ENGR 76: Information Science and Engineering — Winter 2026 — Lectures 5–8

Frequency Domain Representation of Signals

Siddharth Chandak

The idea of representing signals using sines and cosines dates back to the early 19th century, most notably to the work of Joseph Fourier, who introduced these representations while studying heat flow. Since then, Fourier-based representations have become a central tool across science and engineering. In this course, we will use it in two broad contexts: representing and compressing images, and in communication systems. Beyond these, frequency-domain representations also appear throughout quantum mechanics, control systems, audio and speech processing, medical imaging, and data analysis. While the mathematical theory of Fourier analysis is deep and extensive, our focus will be on the representational viewpoint and the intuition needed to use these ideas effectively in practical settings.

1 Signals and Representation

Up to this point in the course, we have primarily viewed an information source as a sequence of random variables,

$$X_1, X_2, X_3, \ldots$$

This viewpoint is natural when studying entropy, coding, and compression. However, many real-world sources are more naturally viewed as *signals*.

Definition 1 (Signal)**.** *A signal is a function that carries information.*

A signal may be a function of time, such as an electrical voltage or an audio waveform, or a function of space, such as an image where pixel intensities vary across a grid. For the first part of this module, we will think of signals as a function of time and use the notation $x(t)$.

Why representation matters? A signal can be described in many different ways, and each of these methods of describing are different representations of the signal. For example, one can describe a signal by listing its values at each time instant. Alternatively, one can describe the same signal by specifying how it is composed from simpler known building blocks. In images, the first representation corresponds to one where the image is described by listing the intensity at each pixel.

A central idea of this module is that: The usefulness of a representation depends on what we want to do with the signal. Some representations make certain properties of a signal immediately visible, while others hide them. Our goal is to develop a representation that reveals important structure in signals.

In this module, we focus on a particular class of representations known as *frequency domain representations*. Instead of describing a signal by listing its values over time, a frequency domain representation describes a signal in terms of the frequencies that make it up and how strongly each frequency contributes. This information is often summarized through what is called the *spectrum* of a signal, which indicates the presence and strength of different frequency components. Viewing a signal through its spectrum can make global trends, and fine-scale variations much more apparent than in the time domain. In the rest of this module, we develop this frequency-domain viewpoint systematically and study how it helps us understand approximation, filtering, and, ultimately, compression.

2 Periodic Functions and Sines and Cosines

We begin by restricting our attention to a special but extremely important class of signals: *periodic functions*. This restriction is not as limiting as it may seem, and it will allow us to develop the key ideas of frequency-domain representation in a clean and intuitive way.

2.1 Periodic functions

Definition 2 (Periodic function)**.** *A signal $x(t)$ is periodic if there exists constant $T > 0$ such that $x(t+T) = x(t)$ for all t. The* smallest *$T > 0$ for which the signal satisfies this property is called the* ***period*** *of the signal.*

In words, a periodic function repeats itself exactly every T units of time. The smallest such T is called the *fundamental period* of the function. Many signals encountered in practice are naturally periodic or can be well-approximated as periodic over a finite time window. Examples include oscillatory electrical signals, sound waves, and repeating patterns in images. By first understanding how to represent periodic signals, we will build tools that can later be adapted to more general settings.

2.2 Sines and cosines as periodic signals

Among all periodic functions, sines and cosines play a particularly important role. They are simple, smooth, and highly structured, and they will eventually serve as the basic building blocks for representing more complicated signals.

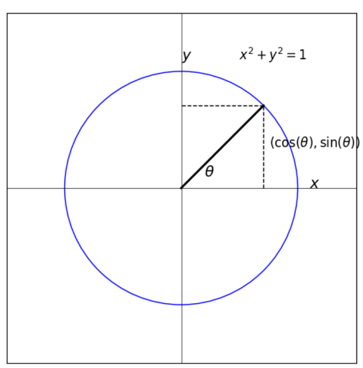

Consider a unit circle centered at the origin. Let a point move along the circle, making an angle θ with the positive x-axis, as shown in the above figure. The coordinates of this point are defined to be $(\cos\theta, \sin\theta)$, i.e., $\cos(\theta)$ is the x-coordinate of the point and $\sin(\theta)$ is the y-coordinate of the point.

This definition extends the familiar right-triangle definitions of sine and cosine. When θ lies between 0 and $\pi/2$, $\sin(\theta)$ and $\cos(\theta)$ agree with the ratios learned from right triangles, i.e.,

$$\sin(\theta) = \frac{\text{opposite}}{\text{hypotenuse}} \quad \text{and} \quad \cos(\theta) = \frac{\text{adjacent}}{\text{hypotenuse}}$$

However, the unit-circle interpretation allows us to define $\sin(\theta)$ and $\cos(\theta)$ for *all* real values of θ, including angles larger than $\pi/2$ and negative angles.

As θ increases, the point moves smoothly around the circle, and the values of $\cos(\theta)$ and $\sin(\theta)$ vary continuously between -1 and 1. After the angle increases by 2π, the point returns to its original position, and both $\sin(\theta)$ and $\cos(\theta)$ repeat their values. This immediately implies that sine and cosine are periodic with period 2π, i.e., $\sin(\theta + 2\pi k) = \sin(\theta)$ and $\cos(\theta + 2\pi k) = \cos(\theta)$ for all integer k.

To use sine and cosine as signals, we need to relate the angle θ to time. Suppose the point on the unit circle completes f full revolutions around the circle every second. In this case, the angle increases linearly with time according to

$$\theta(t) = 2\pi f t.$$

Substituting this into $\sin(\theta)$ and $\cos(\theta)$ gives the time-varying signals

$$\sin(2\pi f t) \quad \text{and} \quad \cos(2\pi f t).$$

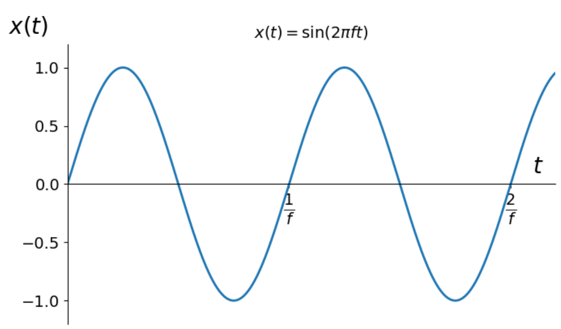

$x(t) = \sin(2\pi f t)$

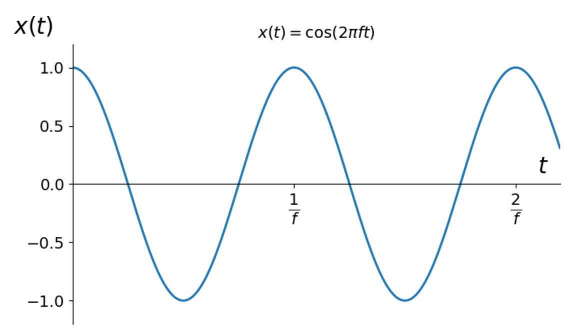

$x(t) = \cos(2\pi f t)$

These functions represent oscillatory signals that repeat regularly over time (as shown in the figure above). To build intuition, we explain the period of the sine and cosine signals in two different but equivalent ways.

Intuitive interpretation. Recall the rotating point on the unit circle. If the point completes f full revolutions around the circle every second, then it takes exactly $1/f$ seconds to complete one full cycle. Since the signal repeats itself every time the point completes a full revolution, the signal must repeat every $1/f$ seconds.

Mathematical verification. We can also verify this directly from the expression $\sin(2\pi ft)$ (and similarly for $\cos(2\pi ft)$). Consider increasing time by $1/f$:

$$\sin\big(2\pi f(t+1/f)\big) = \sin(2\pi ft + 2\pi) = \sin(2\pi ft),$$

where we used the fact that $\sin(\theta + 2\pi) = \sin(\theta)$. This shows that the signal repeats when time is increased by $1/f$, and hence the period is

$$T = \frac{1}{f}.$$

Frequency f and Period T. The parameter f is called the *frequency* of the signal and is measured in cycles per second, or Hertz (Hz). It indicates how many complete oscillations the signal makes in one second. The period T is measured in seconds and represents the amount of time required for the signal to complete one full oscillation. Larger values of f correspond to faster oscillations and shorter periods, while smaller values of f correspond to slower oscillations and longer periods. In the next section, we will build on this understanding to see how sinusoidal signals can be used to represent more general periodic functions.

An interactive visualization helpful in understanding the intuition behind the sine and cosine signal can be found here.

3 Fourier Series

Let $x(t)$ be a periodic signal with period T. The key idea is to represent $x(t)$ as a sum of sine and cosine signals whose frequencies are integer multiples of the fundamental frequency $1/T$.

Fact 1 (Fourier series)**.** *Any well-behaved periodic signal $x(t)$ with period T can be written as*

$$x(t) = b_0 \; + \; \sum_{j=1}^{\infty} \left(a_j \sin\!\left(2\pi \frac{j}{T} t\right) + b_j \cos\!\left(2\pi \frac{j}{T} t\right) \right).$$

Moreover, this representation is unique*: there is exactly one choice of coefficients*

$$b_0, a_1, b_1, a_2, b_2, \ldots$$

that produces the signal $x(t)$.

The phrase *well-behaved* is used informally here; essentially all natural signals satisfy this condition. In this course, we will not worry about this distinction and will simply assume that such a Fourier series representation exists for the signals we study.

The Fourier series gives us two equivalent ways to represent the same periodic signal. The *time-domain representation* describes the signal directly as a function of time, i.e., $x(t)$. The **frequency-domain representation** or **Fourier representation** describes the same signal using the collection of Fourier series coefficients

$$\{b_0, a_1, b_1, a_2, b_2, \ldots\},$$

which specify how much each sinusoidal component contributes to the signal. In general, representing a signal exactly in this way may require infinitely many such coefficients.

3.1 Fourier Series Coefficients and Harmonics

Constant Term: The term b_0 is called the *constant term* or *zero-frequency component*. It equals the average value of the signal over one period:

$$b_0 = \frac{1}{T}\int_0^T x(t)\,dt.$$

Intuitively, b_0 shifts the signal up or down by a constant amount. All other terms in the Fourier series are sines and cosines, and each of those has zero average over a full cycle, so the constant offset of the signal is entirely captured by b_0.

Harmonics: The remaining terms are organized by *harmonics*. For each positive integer j, the Fourier series includes the sine and cosine signals

$$\sin\left(2\pi\frac{j}{T}t\right) \quad \text{and} \quad \cos\left(2\pi\frac{j}{T}t\right),$$

which are called the *jth harmonics*. The first 8 harmonics are presented in the figure below.

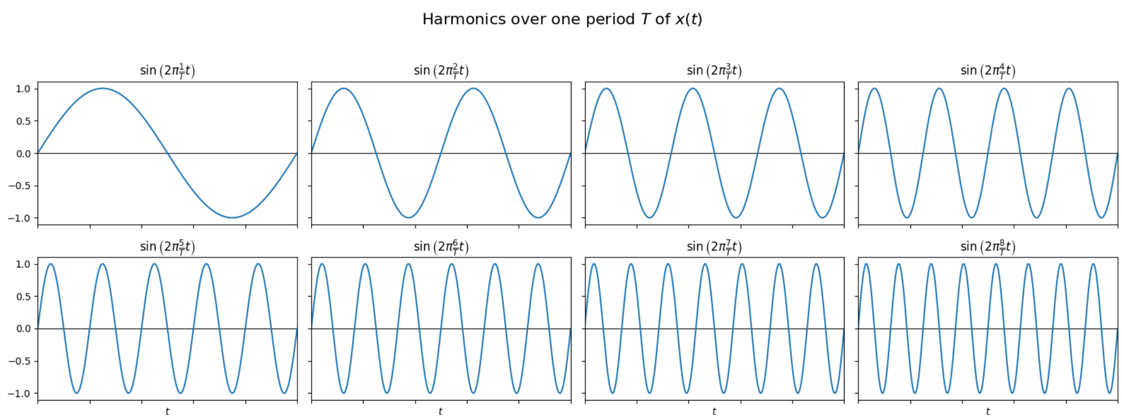

To interpret these, let us compute their frequency and understand what they do over one period of the original signal. Consider the jth harmonic $\sin(2\pi\frac{j}{T}t)$. Its frequency is

$$f = \frac{j}{T} \quad \text{(cycles per second)}.$$

By definition of frequency, a signal with frequency f completes f cycles in one second. Therefore, in T seconds it completes

$$(\#\text{ cycles in } T \text{ seconds}) = f\cdot T = \frac{j}{T}\cdot T = j$$

cycles. In other words, the first harmonic ($j = 1$) completes exactly one cycle over the interval of length T, the second harmonic completes two cycles over the same interval, the third completes three cycles, and so on. The same interpretation holds for the cosine term $\cos(2\pi\frac{j}{T}t)$. This is the reason these particular frequencies appear in the Fourier series: they are precisely the sinusoids that complete an integer number of cycles within one period of $x(t)$, and hence they "fit" the periodic structure of the signal.

Finally, the Fourier coefficients determine how strongly each harmonic contributes. The pair (a_j, b_j) controls the contribution of the jth harmonic: a_j is the weight on the sine term at frequency j/T, and b_j is the weight on the cosine term at the same frequency. There are explicit formulas for computing these coefficients:

$$a_j = \frac{2}{T}\int_0^T x(t)\sin\left(2\pi\frac{j}{T}t\right)\,dt, \qquad b_j = \frac{2}{T}\int_0^T x(t)\cos\left(2\pi\frac{j}{T}t\right)\,dt, \quad j \geq 1.$$

These formulas are NOT required for this course. More details about these formulas, along with a linear algebraic interpretation of the Fourier series is presented in the appendix.

It is often useful to summarize the contribution of the jth harmonic using a single quantity. We define

$$A_j = \sqrt{a_j^2 + b_j^2},$$

which represents the overall strength (or amplitude) of the jth harmonic. Intuitively, A_j measures how much energy the signal contains at frequency j/T, independent of how that contribution is split between the sine and cosine components. A larger value of A_j indicates that oscillations at frequency j/T play a more significant role in the signal.

4 Even Signals and Fourier Cosine Series

In many situations, signals exhibit additional structure that can be exploited to simplify their representation. One particularly important type of structure is symmetry, which we now formalize.

Definition 3 (Even signal)**.** *A signal $x(t)$ is called* even *if*

$$x(t) = x(-t) \quad \textit{for all } t.$$

Intuitively, an even signal symmetric about the y-axis or the left half of the signal can be obtained by taking a mirror image about the y-axis. Another related concept is that of odd signals.

Definition 4 (Odd signal)**.** *A signal $x(t)$ is called* odd *if*

$$x(t) = -x(-t) \quad \textit{for all } t.$$

An odd signal is one which is symmetric about the origin.

Sine and cosine as odd and even signals. Sine and cosine themselves have a simple symmetry structure:

$$\cos(t) = \cos(-t) \quad \text{(cosine is even)}, \qquad \sin(t) = -\sin(-t) \quad \text{(sine is odd)}.$$

This can be seen from the unit-circle interpretation: reflecting a point across the y-axis keeps its x-coordinate the same but flips the sign of its y-coordinate. Since $\cos(t)$ is the x-coordinate, it is symmetric about the y-axis (even), while $\sin(t)$ is the y-coordinate and changes sign (odd). The same symmetry holds for all harmonics: $\cos\left(2\pi\frac{j}{T}t\right)$ is even and $\sin\left(2\pi\frac{j}{T}t\right)$ is odd.

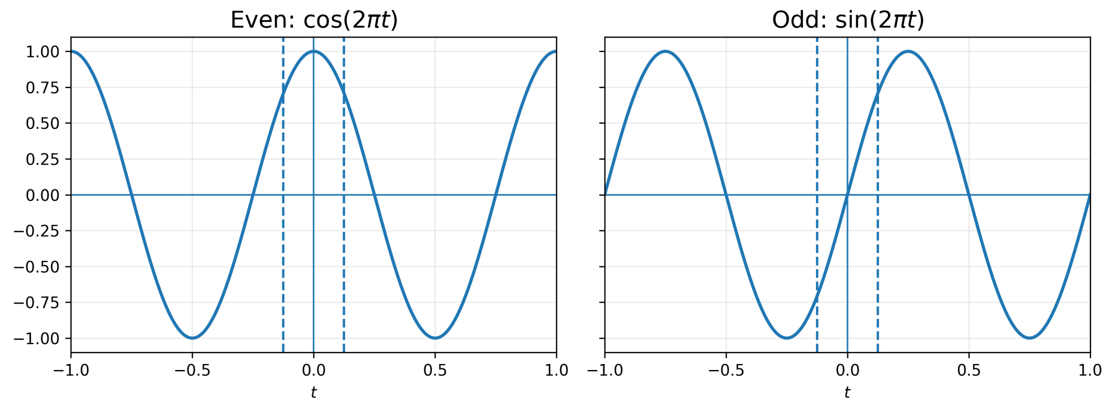

Fourier Cosine Series. This symmetry immediately explains why even signals do not need sine terms in their Fourier representation. If $x(t)$ is even, then its representation must also be even. Cosine terms preserve even symmetry, but any sine term is odd and would introduce an antisymmetric component, making the result no longer even. Therefore, for an even periodic signal, all sine coefficients must be zero, and the signal can be represented using only cosines (formal proof in appendix).

Fact 2 (Fourier cosine series)**.** *Any well-behaved even periodic signal $x(t)$ with period T can be written as*

$$x(t) = b_0 + \sum_{j=1}^{\infty} b_j \cos\left(2\pi \frac{j}{T} t\right),$$

and this representation is unique.

This representation is called the *Fourier cosine series*. It exploits the symmetry of even signals and uses only cosine basis functions, reducing the number of components needed to represent the signal.

5 Understanding What Different Frequencies Represent

Different frequency components in a Fourier (or cosine) series capture different types of structure in a signal. Broadly speaking, low-frequency components describe coarse, global behavior, while high-frequency components describe local, fine-scale variations. The number of frequency components needed to well-approximate the signal depend critically on how smooth the signal is.

- **Low-frequency components: coarse or global structure.** Low-frequency components vary slowly over time. They capture the overall shape or general trend of the signal, such as its average value, envelope, or large-scale variations. The constant term is the simplest example of a low-frequency component, representing the average of the signal, while the first few harmonics refine this average by adding gentle oscillations. If we retain only the low-frequency components of a signal, we preserve its global structure but lose fine detail.

- **High-frequency components: local and fine-scale structure.** High-frequency components vary rapidly over time. They capture local changes and fine details in the signal, including sharp transitions, edges, and noise. These components are essential for accurately representing sudden changes over short time intervals. Signals with significant local variation necessarily require higher-frequency components to capture that behavior.

The role of different frequency components depends strongly on the smoothness of the signal.

- **Discontinuities and sharp transitions.** A jump discontinuity is an extreme example of a local change. Since low-frequency components are smooth and slowly varying, they are fundamentally incapable of representing a sudden jump. Accurately capturing a discontinuity therefore requires contributions from many high-frequency components. As more high-frequency terms are added, the approximation becomes sharper near the jump, but convergence remains slow.

- **Corners and non-smooth points.** Even if a signal is continuous, it may not be smooth. For instance, a signal may have corners where the slope changes abruptly. These non-smooth points also rely heavily on high-frequency components, though typically fewer than required for a true discontinuity. In general, the less smooth a signal is, the more it depends on high-frequency terms.

- **Smooth signals require the fewest terms to represent.** Smooth signals, which change gradually and lack sharp transitions, are dominated by low-frequency components. Their high-frequency coefficients tend to be small, meaning that a good approximation can often be obtained using only a few low-frequency terms. Many natural signals fall into this category, which is why frequency-domain representations are so effective in practice.

An interactive visualization illustrating how signals of each of the above types are approximated using a finite number of frequency components can be found here.

Approximating a signal using a finite number of frequency components. In practice, representing a signal using infinitely many frequency components is neither feasible nor necessary. A central goal of signal representation is therefore to approximate a signal accurately using only a finite number of terms. When a signal can be well approximated by keeping only its lowest-frequency components, it admits a compact representation: most of the signal's essential information is captured using relatively few coefficients. This idea is fundamental to efficient representation and will later be exploited directly for image compression.

Approximating a signal using a finite number of frequency components corresponds to retaining the lowest-frequency terms in the representation and discarding higher-frequency ones. The effect of such an approximation is intuitive:

- the global structure and overall trends of the signal are preserved,

- local variations are suppressed,

- sharp edges and corners become smoother,

- fine details and noise are reduced.

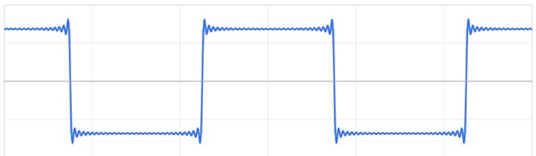

Ringing when a square wave is approximated using a finite number of frequency components

Ringing artifacts in a reconstructed image; observe the oscillations near edges.

A characteristic side effect of approximating a signal using only finitely many frequency components, especially when the original signal contains discontinuities, is the appearance of oscillations near sharp transitions. These oscillations, known as *ringing* or the Gibbs phenomenon, arise because smooth sinusoidal components struggle to approximate sudden jumps. This effect is illustrated in the figures above for both a one-dimensional square wave and a reconstructed image, where oscillations appear near sharp transitions. As more frequency components are included, the oscillations become narrower and more localized near the discontinuity, but their magnitude does not vanish entirely.

6 Finite-Duration Signals

So far, we have focused on periodic signals. In many practical settings, however, we are given signals that are defined only over a finite interval of time. A finite-duration signal is a signal that is defined only on a bounded interval, for example

$$x(t), \quad t \in [0, T].$$

Examples include a recorded audio clip of fixed length, a sensor measurement over a finite time window, or a single row or column of an image. To represent such signals in the frequency domain, we need to relate them to the periodic signals studied earlier.

Periodic extension. A natural approach is to construct a *periodic extension* of the signal. Given $x(t)$ defined on $[0, T]$, we form a new signal $x_p(t)$ by repeating $x(t)$ every T seconds, as shown in the figure below. By construction, $x_p(t)$ is periodic with period T and coincides with the original signal on $[0, T]$. Since $x_p(t)$ is periodic with period T, it admits a Fourier series representation of the form

$$x_p(t) = b_0 + \sum_{j=1}^{\infty} \left(a_j \sin\left(2\pi \frac{j}{T} t\right) + b_j \cos\left(2\pi \frac{j}{T} t\right) \right).$$

However, even if the original signal $x(t)$ is smooth on $[0, T]$, the periodic extension often introduces jump discontinuities at the boundaries (e.g., at $t = 0$ and $t = T$). These artificial discontinuities lead to significant high-frequency content in the representation.

Even periodic extension. An alternative construction is the *even periodic extension*. Starting from $x(t)$ defined on $[0, T]$, we first mirror the signal about the vertical axis to obtain a symmetric signal on $[-T, T]$. This symmetric signal is then repeated periodically, as illustrated in the figure below. Because the signal is defined symmetrically over an interval of length $2T$, the resulting periodic signal has period $2T$. This explains why the fundamental period of the even periodic extension is $2T$, not T.

The even periodic extension is both periodic (with period $2T$) and even. As a result, its frequency-domain representation involves only cosine terms and takes the form

$$x_{ep}(t) = b_0 + \sum_{j=1}^{\infty} b_j \cos\left(2\pi \frac{j}{2T} t\right).$$

The appearance of $2T$ in the denominator reflects the fact that the signal completes one full repetition over an interval of length $2T$.

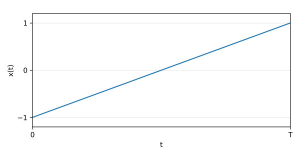

Finite-duration signal $x(t)$

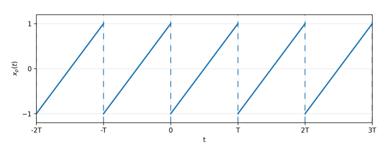

Periodic extension $x_p(t)$

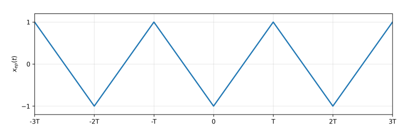

Even periodic extension $x_{ep}(t)$

Which representation to use? A finite-duration signal can be represented using either a Fourier series or a Fourier cosine series, depending on the extension used. The choice of extension is a modeling construct rather than a property of the signal itself, and different choices lead to different frequency-domain representations. This choice is guided by the application at hand. In practice, **even periodic extension is often preferred** due to its smoother behavior and more compact frequency representation.

If the original signal $x(t)$ is continuous on $[0, T]$, then its even periodic extension is also continuous at the boundaries. This avoids the artificial jump discontinuities introduced by ordinary periodic extension. As a result, the frequency-domain representation of the even periodic extension is typically much more compact: most of the signal's energy is concentrated in the lower-frequency cosine coefficients.

7 Discrete-Time Signals and the Discrete Cosine Transform

So far, we have worked with continuous-time signals. In practice, however, most signals we store, transmit, or process digitally are discrete.

Discrete-time signals. A discrete-time signal is a finite sequence of samples

$$y[0], y[1], \ldots, y[L-1],$$

where the index $n = 0, 1, \ldots, L-1$ denotes discrete time. Examples include digital audio samples, pixel intensities along a row or column of an image, and any signal obtained by sampling a continuous-time signal. A discrete-time signal of length L contains exactly L degrees of freedom. Our goal is to represent this signal in a frequency-domain form that captures its structure efficiently.

One useful way to build intuition and to understand why discrete-time signals arise in practice is to view them as samples of an underlying continuous-time signal. In real systems, we do not store or process a continuum of values. Instead, we record measurements at discrete time instants over a finite interval of length T. Interpreting the samples $y[n]$ as values of an underlying continuous-time signal at specific time instants allows us to reuse the ideas developed for finite-duration continuous-time signals while working with a finite, discrete representation.

7.1 Discrete Cosine Transform

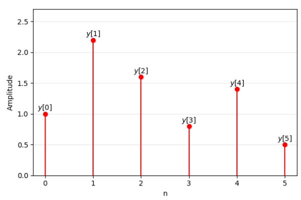

Discrete-time signal

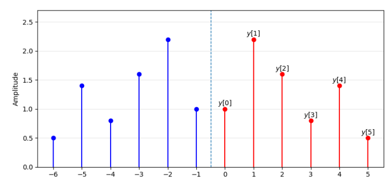

Mirrored signal

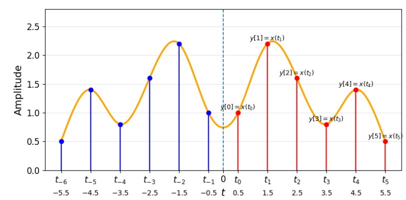

Continuous-time interpretation

As in the continuous-time case, an effective way to represent a finite-length signal is to construct an even extension. For a discrete-time signal $y[0], \ldots, y[L-1]$, this corresponds to mirroring the signal about its endpoints to create a symmetric extension, as illustrated in the figure above. Conceptually, this mirror image reduces boundary artifacts by avoiding artificial jumps at the endpoints, which would otherwise introduce unnecessary high-frequency components.

To make this construction precise, it is helpful to imagine an underlying continuous-time signal from which the discrete samples are obtained (see the continuous-time interpretation figure). Let t_n denote the continuous-time instant at which the nth sample is taken, so that $y[n] = x(t_n)$. Since we are working with discrete-time signals, it is natural to require that the samples be equally spaced, i.e., $t_n - t_{n-1} = 1$.

To enforce symmetry in the even extension (as shown in the mirrored figure), the sampling grid itself must be symmetric about the origin. In particular, the first sample to the right of zero and the first sample to the left of zero should be mirror images of each other, which means we want $t_0 = -t_{-1}$. Together with unit spacing, this implies that $t_0 = \frac{1}{2}$, and hence the sampling instants take the form

$$t_n = n + \tfrac{1}{2}, \qquad n = 0, 1, \ldots, L-1.$$

This half-sample shift is therefore not arbitrary: it follows directly from requiring equally spaced samples and symmetry about the origin. Beyond enforcing symmetry, centering each sample within its interval promotes smoothness in the extended signal. By avoiding samples exactly at the boundaries, this construction reduces artificial discontinuities, leading to a smoother even extension and, consequently, a more compact cosine-based frequency-domain representation.

With this continuous-time interpretation in place, we can now write down the frequency-domain representation of a discrete-time signal. The even periodic extension leads to a cosine-only representation, and evaluating the cosine basis functions of the underlying continuous-time signal at the sampling instants $t_n = n + \frac{1}{2}$ yields the discrete cosine representation.

Specifically, we represent the discrete-time signal $y[n]$ as

$$y[n] = b_0 + \sum_{j=1}^{???} b_j \cos\left(\frac{\pi j(n+\frac{1}{2})}{L}\right), \qquad n = 0, 1, \ldots, L-1.$$

This expression is called the *Discrete Cosine Transform* (DCT). The cosine terms arise from the even periodic extension, while the half-sample shift $(n+\frac{1}{2})$ follows directly from the symmetry and smoothness considerations discussed above. We next discuss the upper limit of the summation, i.e., the number of frequency components required to represent the signal.

A discrete-time signal of length L contains exactly L degrees of freedom, since specifying the signal requires choosing L real numbers. Any representation that captures the signal exactly therefore cannot require more than L independent parameters, and must use at least L parameters to avoid ambiguity. As a result, its frequency-domain representation requires only L terms: one coefficient for each degree of freedom in the signal. We show this formally in the appendix.

Fact 3 (Discrete Cosine Transform)**.** *Given a discrete-time signal $y[0], y[1], \ldots, y[L-1]$, the signal can be represented as*

$$y[n] = b_0 + \sum_{j=1}^{L-1} b_j \cos\left(\frac{\pi j(n+\frac{1}{2})}{L}\right), \qquad n = 0, 1, \ldots, L-1.$$

Moreover, the coefficients $\{b_0, b_1, \ldots, b_{L-1}\}$ are unique for a given signal.

This particular form is called the *DCT-II*. Other variants differ mainly in how the signal is extended at the boundaries and in the associated sampling conventions.

Time-domain and frequency-domain representations. The sequence $\{y[0], \ldots, y[L-1]\}$ is the time-domain representation of the signal. The coefficients $\{b_0, \ldots, b_{L-1}\}$ form its frequency-domain representation. Both representations contain the same information, but reveal different structure: the time-domain view shows sample values directly, while the frequency-domain view describes how much each cosine component contributes.

7.2 A basis interpretation of the DCT

The discrete cosine transform can be understood as expressing a discrete-time signal as a linear combination of fixed cosine *building blocks*. Each building block is a discrete-time cosine pattern of length L, and the signal is formed by adding together these patterns after scaling them by appropriate coefficients.

For each index $j = 0, 1, \ldots, L-1$, define the vector

$$\phi_j = \big[\phi_j[0],\, \phi_j[1],\, \ldots,\, \phi_j[L-1]\big]^\top,$$

where

$$\phi_j[n] = \begin{cases} 1, & j = 0, \\ \cos\left(\dfrac{\pi j(n+\frac{1}{2})}{L}\right), & j \geq 1, \end{cases} \qquad n = 0, 1, \ldots, L-1.$$

Each ϕ_j is itself a discrete-time signal of length L: ϕ_0 is a constant building block, while $\phi_1, \phi_2, \ldots$ oscillate increasingly rapidly as j increases.

The figures below illustrate each cosine building block in two complementary ways. The orange curve shows the underlying continuous-time cosine function from which the discrete pattern is derived, while the blue stems indicate the discrete samples $\phi_j[0], \phi_j[1], \ldots, \phi_j[L-1]$ obtained by evaluating this cosine at the sampling locations $t_n = n + \frac{1}{2}$. This visualization highlights the connection between the continuous-time interpretation and the discrete-time representation: the stems lie exactly on the continuous cosine, and higher-index building blocks correspond to increasingly rapid oscillations.

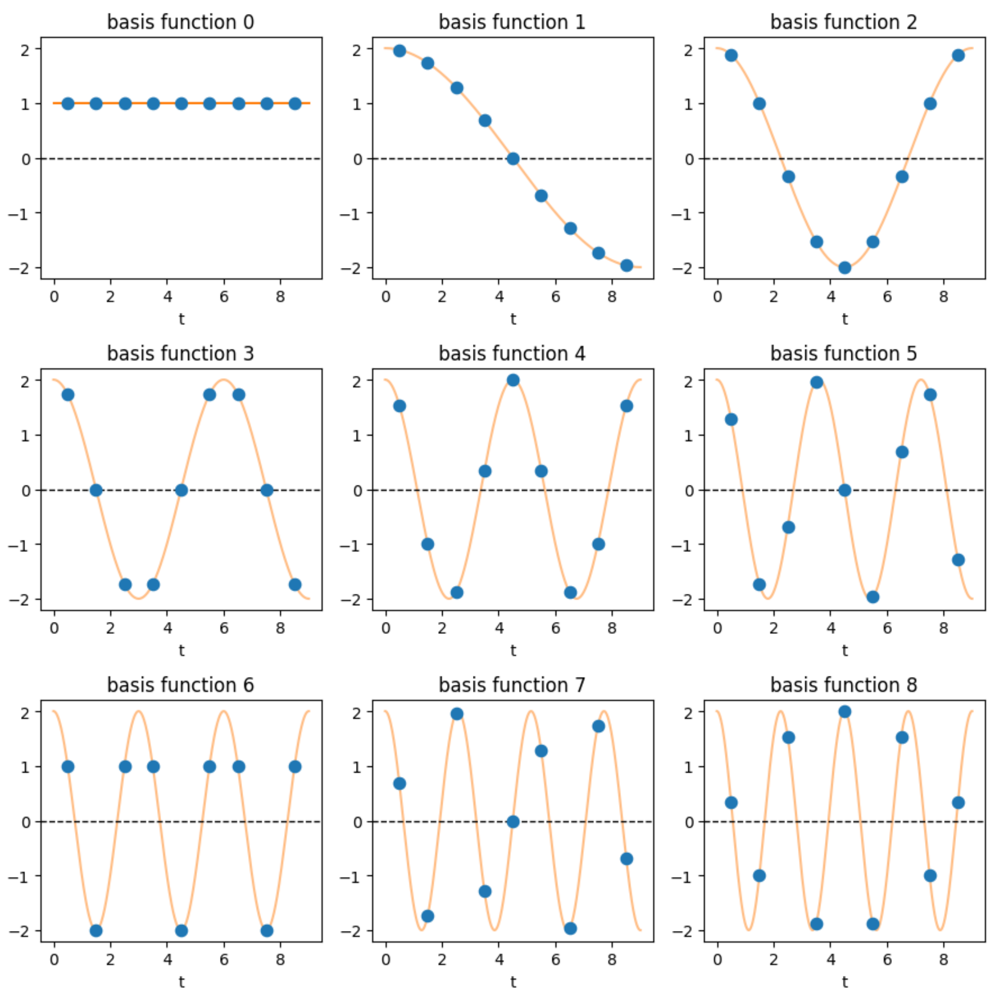

With this notation, the DCT representation can be written compactly as

$$y[n] = \sum_{j=0}^{L-1} b_j \, \phi_j[n],$$

or equivalently,

$$y = \sum_{j=0}^{L-1} b_j \, \phi_j.$$

In this building-block viewpoint, the coefficient b_j specifies how much of the jth cosine pattern is present in the signal. This makes clear how the DCT decomposes a signal into contributions from cosine patterns of different frequencies.

7.3 Computing the DCT coefficients

The constant coefficient is given by

$$b_0 = \frac{1}{L} \sum_{n=0}^{L-1} y[n],$$

which corresponds to the average (DC) component of the signal. The remaining coefficients capture oscillatory behavior at increasing frequencies. For $j \geq 1$, the coefficients are computed as

$$b_j = \frac{2}{L} \sum_{n=0}^{L-1} y[n] \, \cos\left(\frac{\pi j (n + \frac{1}{2})}{L} \right), \qquad j = 1, 2, \ldots, L-1.$$

Intuitively, each coefficient b_j measures how much of the jth cosine pattern is present in the signal.

DCT and IDCT. The operation that maps a discrete-time signal $\{y[n]\}$ to its cosine coefficients $\{b_j\}$ is called the *Discrete Cosine Transform* (DCT). The inverse operation, which reconstructs the signal from its coefficients, is called the *Inverse Discrete Cosine Transform* (IDCT). Together, the DCT and IDCT provide two equivalent representations of the same signal: one in the time domain and one in the frequency domain.

Computational efficiency. A direct implementation of the DCT or IDCT based on the definitions above requires on the order of L^2 operations. In practice, however, both transforms can be computed efficiently using fast Fourier transform (FFT) algorithms, reducing the computational cost to on the order of $L \log L$. This efficiency is one of the key reasons the DCT is widely used in applications such as image and video compression.

8 Two-Dimensional Discrete Cosine Transform

So far, we have focused on one-dimensional discrete-time signals. Many important signals in practice, however, are inherently two-dimensional. The most prominent example is an image, which can be viewed as a two-dimensional signal where pixel intensities vary across both horizontal and vertical directions.

Images as 2D signals. A grayscale image of size $L \times L$ can be represented as a matrix

$$X[m, n], \qquad m, n = 0, 1, \ldots, L-1,$$

where (m, n) denotes the spatial location of a pixel (m-th row and n-th column) and $X[m, n]$ is its intensity value. An $L \times L$ image has L^2 degrees of freedom. Our goal is to represent such a signal in the frequency domain in a way that captures its structure efficiently.

From 1D to 2D. The two-dimensional discrete cosine transform (2D DCT) is obtained by applying the one-dimensional DCT separately along each dimension. Conceptually, we first decompose the image along one direction (for example, rows), and then apply the same decomposition along the other direction (columns). This leads to a representation in terms of two-dimensional cosine patterns.

2D cosine building blocks. In two dimensions, each cosine building block is indexed by a pair of frequency indices (i, j) and has the form

$$\phi_{i,j}[m, n] = \cos\left(\frac{\pi i(m + \frac{1}{2})}{L}\right)\cos\left(\frac{\pi j(n + \frac{1}{2})}{L}\right),$$

where $i, j = 0, 1, \ldots, L-1$. The index i controls oscillations along the horizontal direction, while j controls oscillations along the vertical direction. Low values of (i, j) correspond to slowly varying, smooth patterns, while larger values correspond to finer spatial variations.

The 2D DCT representation. Using these building blocks, the image $X[m, n]$ can be represented exactly as

$$X[m, n] = \sum_{i=0}^{L-1}\sum_{j=0}^{L-1} A[i, j]\, \cos\left(\frac{\pi i(m + \frac{1}{2})}{L}\right)\cos\left(\frac{\pi j(n + \frac{1}{2})}{L}\right), \qquad m, n = 0, 1, \ldots, L-1.$$

The coefficients $\{A[i, j]\}$ constitute the two-dimensional DCT of the image. The coefficient $A[0, 0]$ corresponds to the average intensity of the image.

This representation uses exactly L^2 coefficients—one for each degree of freedom in the image—and is therefore an exact change of representation, not an approximation. The pixel values $\{X[m, n]\}$ constitute the *spatial-domain* (or image-domain) representation, while the coefficients $\{A[i, j]\}$ form the corresponding *frequency-domain* representation. Both contain the same information about the image, but emphasize different structure.

Interpretation of frequency components. The pair (i, j) indicates how rapidly the image varies along each spatial direction. Coefficients with small indices capture coarse structure and smooth intensity variations, while coefficients with larger indices capture edges, textures, and fine details. For most natural images, the energy of the DCT coefficients is concentrated at low frequencies.

8.1 Why the 2D DCT Is Useful for Image Compression?

The two-dimensional DCT plays a central role in image compression because it reorganizes image information in a way that makes compression both effective and visually acceptable. This usefulness stems from two key properties.

Energy compaction. Natural images tend to vary smoothly over space, with large regions of slowly changing intensity punctuated by edges and textures. As a result, when an image is represented using the 2D DCT, most of the signal energy is concentrated in a small number of low-frequency coefficients $A[i, j]$, particularly those with small values of i and j. High-frequency coefficients, which correspond to rapid spatial oscillations, are often much smaller in magnitude.

This phenomenon, known as *energy compaction*, means that a large fraction of the image's information can be captured using only a small subset of the DCT coefficients. The interactive visualization here illustrates this effect by showing that retaining only low-frequency coefficients already produces a recognizable approximation of the image, while many high-frequency coefficients can be discarded with limited impact.

Decorrelation of image data. A second, equally important reason the 2D DCT is useful for compression is that it *decorrelates* image data. In the spatial domain, neighboring pixel values $X[m, n]$ are highly correlated: knowing the value of one pixel gives strong information about the values of nearby pixels. This redundancy limits the effectiveness of compression directly in the image domain.

The 2D DCT transforms the image into a set of coefficients $A[i, j]$ that are much less correlated with one another. Each coefficient captures a distinct spatial pattern, and the value of one coefficient provides little information about the values of others. This reduction in statistical dependence makes it possible to encode the coefficients more efficiently, since redundancy across pixels has been largely removed.

Implication for compression. Together, energy compaction and decorrelation make the 2D DCT particularly well suited for image compression. Energy compaction concentrates most of the image information into a small number of coefficients, while decorrelation reduces redundancy among them. In practical compression systems, this representation also enables perceptually informed quantization, where higher-frequency coefficients (corresponding to fine details to which the human visual system is less sensitive) are quantized more coarsely. While we do not study such quantization strategies in this course, this aspect further contributes to the effectiveness of DCT-based image compression in practice.

9 Discrete Fourier Transform

For completeness, we also study the *Discrete Fourier Transform* (DFT). While the DCT is our primary tool for image compression, the DFT provides an alternative frequency-domain representation that is especially natural for signals that are inherently periodic. Given a signal

$$y[0], y[1], \ldots, y[L-1],$$

the DFT implicitly assumes that this sequence repeats with period L, i.e.,

$$y[n+L] = y[n] \quad \text{for all } n.$$

Under this assumption, the signal can be represented using sinusoidal components that complete an integer number of cycles over one period of length L. Because the signal is periodic, the natural building blocks are discrete-time sine and cosine signals of the form

$$\sin\left(2\pi \frac{j}{L} n\right) \quad \text{and} \quad \cos\left(2\pi \frac{j}{L} n\right),$$

which complete exactly j cycles over L samples. These play the same role for periodic discrete-time signals as the Fourier series does for periodic continuous-time signals.

The DFT representation. Using these building blocks, a real-valued discrete-time signal can be written as

$$y[n] = b_0 + \sum_{k=1}^{???} \left(a_k \sin\left(2\pi \frac{k}{L} n\right) + b_k \cos\left(2\pi \frac{k}{L} n\right) \right), \qquad n = 0, 1, \ldots, L-1.$$

The coefficients $\{a_k, b_k\}$ constitute the discrete Fourier transform of the signal.

A discrete-time signal of length L contains exactly L degrees of freedom. Any exact frequency-domain representation must therefore use exactly L independent coefficients. In the DFT representation, each frequency index j contributes a sine coefficient and a cosine coefficient. Together with the constant term, this means that using harmonics up to $j = (L-1)/2$ provides exactly L independent parameters for a real-valued signal.

Fact 4 (Discrete Fourier Transform)**.** *Let $y[0], y[1], \ldots, y[L-1]$ be a real-valued discrete-time signal. If L is odd, the signal can be represented as*

$$y[n] = b_0 + \sum_{j=1}^{(L-1)/2} \left(a_j \sin\left(2\pi \frac{j}{L} n\right) + b_j \cos\left(2\pi \frac{j}{L} n\right) \right), \qquad n = 0, 1, \ldots, L-1.$$

The coefficients in this representation are unique.

If L is even, an additional cosine term at frequency $k = L/2$ is included, since $\sin(\pi n) = 0$ for all integers n.

Choice of representation. Representing a discrete-time signal using the DFT or the DCT is a modeling choice rather than a property of the signal itself. The DFT is based on a periodic extension of the signal, while the DCT is based on an even extension that promotes smoothness at the boundaries. Both representations are exact and valid. The choice between them depends on the application: the DCT is typically preferred for compression of natural images due to its energy compaction properties, while the DFT is often more natural for analyzing signals where the periodic structure is intrinsic, such as in communications and time-series analysis.

10 Beyond What We Studied

Time-domain representation		**Frequency-domain representation**		
Time type	**Time structure**	**Name**	**Frequency type**	**Frequency-limited**
Continuous	Periodic	Fourier Series $\{b_0, a_1, b_1, \ldots\}$	Discrete	No
Discrete	Finite-length (periodic extension)	Discrete Fourier Transform (DFT) $\{b_0, a_1, \ldots, b_{(L-1)/2}\}$	Discrete	Yes
Continuous	Aperiodic	Fourier Transform $\{X(f)\}$	Continuous	No
Discrete	Aperiodic	Discrete-Time Fourier Transform (DTFT) $\{X(f)\}$ (periodic)	Continuous	Yes

In this module, we studied a subset of frequency-domain representations that are relevant for signal representation and compression in this course. In particular, we covered Fourier series for continuous-time periodic signals, the discrete cosine transform (DCT) for finite-length discrete-time signals and images, and the discrete Fourier transform (DFT) as an alternative representation based on periodic extension.

Fourier analysis is far broader than what we explored here. As summarized in the table above, there are additional representations (highlighted in red) that apply to aperiodic signals in continuous and discrete time. These play a central role in many areas of signal processing and applied mathematics but fall outside the scope of this course.

A Appendices

The following appendices contain supplementary material that is not required for this course and is provided for additional reading. We begin with a brief primer on linear algebra.

B A Primer on Linear Algebra

In this appendix, we briefly review a few linear algebra concepts that are useful for understanding Fourier series, discrete cosine transforms, and frequency-domain representations. This material is not required for the course, but it provides helpful intuition and a unifying viewpoint. Throughout this appendix, we work exclusively in $\mathbb{R}^d$, the space of real-valued vectors of length d.

Vectors in $\mathbb{R}^d$. An element of $\mathbb{R}^d$ is a vector

$$x = (x_1, x_2, \ldots, x_d),$$

which we will typically think of as a column vector. Vectors can be added and scaled by real numbers in the usual component-wise way.

Basis of $\mathbb{R}^d$. A collection of vectors $\{v_1, v_2, \ldots, v_d\}$ in $\mathbb{R}^d$ is called a *basis* if every vector $x \in \mathbb{R}^d$ can be written uniquely as a linear combination of these vectors:

$$x = c_1 v_1 + c_2 v_2 + \cdots + c_d v_d$$

for some real numbers $c_1, \ldots, c_d$. The key points are:

- *Representation*: any vector can be expressed using the basis vectors.
- *Uniqueness*: the coefficients $c_1, \ldots, c_d$ are uniquely determined by x.

Why are there exactly d basis vectors? Intuitively, a vector in $\mathbb{R}^d$ has d degrees of freedom: specifying the vector requires choosing d real numbers. Each basis vector provides one independent "direction" along which we can vary the vector. To represent all possible vectors in $\mathbb{R}^d$ without redundancy or ambiguity, we therefore need exactly d such independent directions. If we had fewer than d vectors, we would not be able to represent every vector in $\mathbb{R}^d$; if we had more than d, the representation would no longer be unique. Thus, a basis for $\mathbb{R}^d$ must consist of exactly d vectors, reflecting the d degrees of freedom of the space.

Example 1. *Let $x \in \mathbb{R}^d$ be a vector*

$$x = (x_1, x_2, \ldots, x_d).$$

With respect to the standard basis $\{e_1, e_2, \ldots, e_d\}$, this vector can be written as

$$x = x_1 e_1 + x_2 e_2 + \cdots + x_d e_d.$$

In this representation, the coefficients are simply the coordinates of the vector. To make this concrete, consider the case $d = 5$ and the vector

$$x = (2, -1, 0, 3, 1).$$

Using the standard basis of $\mathbb{R}^5$, we can write

$$x = 2e_1 - e_2 + 0 \cdot e_3 + 3e_4 + e_5.$$

Thus, the vector is represented uniquely as a linear combination of the five basis vectors, with coefficients given directly by its entries.

Example 2. *Consider $\mathbb{R}^2$. Using the standard basis $e_1 = (1,0)$ and $e_2 = (0,1)$, any vector $x = (x_1, x_2)$ can be written as $x = x_1e_1 + x_2e_2$. For instance, if $x = (3,1)$, then its representation in the standard basis is simply $x = 3e_1 + e_2$.*

Now consider a different basis for $\mathbb{R}^2$, given by the vectors $v_1 = (1,1)$ and $v_2 = (1,-1)$. We can represent the same vector $x = (3,1)$ using this basis as well. Suppose $x = c_1v_1 + c_2v_2$. Writing this out gives

$$(3,1) = c_1(1,1) + c_2(1,-1) = (c_1 + c_2,\ c_1 - c_2).$$

Equating the two components yields the equations $c_1 + c_2 = 3$ and $c_1 - c_2 = 1$, which gives $c_1 = 2$ and $c_2 = 1$. Thus, the vector can also be written as $x = 2v_1 + v_2$.

This example illustrates that while the vector itself does not change, its coefficient representation depends on the choice of basis.

Linear independence. A collection of vectors $\{v_1, v_2, \ldots, v_k\}$ in $\mathbb{R}^d$ is said to be *linearly independent* if the only way to write

$$c_1v_1 + c_2v_2 + \cdots + c_kv_k = 0$$

is by choosing

$$c_1 = c_2 = \cdots = c_k = 0.$$

If there exists a nontrivial choice of coefficients (not all zero) that satisfies the equation above, then the vectors are said to be *linearly dependent*. Intuitively, linear independence means that no vector in the collection can be written as a combination of the others; each vector contributes a genuinely new direction.

Linear independence and bases. In $\mathbb{R}^d$, any collection of d linearly independent vectors automatically forms a basis. Linear independence guarantees uniqueness of representation (easy to verify through contradiction), and having d such vectors ensures that all d degrees of freedom of the space are captured. As a result, every vector in $\mathbb{R}^d$ can be expressed uniquely as a linear combination of these vectors.

Inner product in $\mathbb{R}^d$. To talk about angles, lengths, and orthogonality, we use the *inner product*. For vectors $x, y \in \mathbb{R}^d$, the (standard) inner product is defined as

$$\langle x, y \rangle = \sum_{i=1}^{d} x_i y_i.$$

Intuitively, the inner product measures how much two vectors "point in the same direction." If the inner product is large and positive, the vectors are strongly aligned; if it is negative, they point in largely opposite directions; and if it is zero, the vectors do not align at all.

The inner product also induces a notion of length:

$$\|x\| = \sqrt{\langle x, x \rangle},$$

which corresponds to the usual Euclidean length of a vector.

Finally, the inner product defines orthogonality. Two vectors x and y are said to be orthogonal if $\langle x, y \rangle = 0$. Geometrically, this means the vectors are perpendicular to each other. From a representation perspective, orthogonality means that the two vectors represent independent directions, with no overlap in their contributions.

Orthogonal and orthonormal bases. A basis $\{v_1, \dots, v_d\}$ is called *orthogonal* if

$$\langle v_i, v_j \rangle = 0 \quad \text{for all } i \neq j.$$

It is called *orthonormal* if, in addition,

$$\|v_i\| = 1 \quad \text{for all } i.$$

The standard basis $\{e_1, \dots, e_d\}$ is orthonormal.

Why orthogonality implies linear independence. Orthogonality provides a simple and powerful way to guarantee linear independence. Suppose $\{v_1, \dots, v_k\}$ is a collection of nonzero vectors such that $\langle v_i, v_j \rangle = 0$ for all $i \neq j$. If we had

$$c_1 v_1 + c_2 v_2 + \cdots + c_k v_k = 0,$$

then taking the inner product with v_j gives

$$c_j \|v_j\|^2 = 0.$$

Since $\|v_j\|^2 > 0$, we must have $c_j = 0$. This holds for every j, so all coefficients are zero, and the vectors are linearly independent. Along with the argument that d linear independent vectors form a basis for $\mathbb{R}^d$, this shows that d orthogonal vectors form a basis for $\mathbb{R}^d$.

Representation using orthonormal bases. Orthonormal bases make representation particularly simple. If $\{v_1, \dots, v_d\}$ is an orthonormal basis of $\mathbb{R}^d$, then the coefficient of v_j in the expansion of a vector x is given by a simple inner product:

$$c_j = \langle x, v_j \rangle.$$

Thus, we can write

$$x = \sum_{j=1}^{d} \langle x, v_j \rangle \, v_j.$$

In words, the coefficient along each basis vector is obtained by "projecting" the vector onto that basis vector using the inner product.

Example 3. *Consider the standard basis* $\{e_1, e_2, \ldots, e_d\}$ *of* $\mathbb{R}^d$*, where* e_i *has a* 1 *in the* i*th position and zeros elsewhere. For* $i \neq j$*, we have* $\langle e_i, e_j \rangle = 0$*, since the vectors share no common nonzero entries, and* $\|e_i\| = \sqrt{\langle e_i, e_i \rangle} = 1$*. Thus, the standard basis is orthonormal.*

Now let $x = (x_1, x_2, \ldots, x_d) \in \mathbb{R}^d$*. Using the orthonormal-basis formula, the coefficient corresponding to* e_j *is* $\langle x, e_j \rangle = x_j$*, because* $\langle x, e_j \rangle = x_1 \cdot 0 + \cdots + x_j \cdot 1 + \cdots + x_d \cdot 0$*. Therefore, the representation*

$$x = \sum_{j=1}^{d} \langle x, e_j \rangle \, e_j$$

reduces to the familiar coordinate form $x = x_1 e_1 + x_2 e_2 + \cdots + x_d e_d$*.*

C Fourier Series

C.1 Derivation of Fourier Series Coefficients

In this appendix, we briefly explain how the formulas for the Fourier series coefficients arise. Throughout, we assume that a periodic signal $x(t)$ with period T admits a Fourier series representation of the form

$$x(t) = b_0 + \sum_{j=1}^{\infty} \left(a_j \sin\!\left(2\pi \frac{j}{T} t\right) + b_j \cos\!\left(2\pi \frac{j}{T} t\right) \right),$$

and we show how the coefficients a_j and b_j can be recovered from $x(t)$.

We begin with the constant term. Integrating both sides over one period $[0, T]$ gives

$$\int_0^T x(t)\, dt = \int_0^T b_0 \, dt + \sum_{j=1}^{\infty} \left(a_j \int_0^T \sin\!\left(2\pi \frac{j}{T} t\right) \, dt + b_j \int_0^T \cos\!\left(2\pi \frac{j}{T} t\right) \, dt \right).$$

Each sine and cosine term integrates to zero over a full number of cycles, so all terms in the summation vanish. What remains is

$$\int_0^T x(t)\, dt = b_0 T,$$

which yields

$$b_0 = \frac{1}{T} \int_0^T x(t)\, dt.$$

Next, we derive the formula for a_k for a fixed positive integer k. Multiply both sides of the Fourier series expression by $\sin\!\left(2\pi \frac{k}{T} t\right)$ and integrate over $[0, T]$:

$$\begin{aligned}
\int_0^T x(t) \sin\!\left(2\pi \frac{k}{T} t\right) \, dt = &\int_0^T b_0 \sin\!\left(2\pi \frac{k}{T} t\right) \, dt \\
&+ \sum_{j=1}^{\infty} \Bigg(a_j \int_0^T \sin\!\left(2\pi \frac{j}{T} t\right) \sin\!\left(2\pi \frac{k}{T} t\right) \, dt \\
&\qquad + b_j \int_0^T \cos\!\left(2\pi \frac{j}{T} t\right) \sin\!\left(2\pi \frac{k}{T} t\right) \, dt \Bigg).
\end{aligned}$$

After multiplying both sides by $\sin\bigl(2\pi\frac{k}{T}t\bigr)$ and integrating over $[0,T]$, the right-hand side consists of four types of terms: the constant term, cosine–sine cross terms, sine–sine terms with different frequencies, and the sine–sine term at the matching frequency. Each of these simplifies as follows:

- **Constant term.** The constant term vanishes since integrating a sine over a full number of cycles gives zero:
$$\int_0^T \sin\left(2\pi\frac{k}{T}t\right)\,dt = 0.$$

- **Cosine–sine cross terms.** For any integers j and k,
$$\begin{aligned}\int_0^T \cos\left(2\pi\frac{j}{T}t\right)\sin\left(2\pi\frac{k}{T}t\right)\,dt &= \frac{1}{2}\int_0^T \left[\sin\left(2\pi\frac{j+k}{T}t\right)+\sin\left(2\pi\frac{k-j}{T}t\right)\right]dt\\ &= 0,\end{aligned}$$
since both sine terms complete an integer number of cycles over $[0,T]$.

- **Sine–sine terms with different frequencies.** When $j\neq k$,
$$\begin{aligned}\int_0^T \sin\left(2\pi\frac{j}{T}t\right)\sin\left(2\pi\frac{k}{T}t\right)\,dt &= \frac{1}{2}\int_0^T \left[\cos\left(2\pi\frac{j-k}{T}t\right)-\cos\left(2\pi\frac{j+k}{T}t\right)\right]dt\\ &= 0,\end{aligned}$$
again because each cosine completes an integer number of cycles over one period.

- **Sine–sine term at the matching frequency.** When $j=k$,
$$\begin{aligned}\int_0^T \sin^2\left(2\pi\frac{k}{T}t\right)\,dt &= \frac{1}{2}\int_0^T \left[1-\cos\left(4\pi\frac{k}{T}t\right)\right]dt\\ &= \frac{1}{2}\int_0^T 1\,dt = \frac{T}{2}.\end{aligned}$$

As a result, every term in the summation vanishes except the one corresponding to $j=k$ in the sine–sine product. This leaves
$$\int_0^T x(t)\sin\left(2\pi\frac{k}{T}t\right)\,dt = a_k\cdot\frac{T}{2},$$
and hence
$$a_k = \frac{2}{T}\int_0^T x(t)\sin\left(2\pi\frac{k}{T}t\right)\,dt.$$

An entirely analogous argument, multiplying both sides by $\cos\bigl(2\pi\frac{k}{T}t\bigr)$ and integrating over $[0,T]$, yields
$$b_k = \frac{2}{T}\int_0^T x(t)\cos\left(2\pi\frac{k}{T}t\right)\,dt.$$

The key idea behind all these derivations is that over one full period, different sine and cosine functions are mutually orthogonal: when integrated together, all “cross terms” vanish, leaving only the coefficient of the matching frequency.

C.2 A Linear Algebraic Interpretation of Fourier Series

We now give a linear algebra interpretation of Fourier series, analogous to the discussion for vectors in $\mathbb{R}^d$. Instead of finite-dimensional vectors, we consider the space of all real-valued functions that are periodic with period T.

To define geometric notions such as length, angle, and orthogonality for functions, we introduce an inner product. For two periodic functions $f(t)$ and $g(t)$ with period T, define

$$\langle f, g \rangle = \frac{2}{T} \int_0^T f(t) g(t) \, dt.$$

This inner product measures how much two functions overlap over one period, in direct analogy with the dot product in $\mathbb{R}^d$.

With respect to this inner product, we can make precise what we mean by *orthogonal*: two periodic functions f and g are orthogonal if $\langle f, g \rangle = 0$, i.e., if

$$\frac{2}{T} \int_0^T f(t) g(t) \, dt = 0.$$

Orthogonal Basis. Consider the collection of functions

$$1, \qquad s_j(t) = \sin\left(2\pi \frac{j}{T} t\right), \qquad c_j(t) = \cos\left(2\pi \frac{j}{T} t\right), \quad j = 1, 2, 3, \ldots$$

Then the derivations in the previous subsection directly show that:

- $\langle 1, s_j \rangle = \langle 1, c_j \rangle = 0$ for all $j \geq 0$
- $\langle s_j, s_k \rangle = \langle c_j, c_k \rangle = 0$ for all $j \neq k$
- $\langle s_j, c_k \rangle = 0$ for all $j, k \geq 0$

This shows that these form an orthogonal set of functions. Moreover,

- $\|s_j\|^2 = \langle s_j, s_j \rangle = 1$ and $\|c_j\|^2 = \langle c_j, c_j \rangle = 1$ for all $j \geq 0$
- For the constant function, $\|1\|^2 = \langle 1, 1 \rangle = \frac{2}{T} \int_0^T 1 \, dt = 2.$ And hence $\|1\| = \sqrt{2}$.

As a result, this collection of functions is not strictly orthonormal due to the normalization of the constant function. This distinction is inconsequential for our purposes and is handled explicitly in the definition of the coefficient b_0.

Together, this collection of functions forms an orthogonal basis (with all functions having unit norm except for the constant function). We do not show here that this collection indeed forms a basis for the space of periodic functions with period T; this fact is taken as given. Intuitively, these functions play the same role as an orthogonal basis in $\mathbb{R}^d$: they provide independent directions in which a signal can vary.

Fourier Coefficients as Projections. The Fourier coefficients are obtained by projecting the signal onto these functions using the inner product. Specifically,

$$a_j = \left\langle x(t), \sin\left(2\pi \frac{j}{T} t\right) \right\rangle, \qquad b_j = \left\langle x(t), \cos\left(2\pi \frac{j}{T} t\right) \right\rangle,$$

while the constant coefficient satisfies

$$b_0 = \frac{1}{2} \langle x(t), 1 \rangle,$$

where the factor of $\frac{1}{2}$ compensates for the fact that the constant function has norm $\sqrt{2}$ under the chosen normalization. Thus, computing Fourier series coefficients is directly analogous to computing coordinates of a vector in an orthogonal basis, exactly as in finite-dimensional linear algebra.

D Fourier Cosine Series

In this appendix, we show explicitly why the sine coefficients in the Fourier series of an even signal are zero. Let $x(t)$ be a well-behaved periodic signal with period T that is *even*, i.e.,

$$x(t) = x(-t) \quad \text{for all } t.$$

Recall that the sine coefficients in the Fourier series are given by

$$a_j = \frac{2}{T} \int_0^T x(t) \sin\left(2\pi \frac{j}{T} t\right) \, dt.$$

Because $x(t)$ and $\sin\left(2\pi \frac{j}{T} t\right)$ are both periodic with period T, their product is also periodic with period T. Hence we may integrate over any interval of length T. In particular,

$$a_j = \frac{2}{T} \int_0^T x(t) \sin\left(2\pi \frac{j}{T} t\right) \, dt = \frac{2}{T} \int_{-T/2}^{T/2} x(t) \sin\left(2\pi \frac{j}{T} t\right) \, dt.$$

Define

$$z(t) \ := \ x(t) \sin\left(2\pi \frac{j}{T} t\right).$$

We claim that $z(t)$ is an odd function. To see this, compute $z(-t)$:

$$\begin{aligned} z(-t) &= x(-t) \sin\left(2\pi \frac{j}{T} (-t)\right) \\ &= x(t) \left(-\sin\left(2\pi \frac{j}{T} t\right) \right) \\ &= -\, x(t) \sin\left(2\pi \frac{j}{T} t\right) \\ &= -z(t). \end{aligned}$$

Thus, the signal z is odd. Here the second equality follows from the even nature of $x(t)$ and the odd nature of $\sin(\cdot)$. In general, the product of an even and an odd signal is always odd.

Now, let $z(t)$ be any odd function and let $A > 0$. Consider

$$I := \int_{-A}^{A} z(t)\,dt.$$

We split the integral:

$$I = \int_{-A}^{0} z(t)\,dt + \int_{0}^{A} z(t)\,dt.$$

In the first term, make the change of variables $u = -t$ (so $t = -u$ and $dt = -du$). When $t = -A$, we have $u = A$, and when $t = 0$, we have $u = 0$. Hence,

$$\begin{aligned}
\int_{-A}^{0} z(t)\,dt &= \int_{u=A}^{u=0} z(-u)\,(-du) \\
&= \int_{0}^{A} z(-u)\,du \\
&= \int_{0}^{A} \big(-z(u)\big)\,du \qquad \text{(since } z \text{ is odd)} \\
&= -\int_{0}^{A} z(u)\,du.
\end{aligned}$$

Therefore,

$$I = \left(-\int_{0}^{A} z(u)\,du\right) + \int_{0}^{A} z(t)\,dt = 0.$$

So, the integral of an odd function over $[-A, A]$ is zero. Applying this with $A = T/2$ and $z(t) = x(t)\sin(2\pi\frac{j}{T}t)$, we get

$$a_j = \frac{2}{T}\int_{-T/2}^{T/2} x(t)\sin\!\left(2\pi\frac{j}{T}t\right)\,dt = 0.$$

This shows that an even periodic signal has no sine components in its Fourier series.

E Discrete Cosine Transform

E.1 Why Only L Cosine Terms Are Needed

In this appendix, we justify formally why a discrete-time signal of length L requires only L cosine terms in its discrete cosine transform representation. Consider a discrete-time signal

$$y = [\,y[0], y[1], \ldots, y[L-1]\,]^\top \in \mathbb{R}^L.$$

Thus, the space of all length-L discrete-time signals is the L-dimensional vector space $\mathbb{R}^L$.

For $j = 0, 1, \ldots, L-1$, recall the discrete cosine vectors

$$\phi_j = \big[\phi_j[0],\, \phi_j[1],\, \ldots,\, \phi_j[L-1]\big]^\top,$$

where

$$\phi_j[n] = \begin{cases} 1, & j = 0, \\ \cos\!\left(\dfrac{\pi j(n+\frac{1}{2})}{L}\right), & j \geq 1, \end{cases} \qquad n = 0, 1, \ldots, L-1.$$

Each ϕ_j is a vector in $\mathbb{R}^L$, i.e., a discrete-time signal of length L.

Step 1: The cosine vectors form a basis of $\mathbb{R}^L$. We work with the standard inner product on $\mathbb{R}^L$, defined for vectors $u, v \in \mathbb{R}^L$ as

$$\langle u, v\rangle = \sum_{n=0}^{L-1} u[n]\, v[n].$$

With respect to this inner product, the vectors $\phi_0, \phi_1, \ldots, \phi_{L-1}$ are mutually orthogonal. Since each ϕ_j is nonzero, orthogonality implies linear independence. By the discussion in the linear algebra primer, any collection of L linearly independent vectors in $\mathbb{R}^L$ forms a basis. We omit the explicit orthogonality calculation here; it follows from standard trigonometric identities and the specific sampling locations used in the DCT.

Step 2: Basis size determines the number of terms. Because $\{\phi_0, \phi_1, \ldots, \phi_{L-1}\}$ is a basis of $\mathbb{R}^L$, any discrete-time signal $y \in \mathbb{R}^L$ admits a unique expansion of the form

$$y = \sum_{j=0}^{L-1} b_j\, \phi_j.$$

Each basis vector ϕ_j contributes exactly one scalar coefficient b_j. Thus, using L basis vectors is mathematically equivalent to using L terms in the summation: one term for each independent direction in $\mathbb{R}^L$. No additional terms are required, since any further cosine vectors would be expressible as linear combinations of these L basis vectors, and no fewer terms would suffice to represent all signals uniquely.

E.2 An alternative viewpoint: what happens beyond $j = L - 1$.

There is another way to see why cosine terms beyond index $L - 1$ provide no additional expressive power. This viewpoint works directly at the level of the cosine vectors themselves. Recall that the DCT building blocks are defined as

$$\phi_j[n] = \cos\left(\frac{\pi j (n + \frac{1}{2})}{L}\right), \qquad n = 0, 1, \ldots, L - 1.$$

Consider what happens when $j = L$. In this case,

$$\phi_L[n] = \cos\bigl(\pi(n + \tfrac{1}{2})\bigr) = 0 \quad \text{for all } n,$$

since $\cos(\pi k + \frac{\pi}{2}) = 0$ for all integers k. Thus, ϕ_L is the zero vector and contributes nothing to the representation. Now consider indices beyond L. For $j = L + m$ with $m \geq 1$, we have

$$\begin{aligned}
\phi_{L+m}[n] &= \cos\left(\frac{\pi(L + m)(n + \frac{1}{2})}{L}\right) \\
&= \cos\left(\pi(n + \tfrac{1}{2}) + \frac{\pi m (n + \frac{1}{2})}{L}\right) \\
&= -\cos\left(\frac{\pi m (n + \frac{1}{2})}{L}\right),
\end{aligned}$$

since $\cos(\theta + \pi/2 + \pi k) = -\cos(\theta)$ for integer k. This shows that

$$\phi_{L+m} = -\phi_{L-m}, \qquad m = 1, 2, \ldots, L-1.$$

In other words, cosine vectors beyond index $L-1$ do not generate new patterns: they either vanish entirely or reproduce existing vectors up to a sign. Because cosine vectors with index $j \geq L$ are either identically zero or duplicates (up to sign) of earlier vectors, including them in the summation cannot increase the set of signals that can be represented. Any contribution from such terms can be absorbed into the coefficients of existing basis vectors. This provides a direct, explicit explanation for why summing beyond $j = L-1$ offers no advantage.

F Discrete Fourier Transform

F.1 Why DFT Terms Beyond $(L-1)/2$ Are Not Needed

Recall that the DFT building blocks are

$$\sin\left(2\pi \frac{j}{L} n\right), \qquad \cos\left(2\pi \frac{j}{L} n\right), \qquad n = 0, 1, \ldots, L-1,$$

with frequency indexed by j.

Indices beyond $L-1$ repeat. Let $i = j + L$. Then for every integer n,

$$\cos\left(2\pi \frac{i}{L} n\right) = \cos\left(2\pi \frac{j+L}{L} n\right) = \cos\left(2\pi \frac{j}{L} n + 2\pi n\right) = \cos\left(2\pi \frac{j}{L} n\right),$$

and similarly,

$$\sin\left(2\pi \frac{i}{L} n\right) = \sin\left(2\pi \frac{j}{L} n\right).$$

Thus, using $i = j + L$ produces exactly the same sine/cosine patterns as using j, so indices beyond $L-1$ introduce no new terms.

Indices above $(L-1)/2$ mirror lower ones. Assume L is odd and define

$$i = j + \frac{L-1}{2}, \qquad j = 1, 2, \ldots, \frac{L-1}{2}.$$

Then

$$L - i = L - \left(j + \frac{L-1}{2}\right) = \frac{L+1}{2} - j.$$

Using the identities $\cos(2\pi n - \theta) = \cos(\theta)$ and $\sin(2\pi n - \theta) = -\sin(\theta)$ for integer n, we get for every integer n,

$$\cos\left(2\pi \frac{i}{L} n\right) = \cos\left(2\pi n - 2\pi \frac{L-i}{L} n\right) = \cos\left(2\pi \frac{L-i}{L} n\right),$$

$$\sin\left(2\pi \frac{i}{L} n\right) = \sin\left(2\pi n - 2\pi \frac{L-i}{L} n\right) = -\sin\left(2\pi \frac{L-i}{L} n\right).$$

Therefore, the frequency index $i = j + \frac{L-1}{2}$ produces cosine terms identical to those at index $L - i = \frac{L+1}{2} - j$, and sine terms that differ only by a sign. In particular, these indices lie in $\{1, \ldots, \frac{L-1}{2}\}$:

$$\frac{L+1}{2} - j \in \left\{1, 2, \ldots, \frac{L-1}{2}\right\}.$$

Hence all terms with indices above $(L-1)/2$ are redundant.