s}. \qquad (4.95)$$

In this case, when $\xi = \theta'_1(u)$, the second term of (4.93) can be neglected and the first term generates a ridge point from which we may recover $\theta'_1(u)$ and $a_1(u)$, using (4.92). Similarly, if $\xi = \theta'_2(u)$ the first term can be neglected and we have a second ridge point that characterizes $\theta'_2(u)$ and $a_2(u)$. The ridge points are distributed along two time-frequency lines $\xi(u) = \theta'_1(u)$ and $\xi(u) = \theta'_2(u)$. This result is valid for any number of time varying spectral components, as long as the distance between any two instantaneous frequencies satisfies (4.95). If two spectral lines are too close, they interfere, which destroys the ridge pattern.

Time-Frequency Ridge Support   The number of instantaneous frequencies is typically unknown. The ridge support Λ is thus defined as the set of all (u,ξ) which are local maxima of $|Sf(u,\xi)|^2$ for u fixed and ξ varying and points of stationary phase $\partial\Theta_S(u,\xi)/\partial u \approx 0$. This support is often reduced by removing small ridge amplitudes $|Sf(u,\xi)|$ that are mostly dominated by the noise, or because smaller ridges may be "shadows" of other instantaneous frequencies created by the side-lobes of $\hat{g}(\omega)$.

Let $\{g_{s,u,\xi}\}_{(u,\xi)\in\Lambda}$ be the set of ridge atoms. For discrete signals, there is a finite number of ridge points, which thus define a frame of the space $\mathbf{V}_\Lambda$ they generate. A ridge signal approximation is computed as an orthogonal projection of f on $\mathbf{V}_\Lambda$. Section 5.1.3 shows that it is obtained with the dual frame $\{\tilde{g}_{\Lambda,u,\xi}\}_{(u,\xi)\in\Lambda}$ of $\{g_{s,u,\xi}\}_{(u,\xi)\in\Lambda}$ in $\mathbf{V}_\Lambda$:

$$f_\Lambda = \sum_{(u,\xi)\in\Lambda} Sf(u,\xi)\,\tilde{g}_{\Lambda,u,\xi} \qquad (4.96)$$

The dual synthesis algorithm of Section 5.1.3 computes this orthogonal projection by inverting the symmetric operator

$$Lh = \sum_{(u,\xi)\in\Lambda} \langle h, g_{s,u,\xi}\rangle\, g_{s,u,\xi} \;. \qquad (4.97)$$

The inversion requires to iterate on this operator many times. If there are few ridge points then (4.97) is efficiently computed by evaluating the inner product and the sum only for $(u,\xi) \in \Lambda$. If there many ridge points, it can be more efficient to compute the full windowed Fourier transform $Sh(u,\xi) = \langle h, g_{s,u,\xi}\rangle$ with the FFT algorithm described in Section 4.2.3, set to zero all coefficients for $(u,\xi) \notin \Lambda$, and apply the fast inverse windowed Fourier transform over all coefficients. The normalization factor N^{-1} in (4.28) must be removed (set to 1) to implement (4.97).
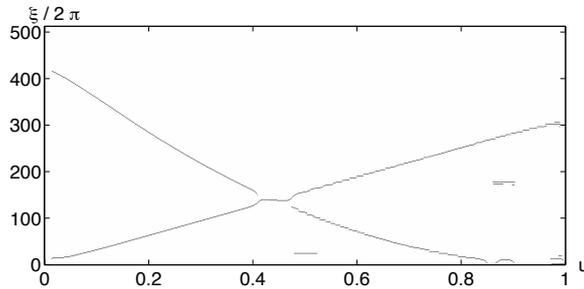
Figure 4.12: Support of larger amplitude ridges calculated from the spectrogram in Figure 4.3. These ridges give the instantaneous frequencies of the linear and quadratic chirps, and of the low and high frequency transients at $t = 0.5$ and $t = 0.87$.

Figure 4.12 displays the ridge support computed from the modulus and phase of the windowed Fourier transform shown in Figure 4.3. For $t \in [0.4, 0.5]$, the instantaneous frequencies of the linear chirp and the quadratic chirps are close and the frequency resolution of the window is not sufficient to discriminate them. As a result, the ridges detect a single average instantaneous frequency.

Time Scaling and Frequency Transpositions A reduction of sound duration by a factor α is implemented according to the deformation model (4.72), by dilating in time the ridge support Λ:

$$\Lambda_\alpha = \{(u, \xi) \ : \ (\alpha u, \xi) \in \Lambda\} \ . \tag{4.98}$$

The windowed-Fourier coefficients $c(u, \xi)$ in Λ_α are derived from the modulus and phase of ridge coefficients

$$\forall (v, \xi) \in \Lambda_\alpha \ , \quad c(v, \xi) = |Sf(\alpha v, \xi)| \, e^{i\Theta_S(\alpha v, \xi)/\alpha} \ . \tag{4.99}$$

The scaled signal is reconstructed from these coefficients, with the dual synthesis algorithm of Section 5.1.3, as in (4.96):

$$f_\alpha = \sum_{(v, \xi) \in \Lambda_\alpha} c(v, \xi) \, \tilde{g}_{\Lambda_\alpha, v, \xi} \ .$$

Similarly, a sound transposition is implemented according to the transposition model (4.73) by dilating in frequency the ridge support Λ:

$$\Lambda_\alpha = \{(u, \xi) \ : \ (u, \alpha\xi) \in \Lambda\} \ . \tag{4.100}$$

The transposed coefficient amplitudes $|c(u, \xi)|$ in Λ_α are calculated with (4.74). At any fixed time u_0, the ridge amplitudes at all frequencies $\{a(u_0, \xi) = |Sf(u_0, \xi)|\}_{(u_0, \xi) \in \Lambda}$ are mapped to a transposed amplitudes $\{b(u_0, \eta)\}_{(u_0, \eta) \in \Lambda_\alpha}$ at frequencies $\eta = \xi/\alpha$, by computing a frequency envelop. The resulting ridge coefficients are

$$\forall (u, \eta) \in \Lambda_\alpha \ , \quad c(u, \eta) = b(u, \eta) \, e^{i\Theta_S(u, \alpha\eta)/\alpha} \ . \tag{4.101}$$

The transposed signal is reconstructed with the dual synthesis algorithm of Section 5.1.3:

$$f_\alpha = \sum_{(u, \eta) \in \Lambda_\alpha} c(u, \eta) \, \tilde{g}_{\Lambda_\alpha, u, \eta} \ .$$

Choice of Window The measurement of instantaneous frequencies at ridge points is valid only if the size s of the window g_s is sufficiently small so that the second order terms $\varepsilon_{a,2}$ and $\varepsilon_{\theta,2}$ in (4.79) and (4.80) are small:

$$\sup_{|t-u| \leqslant s/2} \frac{s^2 \, |a_k''(t)|}{|a_k(u)|} \ll 1 \quad \text{and} \quad \sup_{|t-u| \leqslant s/2} s^2 |\theta_k''(t)| \ll 1 \ . \tag{4.102}$$

On the other hand, the frequency bandwidth $\Delta\omega/s$ must also be sufficiently small to discriminate consecutive spectral components in (4.95). The window scale s must therefore be adjusted as a trade-off between both constraints.

Table 4.1 gives the spectral parameters of several windows of compact support. For instantaneous frequency detection, it is particularly important to ensure that $\hat{g}$ has negligible side-lobes at $\pm\omega_0$, as illustrated by Figure 4.4. The reader can verify with (4.77) that these side-lobes "react" to an instantaneous frequency $\theta'(u)$ by creating shadow maxima of $|Sf(u, \xi)|^2$ at frequencies $\xi = \theta'(u) \pm \omega_0$. The ratio of the amplitude of these shadow maxima to the amplitude of the main local maxima at $\xi = \theta'(u)$ is $|\hat{g}(\omega_0)|^2 \, |\hat{g}(0)|^{-2}$. They can be removed by thresholding or by testing the stationarity of the phase.

Example 4.13. *The sum of two parallel linear chirps*

$$f(t) = a_1 \cos(bt^2 + ct) + a_2 \cos(bt^2) \tag{4.103}$$

has two instantaneous frequencies $\theta_1'(t) = 2bt + c$ and $\theta_2'(t) = 2bt$. Figure 4.13 gives a numerical example. The window g_s has enough frequency resolution to discriminate both chirps if

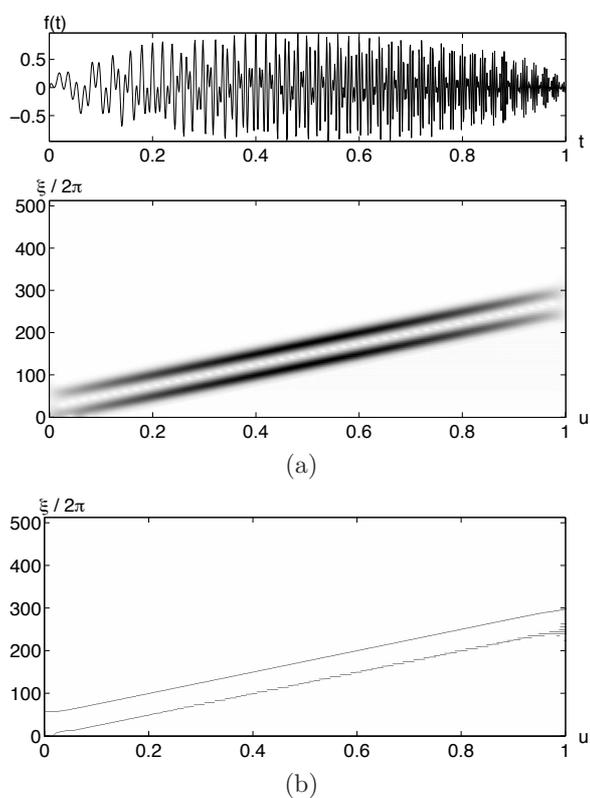$$|\theta_1'(t) - \theta_2'(t)| = |c| \geqslant \frac{\Delta\omega}{s}. \tag{4.104}$$

Figure 4.13: Sum of two parallel linear chirps. (a): Spectrogram $P_S f(u, \xi) = |Sf(u, \xi)|^2$. (b): Ridge support calculated from the spectrogram.

Its time support is small enough compared to their time variation if

$$s^2 \, |\theta_1''(u)| = s^2 \, |\theta_2''(u)| = 2 \, b \, s^2 \ll 1. \tag{4.105}$$

Conditions (4.104) and (4.105) prove that there exists an appropriate window g if and only if

$$\frac{c}{\sqrt{b}} \gg \Delta\omega. \tag{4.106}$$

Since g is a smooth window with a support $[-1/2, 1/2]$, its frequency bandwidth $\Delta\omega$ is on the order of 1. The linear chirps in Figure 4.13 satisfy (4.106). Their ridges are computed with the truncated Gaussian window of Table 4.1, with $s = 0.5$.
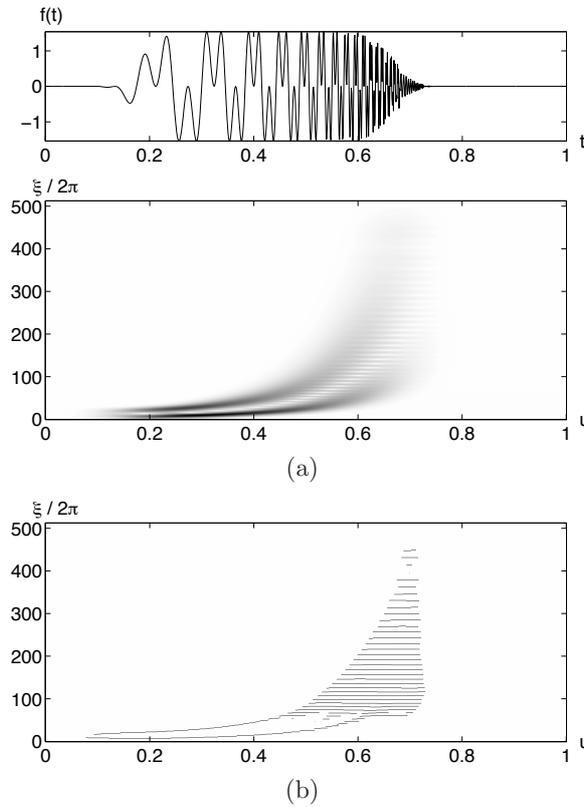
Figure 4.14: Sum of two hyperbolic chirps. (a): Spectrogram $P_S f(u, \xi)$. (b): Ridge support calculated from the spectrogram

Example 4.14. *The hyperbolic chirp*

$$f(t) = \cos\left(\frac{\alpha}{\beta - t}\right)$$

for $0 \leqslant t < \beta$ has an instantaneous frequency

$$\theta'(t) = \frac{\alpha}{(\beta - t)^2},$$

which varies quickly when t is close to β. The instantaneous frequency of hyperbolic chirps goes from 0 to $+\infty$ in a finite time interval. This is particularly useful for radars. These chirps are also emitted by the cruise sonars of bats [203].

The instantaneous frequency of hyperbolic chirps cannot be estimated with a windowed Fourier transform because for any fixed window size the instantaneous frequency varies too quickly at high frequencies. When u is close enough to β then (4.102) is not satisfied because

$$s^2 |\theta''(u)| = \frac{s^2 \alpha}{(\beta - u)^3} > 1.$$

Figure 4.14 shows a signal that is a sum of two hyperbolic chirps:

$$f(t) = a_1 \cos\left(\frac{\alpha_1}{\beta_1 - t}\right) + a_2 \cos\left(\frac{\alpha_2}{\beta_2 - t}\right),\tag{4.107}$$

with $\beta_1 = 0.68$ and $\beta_2 = 0.72$. At the beginning of the signal, the two chirps have close instantaneous frequencies that are discriminated by the windowed Fourier ridge computed with a large size window. When getting close to β_1 and β_2, the instantaneous frequency varies too quickly relative to the window size. The resulting ridges cannot follow these instantaneous frequencies.

4.4.2 Wavelet Ridges

Windowed Fourier atoms have a fixed scale and thus cannot follow the instantaneous frequency of rapidly varying events such as hyperbolic chirps. In contrast, an analytic wavelet transform modifies the scale of its time-frequency atoms. The ridge algorithm of Delprat et al. [203] is extended to analytic wavelet transforms to accurately measure frequency tones that are rapidly changing at high frequencies.

An approximately analytic wavelet is constructed in (4.60) by multiplying a window g with a sinusoidal wave:

$$\psi(t) = g(t) \exp(i\eta t).$$

As in the previous section, g is a symmetric window with a support equal to $[-1/2, 1/2]$, and a unit norm $\|g\| = 1$. Let $\Delta\omega$ be the bandwidth of $\hat{g}$ defined in (4.90). If $\eta > \Delta\omega$ then

$$\forall \omega < 0 \quad , \quad \hat{\psi}(\omega) = \hat{g}(\omega - \eta) \ll 1 .$$

The wavelet ψ is not strictly analytic because its Fourier transform is not exactly equal to zero at negative frequencies.

Dilated and translated wavelets can be rewritten

$$\psi_{u,s}(t) = \frac{1}{\sqrt{s}} \psi\left(\frac{t - u}{s}\right) = g_{s,u,\xi}(t) \exp(-i\xi u),$$

with $\xi = \eta/s$ and

$$g_{s,u,\xi}(t) = \sqrt{s}\, g\left(\frac{t - u}{s}\right) \exp(i\xi t).$$

The resulting wavelet transform uses time-frequency atoms similar to those of a windowed Fourier transform (4.76) but in this case the scale s varies over $\mathbb{R}^+$ while $\xi = \eta/s$:

$$Wf(u, s) = \langle f, \psi_{u,s} \rangle = \langle f, g_{s,u,\xi} \rangle \exp(i\xi u) .$$

Theorem 4.6 computes $\langle f, g_{s,u,\xi} \rangle$ when $f(t) = a(t) \cos \theta(t)$, which gives

$$Wf(u, s) = \frac{\sqrt{s}}{2} a(u) \exp[i\theta(u)]\Big(\hat{g}(s[\xi - \theta'(u)]) + \varepsilon(u, \xi)\Big).\tag{4.108}$$

The corrective term $\varepsilon(u, \xi)$ is negligible if $a(t)$ and $\theta'(t)$ have small variations over the support of $\psi_{u,s}$ and if $\theta'(u) \geqslant \Delta\omega/s$.

Ridge Detection The instantaneous frequency is measured from ridges defined over the wavelet transform. The normalized scalogram defined by

$$\frac{\xi}{\eta} \, P_W f(u, \xi) = \frac{|Wf(u, s)|^2}{s} \quad \text{for } \xi = \eta/s$$

is calculated with (4.108):

$$\frac{\xi}{\eta} \, P_W f(u, \xi) = \frac{1}{4} \, a^2(u) \left| \hat{g}\Big(\eta \Big[1 - \frac{\theta'(u)}{\xi} \Big] \Big) + \varepsilon(u, \xi) \right|^2 .$$

Since $|\hat{g}(\omega)|$ is maximum at $\omega = 0$, if we neglect $\varepsilon(u, \xi)$, this expression shows that the scalogram is maximum at

$$\frac{\eta}{s(u)} = \xi(u) = \theta'(u) \, . \tag{4.109}$$

The corresponding points $(u, \xi(u))$ are called *wavelet ridges*. The analytic amplitude is given by

$$a(u) = \frac{2\,\sqrt{\eta^{-1}\xi \, P_W f(u, \xi)}}{|\hat{g}(0)|}. \tag{4.110}$$

The complex phase of $Wf(u, s)$ in (4.108) is $\Theta_W(u, \xi) = \theta(u)$. At ridge points,

$$\frac{\partial \Theta_W(u, \xi)}{\partial u} = \theta'(u) = \xi. \tag{4.111}$$

When $\xi = \theta'(u)$, the first order term $\varepsilon_{a,1}$ calculated in (4.81) becomes negligible. The corrective term is then dominated by $\varepsilon_{a,2}$ and $\varepsilon_{\theta,2}$. To simplify their expression, we approximate the sup of a'' and θ'' in the neighborhood of u by their value at u. Since $s = \eta/\xi = \eta/\theta'(u)$, (4.79) and (4.80) imply that these second order terms become negligible if

$$\frac{\eta^2}{|\theta'(u)|^2} \frac{|a''(u)|}{|a(u)|} \ll 1 \quad \text{and} \quad \eta^2 \frac{|\theta''(u)|}{|\theta'(u)|^2} \ll 1. \tag{4.112}$$

The presence of θ' in the denominator proves that a' and θ' must have slow variations if θ' is small but may vary much more quickly for large instantaneous frequencies.

Multispectral Estimation Suppose that f is a sum of two spectral components:

$$f(t) = a_1(t) \cos \theta_1(t) + a_2(t) \cos \theta_2(t).$$

As in (4.94), we verify that the second instantaneous frequency θ'_2 does not interfere with the ridge of θ'_1 if the dilated window has a sufficient spectral resolution at the ridge scale $s = \eta/\xi = \eta/\theta'_1(u)$:

$$\hat{g}(s|\theta'_1(u) - \theta'_2(u)|) \ll 1. \tag{4.113}$$

Since the bandwidth of $\hat{g}(\omega)$ is $\Delta\omega$, this means that

$$\frac{|\theta'_1(u) - \theta'_2(u)|}{\theta'_1(u)} \geqslant \frac{\Delta\omega}{\eta}. \tag{4.114}$$

Similarly, the first spectral component does not interfere with the second ridge located at $s = \eta/\xi = \eta/\theta'_2(u)$ if

$$\frac{|\theta'_1(u) - \theta'_2(u)|}{\theta'_2(u)} \geqslant \frac{\Delta\omega}{\eta}. \tag{4.115}$$

To separate spectral lines whose instantaneous frequencies are close, these conditions prove that the wavelet must have a small octave bandwidth $\Delta\omega/\eta$. The bandwidth $\Delta\omega$ is a fixed constant, which is on the order of 1. The frequency η is a free parameter whose value is chosen as a trade-off between the time-resolution condition (4.112) and the frequency bandwidth conditions (4.114) and (4.115).

Figure 4.15 displays the ridges computed from the normalized scalogram and the wavelet phase shown in Figure 4.11. The ridges of the high frequency transient located at $t = 0.87$ have oscillations because of the interferences with the linear chirp above. The frequency separation condition (4.114) is not satisfied. This is also the case in the time interval $[0.35, 0.55]$, where the instantaneous frequencies of the linear and quadratic chirps are too close.
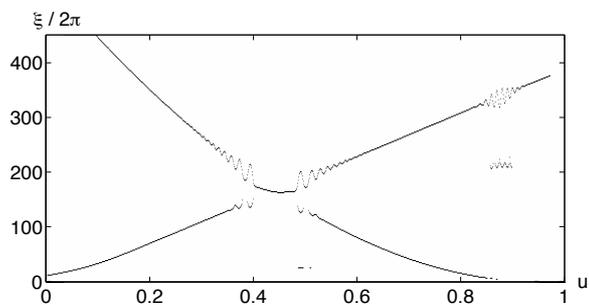
Figure 4.15: Ridge support calculated from the scalogram shown in Figure 4.11. Compare with the windowed Fourier ridges in Figure 4.12.
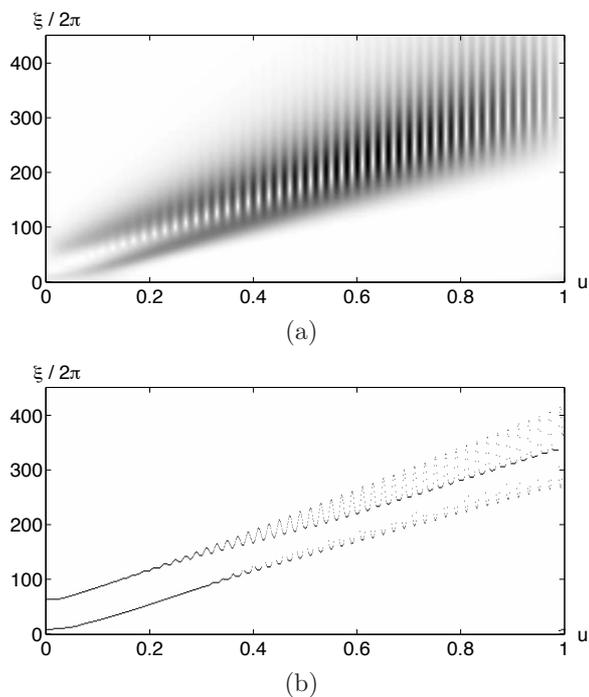

(a)

(b)

Figure 4.16: (a): Normalized scalogram $\eta^{-1}\xi P_W f(u,\xi)$ of two parallel linear chirps shown in Figure 4.13. (b): Wavelet ridges.

Ridge Support and Processing   The wavelet ridge support Λ of f is the set of all ridge points (u, s) in the time-scale plane or $(u, \xi = \eta/s)$ in the time-frequency plane, corresponding to local maxima of $|Wf(u, s)|/s$ for a fixed u and s varying, where the complex phase $\Theta_W(u, s)$ nearly satisfies (4.111).

As in the windowed Fourier case, an orthogonal projection is computed over the space $\mathbf{V}_\Lambda$ generated by the ridge wavelets $\{\psi_{u,s}\}_{(u,s)\in\Lambda}$, by using the dual wavelet frame $\{\tilde{\psi}_{\Lambda,u,s}\}_{(u,s)\in\Lambda}$:

$$f_\Lambda = \sum_{(u,s)\in\Lambda} Wf(u, s)\, \tilde{\psi}_{\Lambda,u,s} \; , \tag{4.116}$$

It is implemented with the dual synthesis algorithm of Section 5.1.3, by inverting the symmetric operator

$$Lh = \sum_{(u,s)\in\Lambda} \langle h, \psi_{u,s}\rangle\, \psi_{u,s} \; , \tag{4.117}$$

which is performed by computing this operator many times. When there are many ridge points, instead of computing this sum only for $(u, s) \in \Lambda$, it may require less operations to compute $Wf(u, s)$ with the fast wavelet transform algorithm of Section 4.3.3. All coefficients $(u, s) \notin \Lambda$ are set to zero, and the fast inverse wavelet transform algorithm is applied. The inverse wavelet transform formula (4.66) must be modified by removing the renormalization factor a^{-j} and a^{-J} in the sum (set them to 1), to implement the operator (4.117).

Like in the windowed Fourier case, modifications of the sound durations or frequency transpositions are computed with by modifying the ridge support. A reduction of sound duration by a factor α transforms the ridge support Λ into:

$$\Lambda_\alpha = \{(u, s) \; : \; (\alpha u, s) \in \Lambda\} \; . \tag{4.118}$$

A sound transposition is implemented by modifying the scales of the time-scale ridge support Λ, which defines:

$$\Lambda_\alpha = \{(u, s) \; : \; (u, s/\alpha) \in \Lambda\} \; . \tag{4.119}$$

The wavelet coefficients over these supports are derived from the deformation model (4.72) or (4.74), similarly to (4.99) and (4.101) for the windowed Fourier transform. Processed signal are recovered from the modified wavelet coefficients and modified supports with the dual synthesis algorithm of Section 5.1.3.

Example 4.15. *The instantaneous frequencies of two linear chirps*

$$f(t) = a_1 \cos(b\,t^2 + c\,t) + a_2 \cos(b\,t^2)$$

are not well measured by wavelet ridges. Indeed

$$\frac{|\theta_2'(u) - \theta_1'(u)|}{\theta_1'(u)} = \frac{c}{b\,t}$$

converges to zero when t increases. When it is smaller than $\Delta\omega/\eta$ the two chirps interact and create interference patterns like those in Figure 4.16. The ridges follow these interferences and do not estimate properly the two instantaneous frequencies, as opposed to the windowed Fourier ridges shown in Figure 4.13.

Example 4.16. *The instantaneous frequency of a hyperbolic chirp*

$$f(t) = \cos\left(\frac{\alpha}{\beta - t}\right)$$

is $\theta'(t) = \alpha\,(1 - t)^{-2}$. Wavelet ridges can measure this instantaneous frequency if the time resolution condition (4.112) is satisfied:

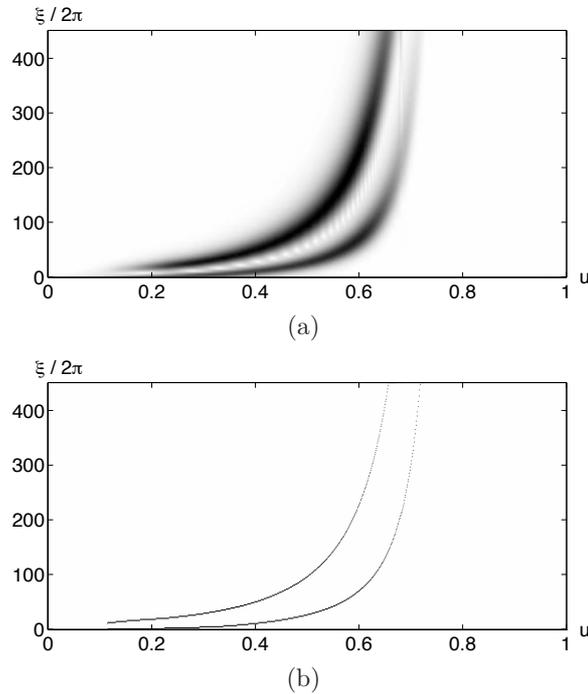$$\eta^2 \ll \frac{\theta'(t)^2}{|\theta''(t)|} = \frac{\alpha}{|t - \beta|}.$$

Figure 4.17: (a): Normalized scalogram $\eta^{-1}\xi P_W f(u,\xi)$ of two hyperbolic chirps shown in Figure 4.14. (b): Wavelet ridges.

This is the case if $|t - \beta|$ is not too large.

   Figure 4.17 displays the scalogram and the ridges of two hyperbolic chirps

$$f(t) = a_1 \cos\left(\frac{\alpha_1}{\beta_1 - t}\right) + a_2 \cos\left(\frac{\alpha_2}{\beta_2 - t}\right),$$

with $\beta_1 = 0.68$ and $\beta_2 = 0.72$. As opposed to the windowed Fourier ridges shown in Figure 4.14, the wavelet ridges follow the rapid time modification of both instantaneous frequencies. This is particularly useful in analyzing the returns of hyperbolic chirps emitted by radars or sonars. Several techniques have been developed to detect chirps with wavelet ridges in presence of noise [150, 455].

Better is More Sparse   The linear and hyperbolic chirp examples show that the best transform depends upon the signal time-frequency property. All examples also show that when the time-frequency transform has a resolution adapted to the signal time-frequency properties then the number of ridge points is reduced. Indeed, if signal structures do not match the dictionary time-frequency atoms then their energy is diffused over many more atoms, which produces more local maxima. Sparsity thus appears as a natural criteria to adjust the resolution of time-frequency transforms.

   Section 12.3.3 studies sparse time-frequency decompositions in very redundant Gabor time-frequency dictionaries, including windowed Fourier atoms and wavelet atoms, with a computationally more intensive matching pursuit algorithm.

4.5 Quadratic Time-Frequency Energy

The wavelet and windowed Fourier transforms are computed by correlating the signal with families of time-frequency atoms. The time and frequency resolution of these transforms is thus limited by the time-frequency resolution of the corresponding atoms. Ideally, one would like to define a density of energy in a time-frequency plane, with no loss of resolution.

The Wigner-Ville distribution is a time-frequency energy density computed by correlating f with a time and frequency translation of itself. Despite its remarkable properties, the application of Wigner-Ville distributions is limited by the existence of interference terms. These interferences can be attenuated by a time-frequency averaging, but this results in a loss of resolution. It is proved that the spectrogram, the scalogram and all squared time-frequency decompositions can be written as a time-frequency averaging of the Wigner-Ville distribution, which gives a common framework to relate these transforms.

4.5.1 Wigner-Ville Distribution

To analyze time-frequency structures, in 1948 Ville [475] introduced in signal processing a quadratic form that had been studied by Wigner [484] in a 1932 article on quantum thermodynamics:

$$P_V f(u, \xi) = \int_{-\infty}^{+\infty} f\left(u + \frac{\tau}{2}\right) f^*\left(u - \frac{\tau}{2}\right) e^{-i\tau\xi} \, d\tau. \tag{4.120}$$

The Wigner-Ville distribution remains real because it is the Fourier transform of $f(u+\tau/2)f^*(u-\tau/2)$, which has a Hermitian symmetry in τ. Time and frequency have a symmetrical role. This distribution can also be rewritten as a frequency integration by applying the Parseval formula:

$$P_V f(u, \xi) = \frac{1}{2\pi} \int_{-\infty}^{+\infty} \hat{f}\left(\xi + \frac{\gamma}{2}\right) \hat{f}^*\left(\xi - \frac{\gamma}{2}\right) e^{i\gamma u} \, d\gamma. \tag{4.121}$$

Time-Frequency Support   The Wigner-Ville transform localizes the time-frequency structures of f. If the energy of f is well concentrated in time around u_0 and in frequency around ξ_0 then $P_V f$ has its energy centered at (u_0, ξ_0), with a spread equal to the time and frequency spread of f. This property is illustrated by the following theorem, which relates the time and frequency support of $P_V f$ to the support of f and $\hat{f}$.

Theorem 4.7.   • *If the support of f is $[u_0 - T/2, u_0 + T/2]$, then for all ξ the support in u of $P_V f(u, \xi)$ is included in this interval.*

• *If the support of $\hat{f}$ is $[\xi_0 - \Delta/2, \xi_0 + \Delta/2]$, then for all u the support in ξ of $P_V f(u, \xi)$ is included in this interval.*

Proof. Let $\bar{f}(t) = f(-t)$. The Wigner-Ville distribution is rewritten

$$P_V f(u, \xi) = \int_{-\infty}^{+\infty} f\left(\frac{\tau + 2u}{2}\right) \bar{f}^*\left(\frac{\tau - 2u}{2}\right) e^{-i\xi\tau} \, d\tau. \tag{4.122}$$

Suppose that f has a support equal to $[u_0 - T/2, u_0 + T/2]$. The supports of $f(\tau/2 + u)$ and $\bar{f}(\tau/2 - u)$ are then respectively

$$[2(u_0 - u) - T, 2(u_0 - u) + T] \quad \text{and} \quad [-2(u_0 + u) - T, -2(u_0 + u) + T].$$

The Wigner-Ville integral (4.122) shows that $P_V f(u, \xi)$ is non-zero if these two intervals overlap, which is the case only if $|u_0 - u| < T$. The support of $P_V f(u, \xi)$ along u is therefore included in the support of f. If the support of $\hat{f}$ is an interval, then the same derivation based on (4.121) shows that the support of $P_V f(u, \xi)$ along ξ is included in the support of $\hat{f}$. ■   ■

Example 4.17. *Theorem 4.7 proves that the Wigner-Ville distribution does not spread the time or frequency support of Diracs or sinusoids, unlike windowed Fourier and wavelet transforms. Direct calculations yield*

$$f(t) = \delta(u - u_0) \quad \Longrightarrow \quad P_V f(u, \xi) = \delta(u - u_0), \tag{4.123}$$

$$f(t) = \exp(i\xi_0 t) \quad \Longrightarrow \quad P_V f(u, \xi) = \frac{1}{2\pi} \delta(\xi - \xi_0). \tag{4.124}$$

Example 4.18. *If f is a smooth and symmetric window then its Wigner-Ville distribution $P_V f(u, \xi)$ is concentrated in a neighborhood of $u = \xi = 0$. A Gaussian $f(t) = (\sigma^2 \pi)^{-1/4} \exp(-t^2/(2\sigma^2))$ is transformed into a two-dimensional Gaussian because its Fourier transform is also a Gaussian (2.32) and one can verify that*

$$P_V f(u, \xi) = \frac{1}{\pi} \exp\left(\frac{-u^2}{\sigma^2} - \sigma^2 \xi^2 \right). \tag{4.125}$$

In this particular case $P_V f(u, \xi) = |f(u)|^2 |\hat{f}(\xi)|^2$.

The Wigner-Ville distribution has important invariance properties. A phase shift does not modify its value:

$$f(t) = e^{i\theta} g(t) \implies P_V f(u, \xi) = P_V g(u, \xi) . \tag{4.126}$$

When f is translated in time or frequency, its Wigner-Ville transform is also translated:

$$f(t) = g(t - u_0) \quad \implies \quad P_V f(u, \xi) = P_V g(u - u_0, \xi) , \tag{4.127}$$

$$f(t) = \exp(i\xi_0 t) g(t) \quad \implies \quad P_V f(u, \xi) = P_V g(u, \xi - \xi_0) . \tag{4.128}$$

If f is scaled by s and thus $\hat{f}$ is scaled by $1/s$ then the time and frequency parameters of $P_V f$ are also scaled respectively by s and $1/s$

$$f(t) = \frac{1}{\sqrt{s}} g\left(\frac{t}{s} \right) \implies P_V f(u, \xi) = P_V f\left(\frac{u}{s}, s\xi \right) . \tag{4.129}$$

Example 4.19. *If g is a smooth and symmetric window then $P_V g(u, \xi)$ has its energy concentrated in the neighborhood of $(0, 0)$. The time-frequency atom*

$$f_0(t) = \frac{a}{\sqrt{s}} \exp(i\theta_0) g\left(\frac{t - u_0}{s} \right) \exp(i\xi_0 t) \quad .$$

has a Wigner-Ville distribution that is calculated with (4.126), (4.127) and (4.128):

$$P_V f_0(u, \xi) = |a|^2 P_V g\left(\frac{u - u_0}{s}, s(\xi - \xi_0) \right) \quad . \tag{4.130}$$

Its energy is thus concentrated in the neighborhood of (u_0, ξ_0), on an ellipse whose axes are proportional to s in time and $1/s$ in frequency.

Instantaneous Frequency   Ville's original motivation for studying time-frequency decompositions was to compute the instantaneous frequency of a signal [475]. Let f_a be the analytic part of f obtained in (4.69) by setting to zero $\hat{f}(\omega)$ for $\omega < 0$. We write $f_a(t) = a(t) \exp[i\theta(t)]$ to define the instantaneous frequency $\omega(t) = \theta'(t)$. The following theorem proves that $\theta'(t)$ is the "average" frequency computed relative to the Wigner-Ville distribution $P_V f_a$.

Theorem 4.8. *If $f_a(t) = a(t) \exp[i\theta(t)]$ then*

$$\theta'(u) = \frac{\int_{-\infty}^{+\infty} \xi\, P_V f_a(u, \xi)\, d\xi}{\int_{-\infty}^{+\infty} P_V f_a(u, \xi)\, d\xi}. \tag{4.131}$$

Proof. To prove this result, we verify that any function g satisfies

$$\iint \xi\, g\left(u + \frac{\tau}{2} \right) g^*\left(u - \frac{\tau}{2} \right) \exp(-i\tau\xi)\, d\xi\, d\tau = -\pi i \left[g'(u)\, g^*(u) - g(u)\, g^{*'}(u) \right]. \tag{4.132}$$

This identity is obtained by observing that the Fourier transform of $i\xi$ is the derivative of a Dirac, which gives an equality in the sense of distributions:

$$\int_{-\infty}^{+\infty} \xi \exp(-i\tau\xi)\, d\xi = -i\, 2\pi\, \delta'(\tau).$$

Since $\int_{-\infty}^{+\infty} \delta'(\tau) h(\tau) d\tau = -h'(0)$, inserting $h(\tau) = g(u + \tau/2) g^*(u - \tau/2)$ proves (4.132). If $g(u) = f_a(u) = a(u) \exp[i\theta(u)]$ then (4.132) gives

$$\int_{-\infty}^{+\infty} \xi \, P_V f_a(u, \xi) \, d\xi = 2\pi \, a^2(u) \, \theta'(u).$$

We will see in (4.136) that $|f_a(u)|^2 = (2\pi)^{-1} \int_{-\infty}^{+\infty} P_V f_a(u, \xi) \, d\xi$, and since $|f_a(u)|^2 = a(u)^2$ we derive (4.131). ■ ■

This theorem shows that for a fixed u the mass of $P_V f_a(u, \xi)$ is typically concentrated in the neighborhood of the instantaneous frequency $\xi = \theta'(u)$. For example, a linear chirp $f(t) = \exp(iat^2)$ is transformed into a Dirac located along the instantaneous frequency $\xi = \theta'(u) = 2au$:

$$P_V f(u, \xi) = \delta(\xi - 2au).$$

Similarly, the multiplication of f by a linear chirp $\exp(iat^2)$ makes a frequency translation of $P_V f$ by the instantaneous frequency $2au$:

$$f(t) = \exp(iat^2) \, g(t) \implies P_V f(u, \xi) = P_V g(u, \xi - 2au) . \tag{4.133}$$

Energy Density The Moyal [378] formula proves that the Wigner-Ville transform is unitary, which implies energy conservation properties.

Theorem 4.9 (Moyal). *For any f and g in $\mathbf{L}^2(\mathbb{R})$*

$$\left| \int_{-\infty}^{+\infty} f(t) \, g^*(t) \, dt \right|^2 = \frac{1}{2\pi} \iint P_V f(u, \xi) \, P_V g(u, \xi) \, du \, d\xi. \tag{4.134}$$

Proof. Let us compute the integral

$$\begin{aligned}
I &= \iint P_V f(u, \xi) \, P_V g(u, \xi) \, du \, d\xi \\
&= \iiiint f\left(u + \frac{\tau}{2}\right) f^*\left(u - \frac{\tau}{2}\right) g\left(u + \frac{\tau'}{2}\right) g^*\left(u - \frac{\tau'}{2}\right) \\
&\qquad\qquad\qquad \exp[-i\xi(\tau + \tau')] \, d\tau \, d\tau' \, du \, d\xi.
\end{aligned}$$

The Fourier transform of $h(t) = 1$ is $\hat{h}(\omega) = 2\pi\delta(\omega)$, which means that we have a distribution equality $\int \exp[-i\xi(\tau + \tau')] d\xi = 2\pi\delta(\tau + \tau')$. As a result,

$$\begin{aligned}
I &= 2\pi \iiint f\left(u + \frac{\tau}{2}\right) f^*\left(u - \frac{\tau}{2}\right) g\left(u + \frac{\tau'}{2}\right) g^*\left(u - \frac{\tau'}{2}\right) \delta(\tau + \tau') \, d\tau \, d\tau' \, du \\
&= 2\pi \iint f\left(u + \frac{\tau}{2}\right) f^*\left(u - \frac{\tau}{2}\right) g\left(u - \frac{\tau}{2}\right) g^*\left(u + \frac{\tau}{2}\right) \, d\tau \, du.
\end{aligned}$$

The change of variable $t = u + \tau/2$ and $t' = u - \tau/2$ yields (4.134). ■ ■

One can consider $|f(t)|^2$ and $|\hat{f}(\omega)|^2/(2\pi)$ as energy densities in time and frequency that satisfy a conservation equation:

$$\|f\|^2 = \int_{-\infty}^{+\infty} |f(t)|^2 \, dt = \frac{1}{2\pi} \int_{-\infty}^{+\infty} |\hat{f}(\omega)|^2 \, d\omega.$$

The following theorem shows that these time and frequency densities are recovered with marginal integrals over the Wigner-Ville distribution.

Theorem 4.10. *For any $f \in \mathbf{L}^2(\mathbb{R})$*

$$\int_{-\infty}^{+\infty} P_V f(u, \xi) \, du = |\hat{f}(\xi)|^2, \tag{4.135}$$

and

$$\frac{1}{2\pi} \int_{-\infty}^{+\infty} P_V f(u, \xi) \, d\xi = |f(u)|^2. \tag{4.136}$$

Proof. The frequency integral (4.121) proves that the one-dimensional Fourier transform of $g_\xi(u) = P_V f(u, \xi)$ with respect to u is

$$\hat{g}_\xi(\gamma) = \hat{f}\left(\xi + \frac{\gamma}{2}\right) \hat{f}^*\left(\xi - \frac{\gamma}{2}\right).$$

We derive (4.135) from the fact that is

$$\hat{g}_\xi(0) = \int_{-\infty}^{+\infty} g_\xi(u) \, du.$$

Similarly, (4.120) shows that $P_V f(u, \xi)$ is the one-dimensional Fourier transform of $f(u+\tau/2)f^*(u-\tau/2)$ with respect to τ, where ξ is the Fourier variable. Its integral in ξ thus gives the value for $\tau = 0$, which is the identity (4.136). ■   ■

This theorem suggests interpreting the Wigner-Ville distribution as a joint time-frequency energy density. However, the Wigner-Ville distribution misses one fundamental property of an energy density: positivity. Let us compute for example the Wigner-Ville distribution of $f = \mathbf{1}_{[-T,T]}$ with the integral (4.120):

$$P_V f(u, \xi) = \frac{2 \sin\left(2(T - |u|)\xi\right)}{\xi} \mathbf{1}_{[-T,T]}(u).$$

It is an oscillating function that takes negative values. In fact, one can prove that translated and frequency modulated Gaussians are the only functions whose Wigner-Ville distributions remain positive. As we will see in the next section, to obtain positive energy distributions for all signals, it is necessary to average the Wigner-Ville transform and thus lose some time-frequency resolution.

4.5.2   Interferences and Positivity

At this point, the Wigner-Ville distribution may seem to be an ideal tool for analyzing the time-frequency structures of a signal. This is however not the case because of interferences created by the quadratic properties of this transform. These interferences can be removed by averaging the Wigner-Ville distribution with appropriate kernels which yield positive time-frequency densities. However, this reduces the time-frequency resolution. Spectrograms and scalograms are examples of positive quadratic distributions obtained by smoothing the Wigner-Ville distribution.

Cross Terms   Let $f = f_1 + f_2$ be a composite signal. Since the Wigner-Ville distribution is a quadratic form,

$$P_V f = P_V f_1 + P_V f_2 + P_V[f_1, f_2] + P_V[f_2, f_1], \tag{4.137}$$

where $P_V[h, g]$ is the cross Wigner-Ville distribution of two signals

$$P_V[h, g](u, \xi) = \int_{-\infty}^{+\infty} h\left(u + \frac{\tau}{2}\right) g^*\left(u - \frac{\tau}{2}\right) e^{-i\tau\xi} \, d\tau. \tag{4.138}$$

The interference term

$$I[f_1, f_2] = P_V[f_1, f_2] + P_V[f_2, f_1]$$

is a real function that creates non-zero values at unexpected locations of the (u, ξ) plane.
Let us consider two time-frequency atoms defined by

$$f_1(t) = a_1 \, e^{i\theta_1} \, g(t - u_1) \, e^{i\xi_1 t} \quad \text{and} \quad f_2(t) = a_2 \, e^{i\theta_2} \, g(t - u_2) \, e^{i\xi_2 t},$$

where g is a time window centered at $t = 0$. Their Wigner-Ville distributions computed in (4.130) are

$$P_V f_1(u, \xi) = a_1^2 P_V g(u - u_1, \xi - \xi_1) \quad \text{and} \quad P_V f_2(u, \xi) = a_2^2 P_V g(u - u_2, \xi - \xi_2).$$

Since the energy of $P_V g$ is centered at $(0, 0)$, the energy of $P_V f_1$ and $P_V f_2$ is concentrated in the neighborhoods of (u_1, ξ_1) and (u_2, ξ_2) respectively. A direct calculation verifies that the interference term is

$$I[f_1, f_2](u, \xi) = 2a_1 a_2 \, P_V g(u - u_0, \xi - \xi_0) \cos\left((u - u_0)\Delta\xi - (\xi - \xi_0)\Delta u + \Delta\theta\right)$$

with

$$u_0 = \frac{u_1 + u_2}{2} \quad , \quad \xi_0 = \frac{\xi_1 + \xi_2}{2}$$
$$\Delta u = u_1 - u_2 \quad , \quad \Delta \xi = \xi_1 - \xi_2$$
$$\Delta \theta = \theta_1 - \theta_2 + u_0 \, \Delta \xi.$$

It is an oscillatory waveform centered at the middle point (u_0, ξ_0). This is quite counter-intuitive since f and $\hat{f}$ have very little energy in the neighborhood of u_0 and ξ_0. The frequency of the oscillations is proportional to the Euclidean distance $\sqrt{\Delta \xi^2 + \Delta u^2}$ of (u_1, ξ_1) and (u_2, ξ_2). The direction of these oscillations is perpendicular to the line that joins (u_1, ξ_1) and (u_2, ξ_2). Figure 4.18 displays the Wigner-Ville distribution of two atoms obtained with a Gaussian window g. The oscillating interference appears at the middle time-frequency point.
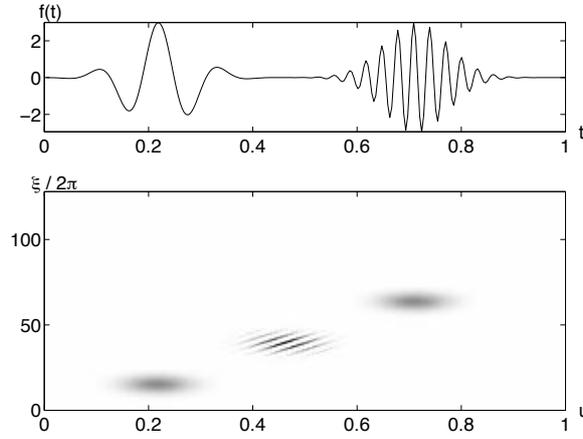

Figure 4.18: Wigner-Ville distribution $P_V f(u, \xi)$ of two Gabor atoms shown at the top. The oscillating interferences are centered at the middle time-frequency location.

This example shows that the interference $I[f_1, f_2](u, \xi)$ has some energy in regions where $|f(u)|^2 \approx 0$ and $|\hat{f}(\xi)|^2 \approx 0$. These interferences can have a complicated structure [25, 301] but they are necessarily oscillatory because the marginal integrals (4.135) and (4.136) vanish:

$$\int_{-\infty}^{+\infty} P_V f(u, \xi) d\xi = 2\pi |f(u)|^2 \quad , \quad \int_{-\infty}^{+\infty} P_V f(u, \xi) du = |\hat{f}(\xi)|^2.$$

Analytic Part   Interference terms also exist in a real signal f with a single instantaneous frequency component. Let $f_a(t) = a(t) \exp[i\theta(t)]$ be its analytic part:

$$f = \text{Re}[f_a] = \frac{1}{2}(f_a + f_a^*).$$

Theorem 4.8 proves that for fixed u, $P_V f_a(u, \xi)$ and $P_V f_a^*(u, \xi)$ have an energy concentrated respectively in the neighborhood of $\xi_1 = \theta'(u)$ and $\xi_2 = -\theta'(u)$. Both components create an interference term at the intermediate zero frequency $\xi_0 = (\xi_1 + \xi_2)/2 = 0$. To avoid this low frequency interference, we often compute $P_V f_a$ as opposed to $P_V f$.

Figure 4.19 displays $P_V f_a$ for a real signal f that includes a linear chirp, a quadratic chirp and two isolated time-frequency atoms. The linear and quadratic chirps are localized along narrow time frequency lines, which are spread on wider bands by the scalogram and the scalogram shown in Figure 4.3 and 4.11. However, the interference terms create complex oscillatory patterns that make it difficult to detect the existence of the two time-frequency transients at $t = 0.5$ and $t = 0.87$, which clearly appear in the spectrogram and the scalogram.
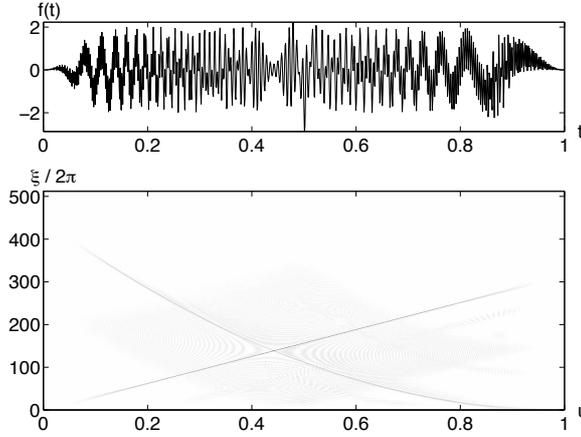
Figure 4.19: The bottom displays the Wigner-Ville distribution $P_V f_a(u, \xi)$ of the analytic part of the top signal.

Positivity   Since the interference terms include positive and negative oscillations, they can be partly removed by smoothing $P_V f$ with a kernel K:

$$P_K f(u, \xi) = \int_{-\infty}^{+\infty} \int_{-\infty}^{+\infty} P_V f(u', \xi') \, K(u, u', \xi, \xi') \, du' \, d\xi'. \tag{4.139}$$

The time-frequency resolution of this distribution depends on the spread of the kernel K in the neighborhood of (u, ξ). Since the interferences take negative values, one can guarantee that all interferences are removed by imposing that this time-frequency distribution remain positive $P_K f(u, \xi) \geqslant 0$ for all $(u, \xi) \in \mathbb{R}^2$.

The spectrogram (4.12) and scalogram (4.55) are examples of positive time-frequency energy distributions. In general, let us consider a family of time-frequency atoms $\{\phi_\gamma\}_{\gamma \in \Gamma}$. Suppose that for any (u, ξ) there exists a unique atom $\phi_{\gamma(u, \xi)}$ centered in time-frequency at (u, ξ). The resulting time-frequency energy density is

$$Pf(u, \xi) = |\langle f, \phi_{\gamma(u, \xi)} \rangle|^2.$$

The Moyal formula (4.134) proves that this energy density can be written as a time-frequency averaging of the Wigner-Ville distribution

$$Pf(u, \xi) = \frac{1}{2\pi} \int \int P_V f(u', \xi') \, P_V \phi_{\gamma(u, \xi)}(u', \xi') \, du' \, d\xi'. \tag{4.140}$$

The smoothing kernel is the Wigner-Ville distribution of the atoms

$$K(u, u', \xi, \xi') = \frac{1}{2\pi} \, P_V \phi_{\gamma(u, \xi)}(u', \xi').$$

The loss of time-frequency resolution depends on the spread of the distribution $P_V \phi_{\gamma(u, \xi)}(u', \xi')$ in the neighborhood of (u, v).

Example 4.20. *A spectrogram is computed with windowed Fourier atoms*

$$\phi_{\gamma(u, \xi)}(t) = g(t - u) \, e^{i\xi t}.$$

The Wigner-Ville distribution calculated in (4.130) yields

$$K(u, u', \xi, \xi') = \frac{1}{2\pi} \, P_V \phi_{\gamma(u, \xi)}(u', \xi') = \frac{1}{2\pi} \, P_V g(u' - u, \xi' - \xi). \tag{4.141}$$

For a spectrogram, the Wigner-Ville averaging (4.140) is therefore a two-dimensional convolution with $P_V g$. If g is a Gaussian window, then $P_V g$ is a two-dimensional Gaussian. This proves that

averaging $P_V f$ with a sufficiently wide Gaussian defines a positive energy density. The general class of time-frequency distributions obtained by convolving $P_V f$ with an arbitrary kernel K is studied in Section 4.5.3.

Example 4.21. *Let ψ be an analytic wavelet whose center frequency is η. The wavelet atom $\psi_{u,s}(t) = s^{-1/2}\psi((t-u)/s)$ is centered at $(u, \xi = \eta/s)$ and the scalogram is defined by*

$$P_W f(u,\xi) = |\langle f, \psi_{u,s}\rangle|^2 \quad \text{for } \xi = \eta/s.$$

Properties (4.127,4.129) prove that the averaging kernel is

$$K(u,u',\xi,\xi') = \frac{1}{2\pi}\, P_V \psi\left(\frac{u'-u}{s}, s\xi'\right) = \frac{1}{2\pi}\, P_V \psi\left(\frac{\xi}{\eta}(u'-u), \frac{\eta}{\xi}\xi'\right).$$

Positive time-frequency distributions totally remove the interference terms but produce a loss of resolution. This is emphasized by the following theorem, due to Wigner [485].

Theorem 4.11 (Wigner). *There is no positive quadratic energy distribution Pf that satisfies*

$$\int_{-\infty}^{+\infty} Pf(u,\xi)\, d\xi = 2\pi\, |f(u)|^2 \quad and \quad \int_{-\infty}^{+\infty} Pf(u,\xi)\, du = |\hat{f}(\xi)|^2. \tag{4.142}$$

Proof. Suppose that Pf is a positive quadratic distribution that satisfies these marginals. Since $Pf(u,\xi) \geqslant 0$, the integrals (4.142) imply that if the support of f is included in an interval I then $Pf(u,\xi) = 0$ for $u \notin I$. We can associate to the quadratic form Pf a bilinear distribution defined for any f and g by

$$P[f,g] = \frac{1}{4}\Big(P(f+g) - P(f-g)\Big).$$

Let f_1 and f_2 be two non-zero signals whose supports are two intervals I_1 and I_2 that do not intersect, so that $f_1 f_2 = 0$. Let $f = a\, f_1 + b\, f_2$:

$$Pf = |a|^2\, Pf_1 + ab^*\, P[f_1, f_2] + a^* b\, P[f_2, f_1] + |b|^2\, Pf_2.$$

Since I_1 does not intersect I_2, $Pf_1(u,\xi) = 0$ for $u \in I_2$. Remember that $Pf(u,\xi) \geqslant 0$ for all a and b so necessarily $P[f_1, f_2](u,\xi) = P[f_2, f_1](u,\xi) = 0$ for $u \in I_2$. Similarly we prove that these cross terms are zero for $u \in I_1$ and hence

$$Pf(u,\xi) = |a|^2\, Pf_1(u,\xi) + |b|^2\, Pf_2(u,\xi).$$

Integrating this equation and inserting (4.142) yields

$$|\hat{f}(\xi)|^2 = |a|^2\, |\hat{f}_1(\xi)|^2 + |b|^2\, |\hat{f}_2(\xi)|^2.$$

Since $\hat{f}(\xi) = a\, \hat{f}_1(\xi) + b\, \hat{f}_2(\xi)$ it follows that $\hat{f}_1(\xi)\, \hat{f}_2^*(\xi) = 0$. But this is not possible because f_1 and f_2 have a compact support in time and Theorem 2.7 proves that $\hat{f}_1$ and $\hat{f}_2$ are $\mathbf{C}^\infty$ functions that cannot vanish on a whole interval. We thus conclude that one cannot construct a positive quadratic distribution Pf that satisfies the marginals (4.142). ∎ ∎

4.5.3 Cohen's Class

While attenuating the interference terms with a smoothing kernel K, we may want to retain certain important properties. Cohen [176] introduced a general class of quadratic time-frequency distributions that satisfy the time translation and frequency modulation invariance properties (4.127) and (4.128). If a signal is translated in time or frequency, its energy distribution is just translated by the corresponding amount. This was the beginning of a systematic study of quadratic time-frequency distributions obtained as a weighted average of a Wigner-Ville distribution [6, 25, 177, 300].

Section 2.1 proves that linear translation invariant operators are convolution products. The translation invariance properties (4.127,4.128) are thus equivalent to imposing that the smoothing kernel in (4.139) be a convolution kernel

$$K(u,u',\xi,\xi') = K(u-u',\xi-\xi'), \tag{4.143}$$

and hence

$$P_K f(u, \xi) = P_V f \star K(u, \xi) = \iint K(u - u', \xi - \xi') \, P_V f(u', \xi') \, du' \, d\xi'. \qquad (4.144)$$

The spectrogram is an example of Cohen's class distribution, whose kernel in (4.141) is the Wigner-Ville distribution of the window

$$K(u, \xi) = \frac{1}{2\pi} P_V g(u, \xi) = \frac{1}{2\pi} \int_{-\infty}^{+\infty} g\left(u + \frac{\tau}{2}\right) g\left(u - \frac{\tau}{2}\right) e^{-i\tau\xi} \, d\tau. \qquad (4.145)$$

Ambiguity Function   The properties of the convolution (4.144) are more easily studied by calculating the two-dimensional Fourier transform of $P_V f(u, \xi)$ with respect to u and ξ. We denote by $Af(\tau, \gamma)$ this Fourier transform

$$Af(\tau, \gamma) = \int_{-\infty}^{+\infty} \int_{-\infty}^{+\infty} P_V f(u, \xi) \, \exp[-i(u\gamma + \xi\tau)] \, du \, d\xi.$$

Note that the Fourier variables τ and γ are inverted with respect to the usual Fourier notation. Since the one-dimensional Fourier transform of $P_V f(u, \xi)$ with respect to u is $\hat{f}(\xi + \gamma/2) \, \hat{f}^*(\xi - \gamma/2)$, applying the one-dimensional Fourier transform with respect to ξ gives

$$Af(\tau, \gamma) = \int_{-\infty}^{+\infty} \hat{f}\left(\xi + \frac{\gamma}{2}\right) \hat{f}^*\left(\xi - \frac{\gamma}{2}\right) e^{-i\tau\xi} \, d\xi. \qquad (4.146)$$

The Parseval formula yields

$$Af(\tau, \gamma) = \int_{-\infty}^{+\infty} f\left(u + \frac{\tau}{2}\right) f^*\left(u - \frac{\tau}{2}\right) e^{-i\gamma u} \, du. \qquad (4.147)$$

We recognize the *ambiguity function* encountered in (4.24) when studying the time-frequency resolution of a windowed Fourier transform. It measures the energy concentration of f in time and in frequency.

Kernel Properties   The Fourier transform of $K(u, \xi)$ is

$$\hat{K}(\tau, \gamma) = \int_{-\infty}^{+\infty} \int_{-\infty}^{+\infty} K(u, \xi) \, \exp[-i(u\gamma + \xi\tau)] \, du \, d\xi.$$

As in the definition of the ambiguity function (4.146), the Fourier parameters τ and γ of $\hat{K}$ are inverted. The following theorem gives necessary and sufficient conditions to ensure that P_K satisfies marginal energy properties like those of the Wigner-Ville distribution. The Wigner Theorem 4.11 proves that in this case $P_K f(u, \xi)$ takes negative values.

Theorem 4.12. *For all $f \in \mathbf{L}^2(\mathbb{R})$*

$$\int_{-\infty}^{+\infty} P_K f(u, \xi) \, d\xi = 2\pi \, |f(u)|^2 \quad , \quad \int_{-\infty}^{+\infty} P_K f(u, \xi) \, du = |\hat{f}(\xi)|^2, \qquad (4.148)$$

if and only if

$$\forall (\tau, \gamma) \in \mathbb{R}^2 \quad , \quad \hat{K}(\tau, 0) = \hat{K}(0, \gamma) = 1. \qquad (4.149)$$

Proof. Let $A_K f(\tau, \gamma)$ be the two-dimensional Fourier transform of $P_K f(u, \xi)$. The Fourier integral at $(0, \gamma)$ gives

$$\int_{-\infty}^{+\infty} \int_{-\infty}^{+\infty} P_K f(u, \xi) \, e^{-iu\gamma} \, d\xi \, du = A_K f(0, \gamma). \qquad (4.150)$$

Since the ambiguity function $Af(\tau, \gamma)$ is the Fourier transform of $P_V f(u, \xi)$, the two-dimensional convolution (4.144) gives

$$A_K(\tau, \gamma) = Af(\tau, \gamma) \, \hat{K}(\tau, \gamma). \qquad (4.151)$$

The Fourier transform of $2\pi|f(u)|^2$ is $\hat{f} \star \overline{\hat{f}}(\gamma)$, with $\overline{\hat{f}}(\gamma) = \hat{f}^*(-\gamma)$. The relation (4.150) shows that (4.148) is satisfied if and only if

$$A_K f(0,\gamma) = A f(0,\gamma)\,\hat{K}(0,\gamma) = \hat{f} \star \overline{\hat{f}}(\gamma). \qquad (4.152)$$

Since $P_V f$ satisfies the marginal property (4.135), we similarly prove that

$$A f(0,\gamma) = \hat{f} \star \overline{\hat{f}}(\gamma).$$

Requiring that (4.152) be valid for any $\hat{f}(\gamma)$, is equivalent to requiring that $\hat{K}(0,\gamma) = 1$ for all $\gamma \in \mathbb{R}$. The same derivation applied to the other marginal integration yields $\hat{K}(\xi, 0) = 1$. ∎   ∎

In addition to requiring time-frequency translation invariance, it may be useful to guarantee that P_K satisfies the same scaling property as a Wigner-Ville distribution:

$$g(t) = \frac{1}{\sqrt{s}} f\left(\frac{t}{s}\right) \implies P_K g(u,\xi) = P_K f\left(\frac{u}{s}, s\xi\right).$$

Such a distribution P_K is *affine* invariant. One can verify that affine invariance is equivalent to imposing that

$$\forall s \in \mathbb{R}^+ \quad, \quad K\left(s\,u, \frac{\xi}{s}\right) = K(u,\xi), \qquad (4.153)$$
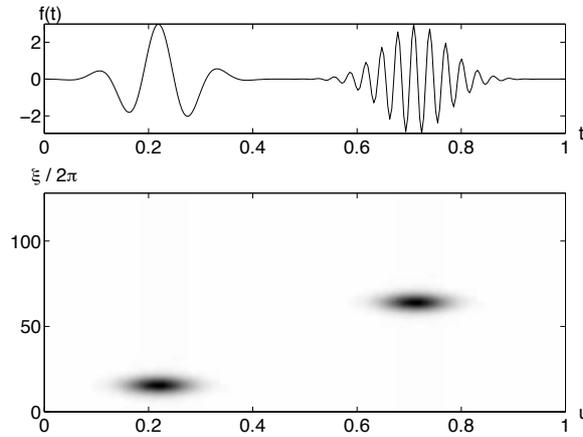
and hence

$$K(u,\xi) = K(u\,\xi, 1) = \beta(u\,\xi).$$

Figure 4.20: Choi-William distribution $P_K f(u,\xi)$ of the two Gabor atoms shown at the top. The interference term that appears in the Wigner-Ville distribution of Figure 4.18 has nearly disappeared.

Example 4.22. *The Rihaczek distribution is an affine invariant distribution whose convolution kernel is*

$$\hat{K}(\tau,\gamma) = \exp\left(\frac{i\,\tau\,\gamma}{2}\right). \qquad (4.154)$$

A direct calculation shows that

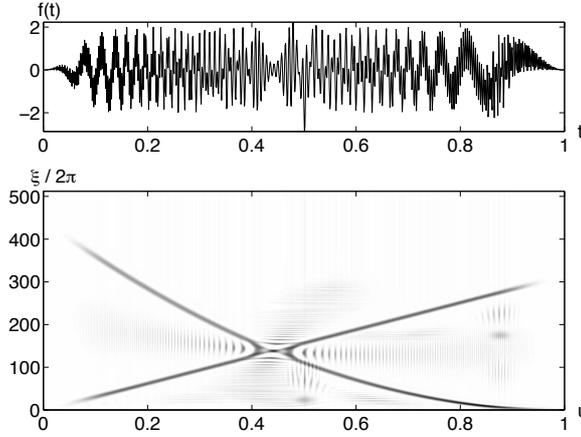$$P_K f(u,\xi) = f(u)\,\hat{f}^*(\xi)\,\exp(-iu\xi). \qquad (4.155)$$

Figure 4.21: Choi-William distribution $P_K f_a(u, \xi)$ of the analytic part of the signal shown at the top. The interferences remain visible.

Example 4.23. *The kernel of the Choi-William distribution is [160]*

$$\hat{K}(\tau, \gamma) = \exp(-\sigma^2 \, \tau^2 \, \gamma^2) \,. \tag{4.156}$$

It is symmetric and thus corresponds to a real function $K(u, \xi)$. This distribution satisfies the marginal conditions (4.149). Since $\lim_{\sigma \to 0} \hat{K}(\tau, \gamma) = 1$, when σ is small the Choi-William distribution is close to a Wigner-Ville distribution. Increasing σ attenuates the interference terms, but spreads $K(u, \xi)$, which reduces the time-frequency resolution of the distribution.

Figure 4.20 shows that the interference terms of two modulated Gaussians nearly disappear when the Wigner-Ville distribution of Figure 4.18 is averaged by a Choi-William kernel having a sufficiently large σ. Figure 4.21 gives the Choi-William distribution of the analytic signal whose Wigner-Ville distribution is in Figure 4.19. The energy of the linear and quadratic chirps are spread over wider time-frequency bands but the interference terms are attenuated, although not totally removed. It remains difficult to isolate the two modulated Gaussians at $t = 0.5$ and $t = 0.87$, which clearly appear in the spectrogram of Figure 4.3.

4.5.4 Discrete Wigner-Ville Computations

The Wigner integral (4.120) is the Fourier transform of $f(u + \tau/2)f^*(u - \tau/2)$:

$$P_V f(u, \xi) = \int_{-\infty}^{+\infty} f\left(u + \frac{\tau}{2}\right) f^*\left(u - \frac{\tau}{2}\right) e^{-i\tau\xi} \, d\tau. \tag{4.157}$$

For a discrete signal $f[n]$ defined over $0 \leqslant n < N$, the integral is replaced by a discrete sum:

$$P_V f[n, k] = \sum_{p=-N}^{N-1} f\left[n + \frac{p}{2}\right] f^*\left[n - \frac{p}{2}\right] \exp\left(\frac{-i2\pi kp}{N}\right). \tag{4.158}$$

When p is odd, this calculation requires knowing the value of f at half integers. These values are computed by interpolating f, with an addition of zeroes to its Fourier transform. This is necessary to avoid the aliasing produced by the discretization of the Wigner-Ville integral [164].

The interpolation $\tilde{f}$ of f is a signal of size $2N$ whose discrete Fourier transform $\widehat{\tilde{f}}$ is defined from the discrete Fourier transform $\hat{f}$ of f by

$$\widehat{\tilde{f}}[k] = \begin{cases} 2\hat{f}[k] & \text{if } 0 \leqslant k < N/2 \\ 0 & \text{if } N/2 < k < 3N/2 \\ 2\hat{f}[k - N] & \text{if } 3N/2 < k < 2N \\ \hat{f}[N/2] & \text{if } k = N/2, 3N/2 \end{cases} \,.$$

Computing the inverse discrete Fourier transform shows that $\tilde{f}[2n] = f[n]$ for $n \in [0, N-1]$. When $n \notin [0, 2N-1]$, we set $\tilde{f}[n] = 0$. The Wigner summation (4.158) is calculated from $\tilde{f}$:

$$
\begin{aligned}
P_V f[n, k] &= \sum_{p=-N}^{N-1} \tilde{f}[2n+p]\, \tilde{f}^*[2n-p] \, \exp\left(\frac{-i2\pi kp}{N}\right) \\
&= \sum_{p=0}^{2N-1} \tilde{f}[2n+p-N]\, \tilde{f}^*[2n-p+N]\, \exp\left(\frac{-i2\pi(2k)p}{2N}\right).
\end{aligned}
$$

For $0 \leqslant n < N$ fixed, $P_V f[n, k]$ is the discrete Fourier transform of size $2N$ of $g[p] = \tilde{f}[2n+p-N]\,\tilde{f}^*[2n-p+N]$ at the frequency $2k$. The discrete Wigner-Ville distribution is thus calculated with N FFT procedures of size $2N$, which requires $O(N^2 \log N)$ operations. To compute the Wigner-Ville distribution of the analytic part f_a of f, we use (4.48).

Cohen's Class A Cohen's class distribution is calculated with a circular convolution of the discrete Wigner-Ville distribution with a kernel $K[p, q]$:

$$
P_K[n, k] = P_V \circledast K[n, k]. \tag{4.159}
$$

Its two-dimensional discrete Fourier transform is therefore

$$
A_K[p, q] = Af[p, q]\, \hat{K}[p, q]. \tag{4.160}
$$

The signal $Af[p, q]$ is the discrete ambiguity function, calculated with a two-dimensional FFT of the discrete Wigner-Ville distribution $P_V f[n, k]$. As in the case of continuous time, we have inverted the index p and q of the usual two-dimensional Fourier transform. The Cohen's class distribution (4.159) is obtained by calculating the inverse Fourier transform of (4.160). This also requires a total of $O(N^2 \log N)$ operations.

4.6 Exercises

4.1. [2] *Instantaneous frequency* Let $f(t) = \exp[i\phi(t)]$.
   (a) Prove that $\int_{-\infty}^{+\infty} |Sf(u, \xi)|^2 \, d\xi = 2\pi$. Hint: $Sf(u, \xi)$ is a Fourier transform; use the Parseval formula.
   (b) Similarly, show that

   $$
   \int_{-\infty}^{+\infty} \xi \, |Sf(u, \xi)|^2 \, d\xi = 2\pi \int_{-\infty}^{+\infty} \phi'(t)\, |g(t-u)|^2 \, dt,
   $$

   and interpret this result.

4.2. [1] When $g(t) = (\pi\sigma^2)^{-1/4} \exp(-t^2/(2\sigma^2))$, compute the ambiguity function $Ag(\tau, \gamma)$.

4.3. [1] Prove that the approximate reconstruction formula (4.66) is exact if and only if

$$
\frac{\log_e a}{C_\psi} \sum_{j=I}^{J} \frac{1}{a^j}\, \hat{\psi}_j[k] + \frac{1}{C_\psi\, a^J}\, \hat{\phi}_J[k] = 1 \ .
$$

Compute numerically the left equation value for different a when $\psi_j[n]$ is constructed from the Gabor wavelet (4.60) and (4.62) with $\sigma = 1$ and $2^J < N/4$.

4.4. [1] Let $g[n]$ be a window with L non-zero coefficients. For signals of size N, describe a fast algorithm that computes the discrete windowed Fourier transform (4.27) with $O(N \log_2 L)$ operations.

4.5. [3] Let K be the reproducing kernel (4.21) of a windowed Fourier transform: $K(u_0, u, \xi_0, \xi) = \langle g_{u,\xi}, g_{u_0,\xi_0} \rangle$.
   (a) For any $\Phi \in \mathbf{L}^2(\mathbb{R}^2)$ we define:

   $$
   T\Phi(u_0, \xi_0) = \frac{1}{2\pi} \int_{-\infty}^{+\infty} \int_{-\infty}^{+\infty} \Phi(u, \xi)\, K(u_0, u, \xi_0, \xi) \, du \, d\xi.
   $$

   Prove that T is an orthogonal projector on the space of functions $\Phi(u, \xi)$ that are windowed Fourier transforms of functions in $\mathbf{L}^2(\mathbb{R})$.

(b) Suppose that for all $(u, \xi) \in \mathbb{R}^2$ we are given $\tilde{S}f(u, \xi) = Q\big(Sf(u, \xi)\big)$, which is a quantization of the windowed Fourier coefficients. How can we reduce the norm $\mathbf{L}^2(\mathbb{R}^2)$ of the quantification error $\varepsilon(u, \xi) = Sf(u, \xi) - Q\big(Sf(u, \xi)\big)$?

4.6. [3] Prove the wavelet reconstruction formula (4.45).

4.7. [3] Prove that a scaling function ϕ defined by (4.42) satisfies $\|\phi\| = 1$.

4.8. [2] Let ψ be a real and even wavelet such that $C = \int_0^{+\infty} \omega^{-1}\, \hat{\psi}(\omega)\, d\omega < +\infty$. Prove that

$$\forall f \in \mathbf{L}^2(\mathbb{R}) \quad , \quad f(t) = \frac{1}{C} \int_0^{+\infty} Wf(t, s)\, \frac{ds}{s^{3/2}} \ . \tag{4.161}$$

4.9. [3] *Analytic Continuation* Let $f \in \mathbf{L}^2(\mathbb{R})$ be a function such that $\hat{f}(\omega) = 0$ for $\omega < 0$. For any complex $z \in \mathbb{C}$ such that $\mathrm{Im}(z) \geqslant 0$, we define

$$f^{(p)}(z) = \frac{1}{\pi} \int_0^{+\infty} (i\omega)^p\, \hat{f}(\omega)\, \mathrm{e}^{iz\omega}\, d\omega \ .$$

(a) Verify that if f is $\mathbf{C}^p$ then $f^{(p)}(t)$ is the derivative of order p of $f(t)$.
(b) Prove that if $\mathrm{Im}(z) > 0$, then $f^{(p)}(z)$ is differentiable relatively to the complex variable z. Such a function is said to be *analytic* on the upper half complex plane.
(c) Prove that this analytic extension can be written as a wavelet transform

$$f^{(p)}(x + iy) = y^{-p-1/2}\, Wf(x, y) \ ,$$

calculated with an analytic wavelet ψ that you will specify.

4.10. [2] Let $f(t) = \cos(a \cos bt)$. We want to compute precisely the instantaneous frequency of f from the ridges of its windowed Fourier transform. Find a necessary condition on the window support as a function of a and b. If $f(t) = \cos(a \cos bt) + \cos(a \cos bt + ct)$, find a condition on a, b and c in order to measure both instantaneous frequencies with the ridges of a windowed Fourier transform. Verify your calculations with a numerical implementation.

4.11. [4] *Noise removal* We want to suppress noise from audio signals by thresholding ridge coefficients. Implement a dual synthesis algorithm that reconstruct audio signal approximations from windowed Fourier ridge points (4.96) or wavelet ridge points (4.116), with the conjugate gradient inverse frame algorithm of Theorem 5.8. Study the SNR of audio denoising by thresholding the ridge coefficients. Try to improve this SNR by averaging the spectrogram values along ridges. Compare the SNR with a linear filtering estimator.

4.12. [4] *Sound duration* Make a program that modifies the sound duration with the formula (4.72) by modifying the ridges of a window Fourier transform with (4.99) or of a wavelet transform with (4.118), and by reconstructing a signal with a dual synthesis.

4.13. [4] *Sound transposition* Implement a sound transposition, with windowed Fourier or wavelet ridges, with the transposition model (4.73). The resulting modifications of the ridge supports are specified by (4.100) and (4.119). The amplitude of the transposed harmonics can be computed with the auto-regressive model (4.75). A signal is restored with a dual synthesis algorithm.

4.14. [4] The sinusoidal model (4.71) is improved for speech signals by adding a non-harmonic component $B(t)$ to the partials [338]:

$$F(t) = \sum_{k=1}^{K} a_k(t) \cos \phi_k(t) + B(t). \tag{4.162}$$

Given a signal $f(t)$ that is considered to be a realization of $F(t)$, compute the ridges of a windowed Fourier transform, find the "main" partials and compute their amplitude a_k and phase ϕ_k. These partials are subtracted from the signal. Over intervals of fixed size, the residue is modeled as the realization of an autoregressive process $B(t)$, of order 10 to 15. Use a standard algorithm to compute the parameters of this autoregressive process [56]. Evaluate the audio quality of the sound restored from the calculated model (4.162). Study an application to audio compression by quantizing and coding the parameters of the model.

4.15. [2] Prove that $Pf(u, \xi) = \|f\|^{-2}\, |f(u)|^2\, |\hat{f}(\xi)|^2$ satisfies the marginal properties (4.135) and (4.136). Why can't we apply the Wigner Theorem 4.11?

4.16. [1] Let g_σ be a Gaussian of variance σ^2. Prove that $P_\theta f(u, \xi) = P_V f \star \theta(u, \xi)$ is a positive distribution if $\theta(u, \xi) = g_\sigma(u) g_\beta(\xi)$ with $\sigma \beta \geqslant 1/2$. Hint: consider a spectrogram calculated with a Gaussian window.

4.17. [3] Let $\{g_n(t)\}_{n \in \mathbb{N}}$ be an orthonormal basis of $\mathbf{L}^2(\mathbb{R})$. Prove that

$$\forall (t, \omega) \in \mathbb{R}^2 \ , \ \sum_{n=0}^{+\infty} P_V g_n(t, \omega) = 1 \ .$$

4.18. [2] Let $f_a(t) = a(t) \exp[i\phi(t)]$ be the analytic part of $f(t)$. Prove that

$$\int_{-\infty}^{+\infty} \left(\xi - \phi'(t)\right)^2 P_V f_a(t, \xi) \, d\xi = -\pi \, a^2(t) \, \frac{d^2 \log a(t)}{dt^2} .$$

4.19. [4] To avoid the time-frequency resolution limitations of a windowed Fourier transform, we want to adapt the window size to the signal content. Let $g(t)$ be a window supported in $[-\frac{1}{2}, \frac{1}{2}]$. We denote by $S_j f(u, \xi)$ the windowed Fourier transform calculated with the dilated window $g_j(t) = 2^{-j/2} g(2^{-j} t)$. Find a procedure that computes a single map of ridges by choosing a "best" window size at each (u, ξ). One approach is to choose the scale 2^l for each (u, ξ) such that $|S_l f(u, \xi)|^2 = \sup_j |S_j f(u, \xi)|^2$. Test your algorithm on the linear and hyperbolic chirp signals (4.103,4.107). Test it on the Tweet and Greasy signals in WAVELAB.