Chapter 1

On rainbows and spectra

"One can enjoy a rainbow without necessarily forgetting the forces that made it."

— *Mark Twain*

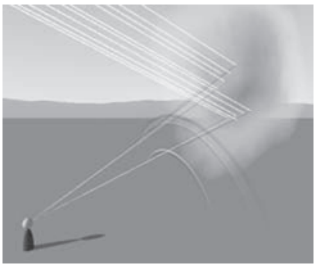

In the late thirteenth century, Theodoric of Freiberg, a Dominican monk, theologian, and physicist, performed a simple experiment: with his back to the sun, he held a spherical bottle filled with water in the sunlight. By following the trajectory of the refracted and reflected light and having the bottle play the same role as a single water drop, he gave a scientific explanation of rainbows, including secondary rainbows with weaker, reversed colors. His geometric analysis, described in his famous treatise *De iride* (*On the Rainbow*, *c.* 1310), was "perhaps the most dramatic development of fourteenth- and fifteenth-century optics" [60].

Theodoric of Freiberg fell short of a complete understanding of the rainbow phenomenon because, like many of his contemporaries, he believed that colors were simply intensities between black and white. A full understanding emerged three hundred years later when René Descartes and Isaac Newton explained that dispersion decomposes white light into spectral components of different wavelengths – the colors of the rainbow. In 1730, Newton, in his landmark book on optics [70], describes what is often called the *experimentum crucis* (crucial experiment) to prove that white light can be decomposed into constituent colors and then recombined into white light. This experiment is a physical implementation of decomposing light into its Fourier components – pure frequencies or colors of the rainbow, followed by a synthesis to recover the original; the cover photograph of the book depicts this experiment.[4]

This bit of history evokes two central themes of this book: geometric thinking

[4]This is a realization of Figure 7 in Part II of Newton's *First Book of Opticks*.

is a great tool in deducing explanations of phenomena; and decomposing an entity into its constituent components can be a key step in understanding its essential character, as well as an enabling tool in modifying these components prior to recombination. The rainbow's appearance is explained by the fact that sunlight contains a combination of all wavelengths within the visible range; separation of white light by wavelength, as with a prism, enables modifications prior to recombination. The collection of wavelengths is, as we will see, the spectrum.

A French physicist and mathematician, Joseph Fourier, formalized the notion of the spectrum in the early nineteenth century. He was interested in the heat equation – the differential equation governing the diffusion of heat. Fourier's key insight was to decompose a periodic function $x(t) = x(t+T)$ into an infinite sum of sines and cosines of periods T/k, $k \in \mathbb{Z}^+$. Since these sine and cosine components are eigenfunctions of the heat equation, the solution of the problem is simplified: one can analyze the differential equation for each component separately and combine the intermediate results, thanks to the linearity of the system. Fourier's decomposition earned him a coveted prize from the French Academy of Sciences, but with a mention that his work lacked rigor. Indeed, the question of which functions admit a Fourier decomposition is a deep one, and it took many years to settle. Fourier's work is one of the foundational blocks of signal processing and at the heart of the present book as well as its companion volume [57]. Fourier techniques have been joined in the past two decades by new tools such as wavelets, the other pillar we cover.

Signal representations The idea of a decomposition and a possible modification in the decomposed state leads to signal representations, where signals can be sequences (discrete domain) or functions (continuous domain). Similarly to what Fourier did, where he used sines and cosines for decomposition, we can imagine using other functions with particular properties. Call these basis vectors and denote them by φ_k, $k \in \mathbb{Z}$. Then

$$x \;=\; \sum_{k\in\mathbb{Z}} X_k \varphi_k \tag{1.1}$$

is called an expansion of x with respect to $\{\varphi_k\}_{k\in\mathbb{Z}}$, with $\{X_k\}$ the expansion coefficients.

Orthonormal bases When the basis vectors form an orthonormal set, the coefficients X_k are obtained from the function x and the basis vectors φ_k through an inner product

$$X_k \;=\; \langle x,\, \varphi_k\rangle. \tag{1.2}$$

For example, Fourier's construction of a series representation for periodic functions with period $T = 1$ can be written as

$$x(t) \;=\; \sum_{k\in\mathbb{Z}} X_k e^{j2\pi kt}, \tag{1.3a}$$

where

$$X_k \;=\; \int_0^1 x(t) e^{-j2\pi kt}\, dt. \tag{1.3b}$$

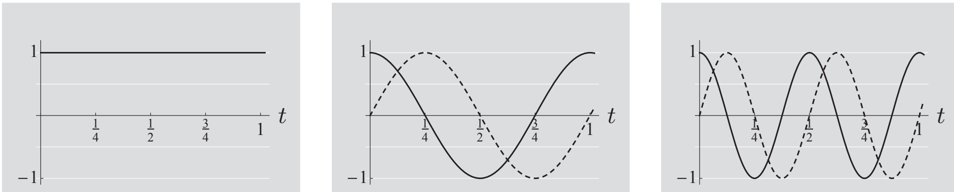

(a) $\varphi_0(t) = 1$. (b) $\varphi_1(t) = e^{j2\pi t}$. (c) $\varphi_2(t) = e^{j4\pi t}$.

Figure 1.1 Example Fourier series basis functions for the interval $[0, 1)$. Real parts are shown with solid lines and imaginary parts are shown with dashed lines.

We can define basis vectors φ_k, $k \in \mathbb{Z}$, on the interval $[0, 1)$, as

$$\varphi_k(t) \;=\; e^{j2\pi kt}, \qquad 0 \le t < 1, \tag{1.4}$$

and the Fourier series coefficients as

$$X_k \;=\; \langle x, \varphi_k \rangle \;=\; \int_0^1 x(t)\varphi_k^*(t)\,dt \;=\; \int_0^1 x(t)e^{-j2\pi kt}\,dt,$$

exactly the same as (1.3b). The basis vectors form an orthonormal set (the first few are shown in Figure 1.1):

$$\langle \varphi_k, \varphi_i \rangle \;=\; \int_0^1 e^{j2\pi kt}e^{-j2\pi it}\,dt \;=\; \begin{cases} 1, & \text{for } i = k; \\ 0, & \text{otherwise.} \end{cases} \tag{1.5}$$

While the Fourier series is certainly a key orthonormal basis with many outstanding properties, other bases exist, some of which have their own favorable properties. Early in the twentieth century, Alfred Haar proposed a basis which looks quite different from Fourier's. It is based on a function $\psi(t)$ defined as

$$\psi(t) \;=\; \begin{cases} 1, & \text{for } 0 \le t < \frac{1}{2}; \\ -1, & \text{for } \frac{1}{2} \le t < 1; \\ 0, & \text{otherwise.} \end{cases} \tag{1.6}$$

For the interval $[0, 1)$, we can build an orthonormal system by scaling $\psi(t)$ by powers of 2, and then shifting the scaled versions appropriately, yielding

$$\psi_{m,n}(t) \;=\; 2^{-m/2}\psi\left(\frac{t - n2^m}{2^m}\right), \tag{1.7}$$

with $m \in \{0, -1, -2, \ldots\}$ and $n \in \{0, 1, \ldots, 2^{-m} - 1\}$ (a few are shown in Figure 1.2). It is quite clear from the figure that the various basis functions are indeed orthogonal to each other, as they either do not overlap, or when they do, one changes sign over the constant span of the other. We will spend a considerable amount of time studying this system in the companion volume to this book, [57].

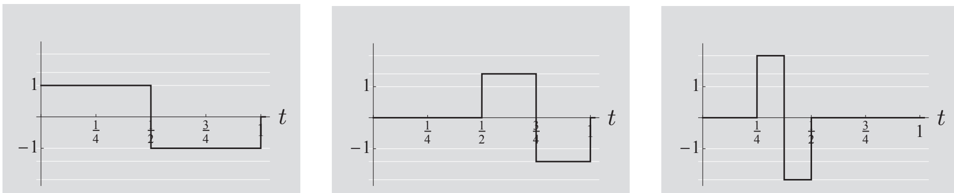

(a) $\psi_{0,0}(t)$. (b) $\psi_{-1,1}(t)$. (c) $\psi_{-2,1}(t)$.

Figure 1.2 Example Haar series basis functions for the interval $[0, 1)$. The prototype function is $\psi(t) = \psi_{0,0}(t)$.

While the system (1.7) is orthonormal, it cannot be a basis for all functions on $[0, 1)$; for example, there would be no way to reconstruct a constant 1. We remedy that by adding the function

$$\varphi_0(t) \;=\; \begin{cases} 1, & \text{for } 0 \leq t < 1; \\ 0, & \text{otherwise,} \end{cases} \tag{1.8}$$

into the mix, yielding an orthonormal basis for the interval $[0, 1)$. This is a very different basis from the Fourier one; for example, instead of being infinitely differentiable, no $\psi_{m,n}$ is even continuous. We can now define an expansion as in (1.3),

$$x(t) \;=\; \langle x, \varphi_0 \rangle \varphi_0(t) + \sum_{m=-\infty}^{0} \sum_{n=0}^{2^{-m}-1} X_{m,n} \psi_{m,n}(t), \tag{1.9a}$$

where

$$X_{m,n} \;=\; \int_0^1 x(t)\, \psi_{m,n}(t)\, dt. \tag{1.9b}$$

It is natural to ask which basis is better. Such a question does not have a simple answer, and the answer will depend on the class of functions or sequences we wish to represent, as well as our goals in the representation. Furthermore, we will have to be careful in describing what we mean by equality in an expansion such as (1.3a); otherwise we could be misled the same way Fourier was.

Approximation One way to assess the quality of a basis is to see how well it can approximate a given function with a finite number of terms. History is again enlightening. Fourier series became such a useful tool during the nineteenth century that researchers built elaborate mechanical devices to compute a function based on Fourier series coefficients. They built analog computers, based on harmonically related rotating wheels, where amplitudes of Fourier coefficients could be set and the sum computed. One such machine, the Harmonic Integrator, was designed by the physicists Albert Michelson and Samuel Stratton, and it could compute a series with 80 terms. To the designers' dismay, the synthesis of a square wave from its Fourier series led to oscillations around the discontinuity that would not go away

Cambridge University Press
978-1-107-03860-8 - Foundations of Signal Processing
Martin Vetterli, Jelena Kovačević and Vivek K Goyal
Excerpt
More information

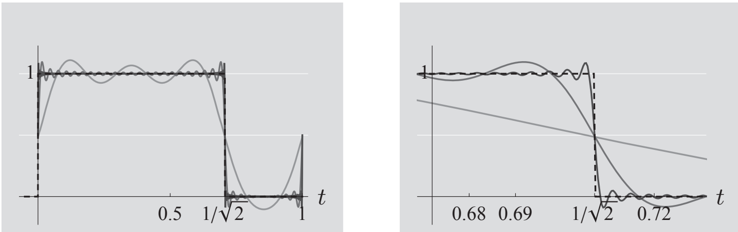

(a) Series with 9, 65, and 513 terms. (b) Detail of (a).

Figure 1.3 Approximations of a box function (dashed lines) with a Fourier series basis using 9, 65, and 513 terms (solid lines, from lightest to darkest). The plots illustrate the Gibbs phenomenon – oscillations that do not diminish in amplitude when approximating a discontinuous function with truncated Fourier series.

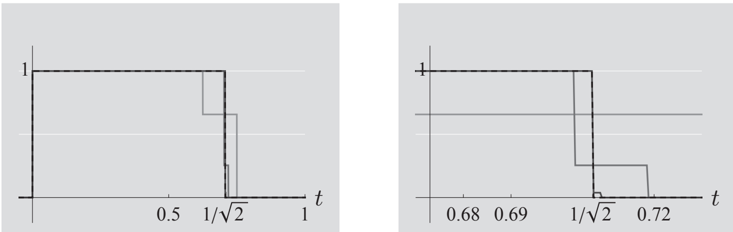

(a) Series with 8, 64, and 512 terms. (b) Detail of (a).

Figure 1.4 Approximation of a box function (dashed lines) with a Haar basis using the first 8 ($m = 0, -1, -2$), 64 ($m = 0, -1, \ldots, -5$), and 512 ($m = 0, -1, \ldots, -8$), terms (solid lines, from lightest to darkest), with $n = 0, 1, \ldots, 2^{-m} - 1$. The discontinuity is at the irrational point $1/\sqrt{2}$.

even as they increased the number of terms; they concluded that a mechanical problem was at fault. Not until 1899, when Josiah Gibbs proved that Fourier series of discontinuous functions cannot converge uniformly, was this myth dispelled. The phenomenon was termed the *Gibbs phenomenon*, referring to the oscillations appearing around the discontinuity when using any finite number of terms. Figure 1.3 shows approximations of a box function with a Fourier series basis (1.3a) using X_k, $k = -K, -K+1, \ldots, K-1, K$.

So what would the Haar basis provide in this case? Surely, it seems more appropriate for a box function. Unfortunately, taking the first 2^{-m} terms in the natural ordering (the term corresponding to the function $\varphi_0(t)$ plus 2^{-m} terms corresponding to each scale $m = 0, -1, -2, \ldots$) leads to a similarly poor performance, shown in Figure 1.4. This poor performance is dependent on the position of the discontinuity; approximating a box function with a discontinuity at an integer multiple of 2^{-k} for some $k \in \mathbb{Z}$ would lead to a much better performance.

However, changing the approximation procedure slightly makes a big differ-

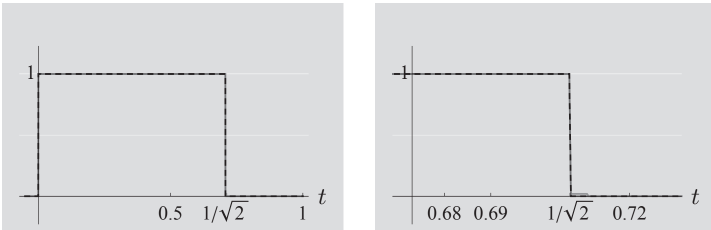

(a) Series with 8 and 15 terms. (b) Detail of (a).

Figure 1.5 Approximation of a box function (dashed lines) with a Haar basis using the 8 (light) and 15 (dark) largest-magnitude terms. The 15-term approximation is visually indistinguishable from the target function.

ence. Upon retaining the largest coefficients in absolute value instead of simply keeping a fixed set of terms, the approximation quality changes drastically, as seen in Figure 1.5. In this admittedly extreme example, for each m, there is only one n such that $X_{m,n}$ is nonzero (that for which the corresponding Haar wavelet straddles the discontinuity). Thus, approximating using coefficients largest in absolute value allows many more values of m to be included.

Through this comparison, we have illustrated how the quality of a basis for approximation can depend on the method of approximation. Retaining a predefined set of terms, as in the Fourier example case (Figure 1.3) or the first Haar example (Figure 1.4) is called linear approximation. Retaining an adaptive set of terms instead, as in the second Haar example (Figure 1.5), is called nonlinear approximation and leads to a superior approximation quality.

Overview of the book The purpose of this book is to develop the framework for the methods just described, namely expansions and approximations, as well as to show practical examples where these methods are used in engineering and applied sciences. In particular, we will see that expansions and approximations are closely related to the essential signal processing tasks of sampling, filtering, estimation, and compression.

Chapter 2, From Euclid to Hilbert, introduces the basic machinery of Hilbert spaces. These are vector spaces endowed with operations that induce intuitive geometric properties. In this general setting, we develop the notion of signal representations, which are essentially coordinate systems for the vector space. When a representation is complete and not redundant, it provides a *basis* for the space; when it is complete and redundant, it provides a *frame* for the space. A key virtue for a basis is orthonormality; its counterpart for a frame is tightness.

Chapters 3 and 4 focus our attention on sequence and function spaces for which the domain can be associated with *time*, leading to an inherent ordering not necessarily present in a general Hilbert space. In **Chapter 3, Sequences and discrete-time systems**, a vector is a sequence that depends on *discrete time*, and an important class of linear operators on these vectors is those that are invariant to time shifts; these are convolution operators. These operators lead naturally to signal representations using the discrete-time Fourier transform and, for circularly extended finite-length sequences, the discrete Fourier transform.

Chapter 4, Functions and continuous-time systems, parallels Chapter 3; a vector is now a function that depends on *continuous time*, and an important class of linear operators on these vectors are again those that are invariant to time shifts; these are convolution operators. These operators lead naturally to signal representations using the Fourier transform and, for circularly extended finite-length functions, or periodic functions, the Fourier series. The four Fourier representations from these two chapters exemplify the diagonalization of linear, shift-invariant operators, or convolutions, in the various domains.

Chapter 5, Sampling and interpolation, makes fundamental connections between Chapters 3 and 4. Associating a discrete-time sequence with a given continuous-time function is *sampling*, and the converse is *interpolation*; these are central concepts in signal processing since digital computations on continuous-domain phenomena must be performed in a discrete domain.

Chapter 6, Approximation and compression, introduces many types of approximations that are central to making computationally practical tools. Approximation by polynomials and by truncations of series expansions are studied, along with the basic principles of compression.

Chapter 7, Localization and uncertainty, introduces time, frequency, scale, and resolution properties of individual vectors; these properties build our intuition for what might or might not be captured by a single representation coefficient. We then study these properties for sets of vectors used to represent signals. In particular, time and frequency localization lead to the concept of a time–frequency plane, where essential differences between Fourier techniques and wavelet techniques become evident: Fourier techniques use vectors with equal spacing in frequency while wavelet techniques use vectors with power-law spacing in frequency; furthermore, Fourier techniques use vectors at equal scale while wavelet techniques use geometrically spaced scales. We end with examples with real signals to develop intuition about various signal representations.

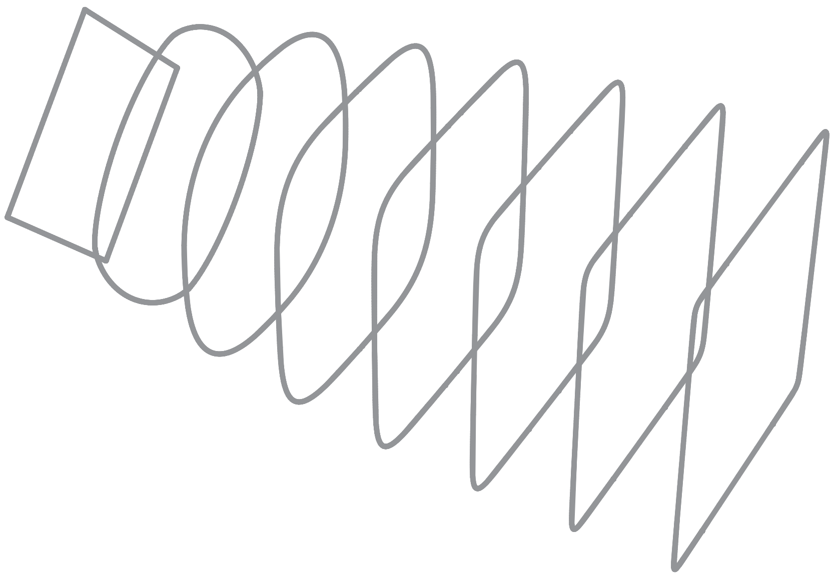

Chapter 2

From Euclid to Hilbert

> "Mathematics is the art of giving the same name to different things."
>
> — *Henri Poincaré*

Contents

We start our journey into signal processing with different backgrounds and perspectives. This chapter aims to establish a common language, develop the foundations for our study, and begin to draw out key themes.

There will be more formal definitions in this chapter than in any other, to approach the ideal of a self-contained treatment. However, we must assume some background in common: On the one hand, we expect the reader to be familiar with linear algebra at the level of [93, Ch. 1–5] (see also Appendix 2.B) and probability at the level of [6, Ch. 1–4] (see also Appendix 2.C). (The textbooks we have cited are just examples; nothing unique to those books is necessary.) On the other hand, we are not assuming prior knowledge of general vector space abstractions or mathematical analysis beyond basic calculus; we develop these topics here to extend geometric intuition from ordinary Euclidean space to spaces of sequences and functions. For more details on abstract vector spaces, we recommend books by Kreyszig [59], Luenberger [64], and Young [111].

2.1 Introduction

This section introduces many topics of the chapter through the familiar setting of the real plane. In the more general treatment of subsequent sections, the intuition we have developed through years of dealing with the Euclidean spaces around us ($\mathbb{R}^2$ and $\mathbb{R}^3$) will generalize to some not-so-familiar spaces. Readers comfortable with vector spaces, inner products, norms, projections, and bases may skip this section; otherwise, this will be a gentle introduction to Euclid's world.

Real plane as a vector space

Let us start with a look at the familiar setting of $\mathbb{R}^2$, that is, real vectors with two coordinates. We adopt the convention of vectors being columns and often write them compactly as transposes of rows, such as $x = \begin{bmatrix} x_0 & x_1 \end{bmatrix}^\top$. The first entry is the horizontal component and the second entry is the vertical component.

Adding two vectors in the plane produces a third one also in the plane; multiplying a vector by a real scalar produces a second vector also in the plane. These two ingrained facts make the real plane be a *vector space*.

Inner product and norm

The *inner product* of vectors $x = \begin{bmatrix} x_0 & x_1 \end{bmatrix}^\top$ and $y = \begin{bmatrix} y_0 & y_1 \end{bmatrix}^\top$ in the real plane is

$$\langle x, y \rangle \;=\; x_0 y_0 + x_1 y_1. \tag{2.1}$$

Other names for inner product are *scalar product* and *dot product*. The inner product of a vector with itself is simply

$$\langle x, x \rangle \;=\; x_0^2 + x_1^2,$$

a nonnegative quantity that is zero when $x_0 = x_1 = 0$. The *norm* of a vector x is

$$\|x\| \;=\; \sqrt{\langle x, x \rangle} \;=\; \sqrt{x_0^2 + x_1^2}. \tag{2.2}$$

While the norm is sometimes called the *length*, we avoid this usage because length can also refer to the number of components in a vector. A vector of norm 1 is called a *unit vector*.

In (2.1), the inner product computation depends on the choice of coordinate axes. Let us now derive an expression in which the coordinates disappear. Consider