# The Frequency Domain

## 1 Motivation

In data science we call any structured object of interest– an image, an audio sequence, a video. . . – a *signal*. Signals are usually modeled as functions of time, space, etc. This paradigm makes it possible to analyze and manipulate such objects using mathematics, an approach known as signal processing. In these notes we describe a fundamental tool in signal processing: frequency representations. We begin by describing a few motivating applications.

### 1.1 Processing periodic signals

Signals often have periodic structure due to repeating patterns. This is evident in the electrocardiogram shown in Figure 1. Electrocardiography is a method to record heart activity by using electrodes to measure electrical changes on the skin. Figure 1 also shows the standard features used by doctors to perform diagnostics using ECGs[1]. Unfortunately, the data contains other periodic patterns that make it difficult to visualize these features. Decomposing the signal in components with different frequencies is very useful to remove such perturbations.

### 1.2 Sampling

Signals are often modeled as continuous functions. However, arbitrary continuous functions cannot be manipulated in a computer. In order to be able to process signals digitally, we need to represent them with a finite number of parameters. For smooth signals, which vary slowly in time or space, an intuitive way to achieve this is by *sampling*. Sampling consists of measuring the value of the signal on a discrete grid of points and then using the measurements as a finite-dimensional representation. A crucial question is to what extent the samples preserve the information in the signal. The celebrated sampling theorem provides an answer building upon frequency representations.

### 1.3 Signal denoising

In signal processing, unwanted perturbations occurring during acquisition, storage, transmission or processing are known as *noise*. The goal of denoising is to recover a clean signal– represented by an $N$-dimensional vector $\vec{x} \in \mathbb{R}^N$– from a vector of perturbed data $\vec{y} \in \mathbb{R}^N$. This is a regression function, where the aim is to obtain a function mapping from multidimensional features (the noisy signal) to a multidimensional response (the clean signal).

[1]The image is borrowed from https://commons.wikimedia.org/wiki/File:SinusRhythmLabels.svg

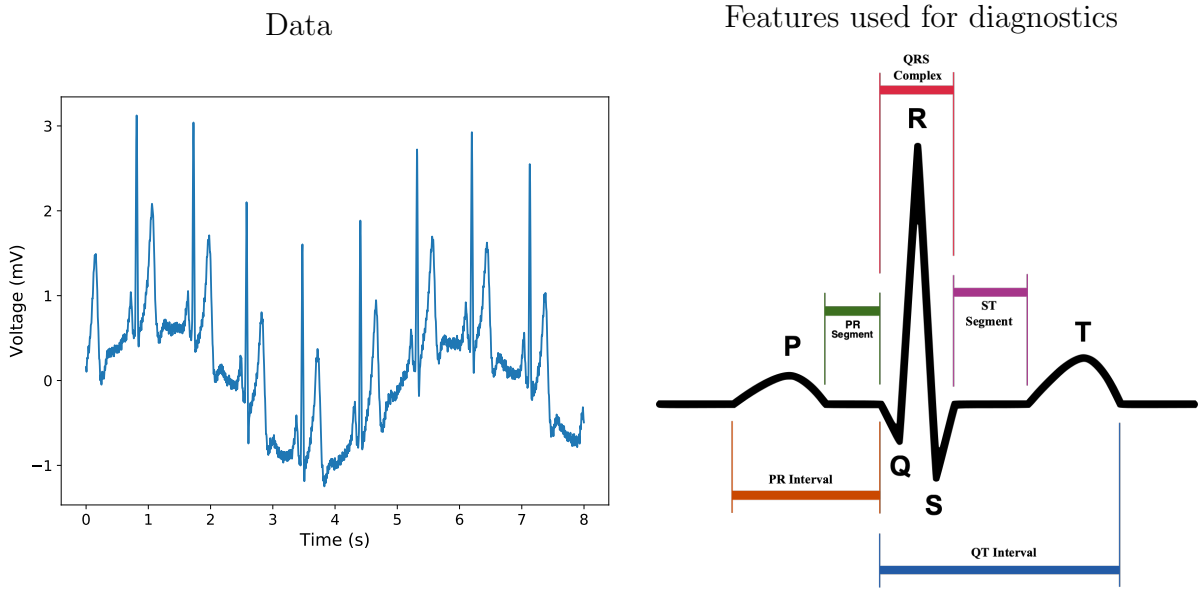


Figure 1: The left image shows an electrocardiogram (ECG) with a duration of 8 s. The top right image shows the features that are relevant for medical diagnostics in a typical ECG.

Let us consider a linear model, where the signal is estimated as a linear transformation of the data parametrized by a matrix $A \in \mathbb{R}^{N \times N}$,

$$\vec{x} \approx A\vec{y}. \tag{1}$$

The linear model has $N^2$ parameters. This is problematic: $N^2$ is *huge* for many types of signals. For sequences of audio of a few seconds, $N \approx 10^5$ (as we will see, audio is sampled at rates of more than 40 kHz). For images, $N$ is the number of pixels, i.e. at least $10^4$. If $N = 10^4$, the linear model has $10^8$ parameters. At the very least, we would need need $N$ examples (of dimension $N$ each), i.e. 100 million images or sound clips (!), to fit the model. Learning an arbitrary linear model is therefore completely impractical, even without taking into account the cost in terms of computational and memory resources. In these notes we will show how to exploit the translation invariance of images and other signals to learn tractable linear models for denoising.

# 2 The frequency domain

## 2.1 The Fourier series

In order to analyze and process signals, it is useful to represent them as linear combinations of simpler functions. Let us consider the Hilbert space of complex-valued square-integrable functions defined on an interval $[a, b] \subset \mathbb{R}$ with the standard inner product,

$$\langle x, y \rangle := \int_a^b x(t)\, \overline{y(t)}\, \mathrm{d}t. \tag{2}$$

To express a function $x : [a, b] \to \mathbb{R}$ in terms of an orthonormal basis of functions $g_1$, $g_2$ ..., each defined from $[a, b]$ to $\mathbb{R}$ all we need to do is compute the coefficients

$$c[j] := \langle x, g_j \rangle . \tag{3}$$

Indeed, by the properties of orthonormal bases,

$$x = \sum_j c[j] g_j. \tag{4}$$

In this section, we describe a representation that decomposes a signal into components that oscillate at different *frequencies*. This often uncovers structure that is otherwise difficult to interpret or manipulate.

The frequency of a periodic signal is the inverse of its period. By far the most popular basis functions with fixed frequencies are sinusoids. Sinusoids are smooth oscillating functions of the form

$$a \cos(2\pi f t + \theta), \tag{5}$$

where $a \in \mathbb{R}$ is the amplitude, $f \in \mathbb{R}$ the frequency, $t$ is the time index, and $\theta \in [0, 2\pi]$ is the phase. A sinusoid with frequency $f$ is periodic with period $T := 1/f$. Expressing a signal in terms of the cosine (or sine) function as in Eq. (5) is not very convenient. The function depends nonlinearly on the phase $\theta$, so there are infinite possible basis functions associated to each frequency. An elegant solution is to instead use complex-valued sinusoids as basis functions.

**Definition 2.1** (Complex sinusoid)**.** *The complex sinusoid with frequency $f \in R$ is given by*

$$\exp(i2\pi f t) := \cos(2\pi f t) + i \sin(2\pi f t). \tag{6}$$

Figure 2 shows a complex sinusoid, along with its real and imaginary parts. Note that complex sinusoids can have negative frequencies. This just changes their imaginary component to $-\sin(2\pi f t)$, instead of $\sin(2\pi f t)$. Any real sinusoid with frequency $f$ can be represented as the sum of two complex sinusoids with frequencies $f$ and $-f$ respectively:

$$\cos(2\pi f t + \theta) = \frac{\exp(i2\pi f t + i\theta) + \exp(-i2\pi f t - i\theta)}{2} \tag{7}$$

$$= \frac{\exp(i\theta)}{2} \exp(i2\pi f t) + \frac{\exp(-i\theta)}{2} \exp(-i2\pi f t). \tag{8}$$

Crucially, the phase is now encoded in the complex amplitude of the sinusoid. As a result, from a linear-algebra perspective, the subspace spanned by the two complex sinusoids with frequencies $f$ and $-f$ contains all possible real sinusoids with frequency $f$. In particular, if we add two sinusoids with the same frequency, but different amplitudes and phases, the result is a sinusoid with that frequency. It therefore makes sense to interpret sinusoids as basis functions, each representing a particular frequency.

If we are interested in obtaining a representation for functions restricted to an interval, working with orthogonal basis functions makes life much easier. The following lemma shows that for any fixed positive $T \in \mathbb{R}$, complex sinusoids with frequency equal to $k/T$– where $k$ is an integer– are all orthogonal.

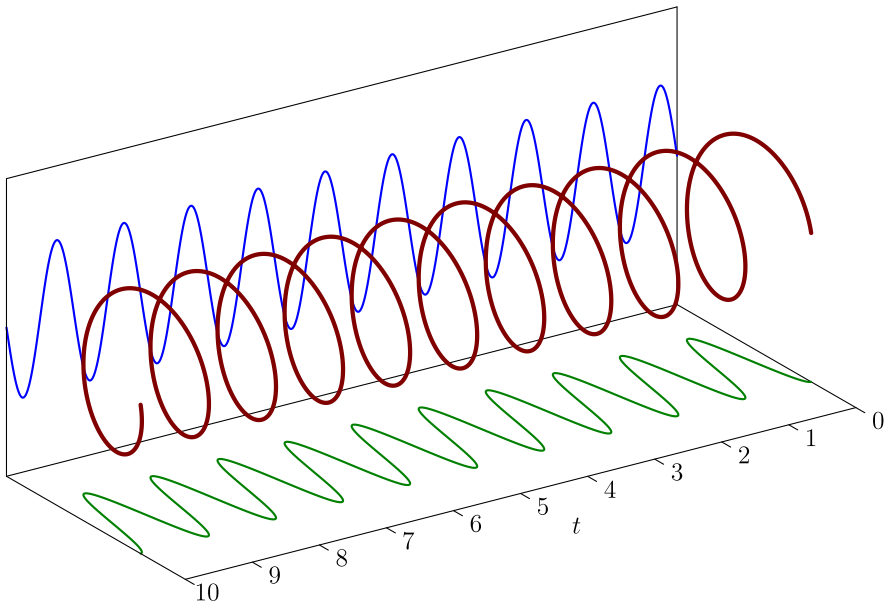


Figure 2: Complex sinusoid (dark red) as a function of time. The real part (green) is a cosine function. The imaginary part (blue) is a sine function.

**Lemma 2.2** (Orthogonality of complex sinusoids)**.** *The family of complex sinusoids with integer frequencies*

$$\phi_k(t) := \exp\left(\frac{i2\pi kt}{T}\right), \qquad k \in \mathbb{Z}, \tag{9}$$

*is an orthogonal set of functions on any interval of the form* $[a, a+T]$*, where* $a, T \in \mathbb{R}$ *and* $T > 0$*.*

*Proof.* We have

$$\begin{aligned}
\langle \phi_k, \phi_j \rangle &= \int_a^{a+T} \phi_k(t)\, \overline{\phi_j(t)}\, \mathrm{d}t && (10)\\
&= \int_a^{a+T} \exp\left(\frac{i2\pi\,(k-j)\,t}{T}\right) \mathrm{d}t && (11)\\
&= \frac{T}{i2\pi\,(k-j)} \left(\exp\left(\frac{i2\pi\,(k-j)\,(a+T)}{T}\right) - \exp\left(\frac{i2\pi\,(k-j)\,a}{T}\right)\right) && (12)\\
&= 0 && (13)
\end{aligned}$$

as long as $j \neq k$. □

In words, this family includes all complex sinusoids with positive and negative frequencies whose period is an integer fraction of the length of the interval. The Fourier series is a decomposition of signals as a sum of these basis functions.

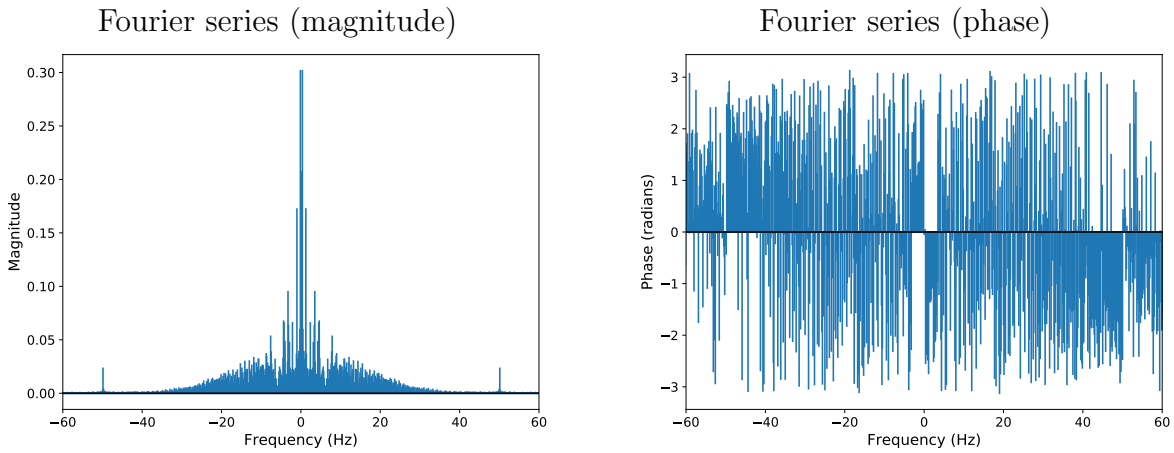


Figure 3: Magnitude (left) and phase (right) of the Fourier coefficients of the ECG data in the left image of Figure 1.

**Definition 2.3** (Fourier series)**.** *The Fourier series coefficients of a function* $x \in \mathcal{L}_2\left[a, a+T\right]$*,* $a, T \in \mathbb{R}$*,* $T > 0$*, are given by*

$$\hat{x}[k] := \langle x, \phi_k \rangle = \int_a^{a+T} x(t) \exp\left(-\frac{i2\pi kt}{T}\right) dt. \tag{14}$$

*The Fourier series of order* $k_c$ *is defined as*

$$\mathcal{F}_{k_c}\left\{x\right\} := \frac{1}{T} \sum_{k=-k_c}^{k_c} \hat{x}[k]\phi_k. \tag{15}$$

*The Fourier series of a signal* $x$ *is defined as* $\lim_{k_c \to \infty} \mathcal{F}_{k_c}\left\{x\right\}$*.*

The basis functions have norm $\sqrt{T}$. We can interpret the Fourier series of order $k_c$ as a projection of the function onto the span of the complex sinusoids with frequencies up to $k_c/T$. By Lemma 2.2 the sinusoids scaled by $1/\sqrt{T}$ are orthonormal so

$$\mathcal{P}_{\text{span}(\{\phi_{-k_c}, \phi_{-k_c+1}, \ldots, \phi_{k_c}\})}\, x = \sum_{k=-k_c}^{k_c} \left\langle x, \frac{1}{\sqrt{T}}\phi_k \right\rangle \frac{1}{\sqrt{T}}\phi_k \tag{16}$$

$$= \mathcal{F}_{k_c}\left\{x\right\}. \tag{17}$$

Figure 3 shows the Fourier series coefficients of the electrocardiogram signal in Figure 1. Figure 4 shows the $k$th-order partial Fourier series of the signal for different values of $k_c$, together with the corresponding real sinusoidal components obtained by combining $\hat{x}_{k_c}\phi_{k_c}$ and $\hat{x}_{-k_c}\phi_{-k_c}$. As $k_c$ increases, the approximation improves. Remarkably, if the function is integrable, the approximation eventually converges to the function. We omit the proof of this result, which is beyond the scope of these notes.

**Theorem 2.4** (Convergence of Fourier series)**.** *For any function* $x \in \mathcal{L}_2[0, T)$*, where* $a, T \in \mathbb{R}$ *and* $T > 0$*,*

$$\lim_{k \to \infty} ||x - \mathcal{F}_k\left\{x\right\}||_{\mathcal{L}_2} = 0. \tag{18}$$

By Theorem 2.4 we can represent any square-integrable function defined on an interval by using its Fourier coefficients, which are often known as the *spectrum* of the function. It is worth noting that one can generalize this representation to functions defined on the whole real line by considering real-valued frequencies, which yields the Fourier transform.

**Example 2.5** (Filtering an electrocardiogram)**.** The low-frequency components of the ECG signal in Figure 1 produce significant *baseline wandering*, which are slowly-varying fluctuations typically caused by motion of the recording apparatus or the patient, or by breathing. In addition, the Fourier series coefficients of the data, shown in Figure 3 reveal a strong component at 50 Hz. This is caused by interference from the electric grid, which transmits electricity exactly at that frequency. In order to use the electrocardiogram for diagnostics, we would like to eliminate these effects, since they are not associated to heart activity. We can be achieve this by removing the corresponding Fourier components from the Fourier series representation, an approach known as *filtering* in the signal processing literature. Removing the low-frequency components gets rid of baseline wandering. Removing the component at 50 Hz removes the electric-grid interference. Figure 5 shows that the approach succeeds in highlighting the features related to heart activity. △

## 2.2 Sampling bandlimited functions

Modeling signals as functions defined on continuous intervals makes it challenging to store and process them in a computer. However, as illustrated by Figure 4, signals are often well approximated by a finite number of Fourier coefficients. Signals with this property are said to be *bandlimited*.

**Definition 2.6** (Bandlimited signal)**.** *A signal is bandlimited with a cut-off frequency $k_c$ if it is equal to its Fourier series representation of order $k_c$, i.e.*

$$x(t) = \sum_{k=-k_c}^{k_c} \hat{x}[k] \exp\left(\frac{i2\pi kt}{T}\right). \tag{19}$$

Sampling is a fundamental tool in signal processing, which consists of evaluating a function at a finite number of fixed locations. The hope is that the samples preserve as much information as possible. In the case of bandlimited signals, sampling can in fact capture all of the information in the signal.

Let us consider a uniform sampling pattern with $N$ sampling locations that divide an interval of length $T$ into $N$ segments of length $T/N$. For simplicity, we assume that the interval is $[0, T)$ without loss of generality. We measure a bandimited signal $x$ with cut-off frequency $k_c$ at those locations to obtain the samples: $x\left(\frac{0}{N}\right)$, $x\left(\frac{T}{N}\right)$, $x\left(\frac{2T}{N}\right)$, $\ldots$, $x\left(\frac{(N-1)T}{N}\right)$. The Fourier series of the bandlimited signal provides a linear equation relating the samples and the Fourier coefficients of the signal:

$$x\left(\frac{jT}{N}\right) = \sum_{k=-k_c}^{k_c} \hat{x}_k \exp\left(\frac{i2\pi kjT}{NT}\right) \tag{20}$$

$$= \sum_{k=-k_c}^{k_c} \hat{x}_k \exp\left(\frac{i2\pi kj}{N}\right). \tag{21}$$

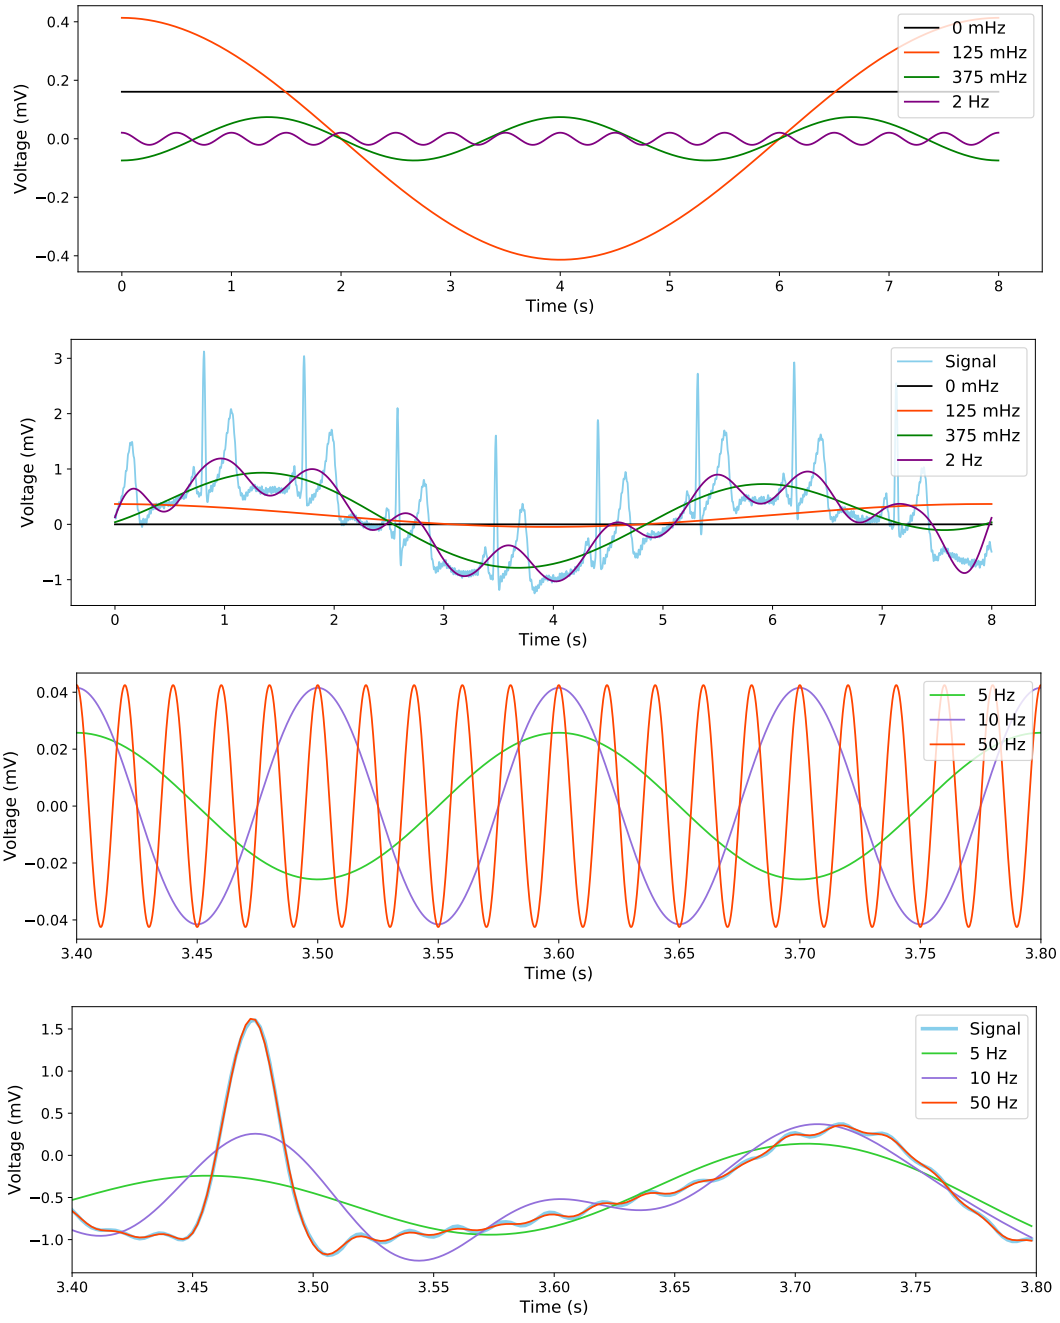


Figure 4: The top image shows components of the Fourier series of the ECG signal from Figure 1. The third image shows zoomed components corresponding to higher frequencies. The second and fourth image show the Fourier series of the signal truncated at different frequencies.

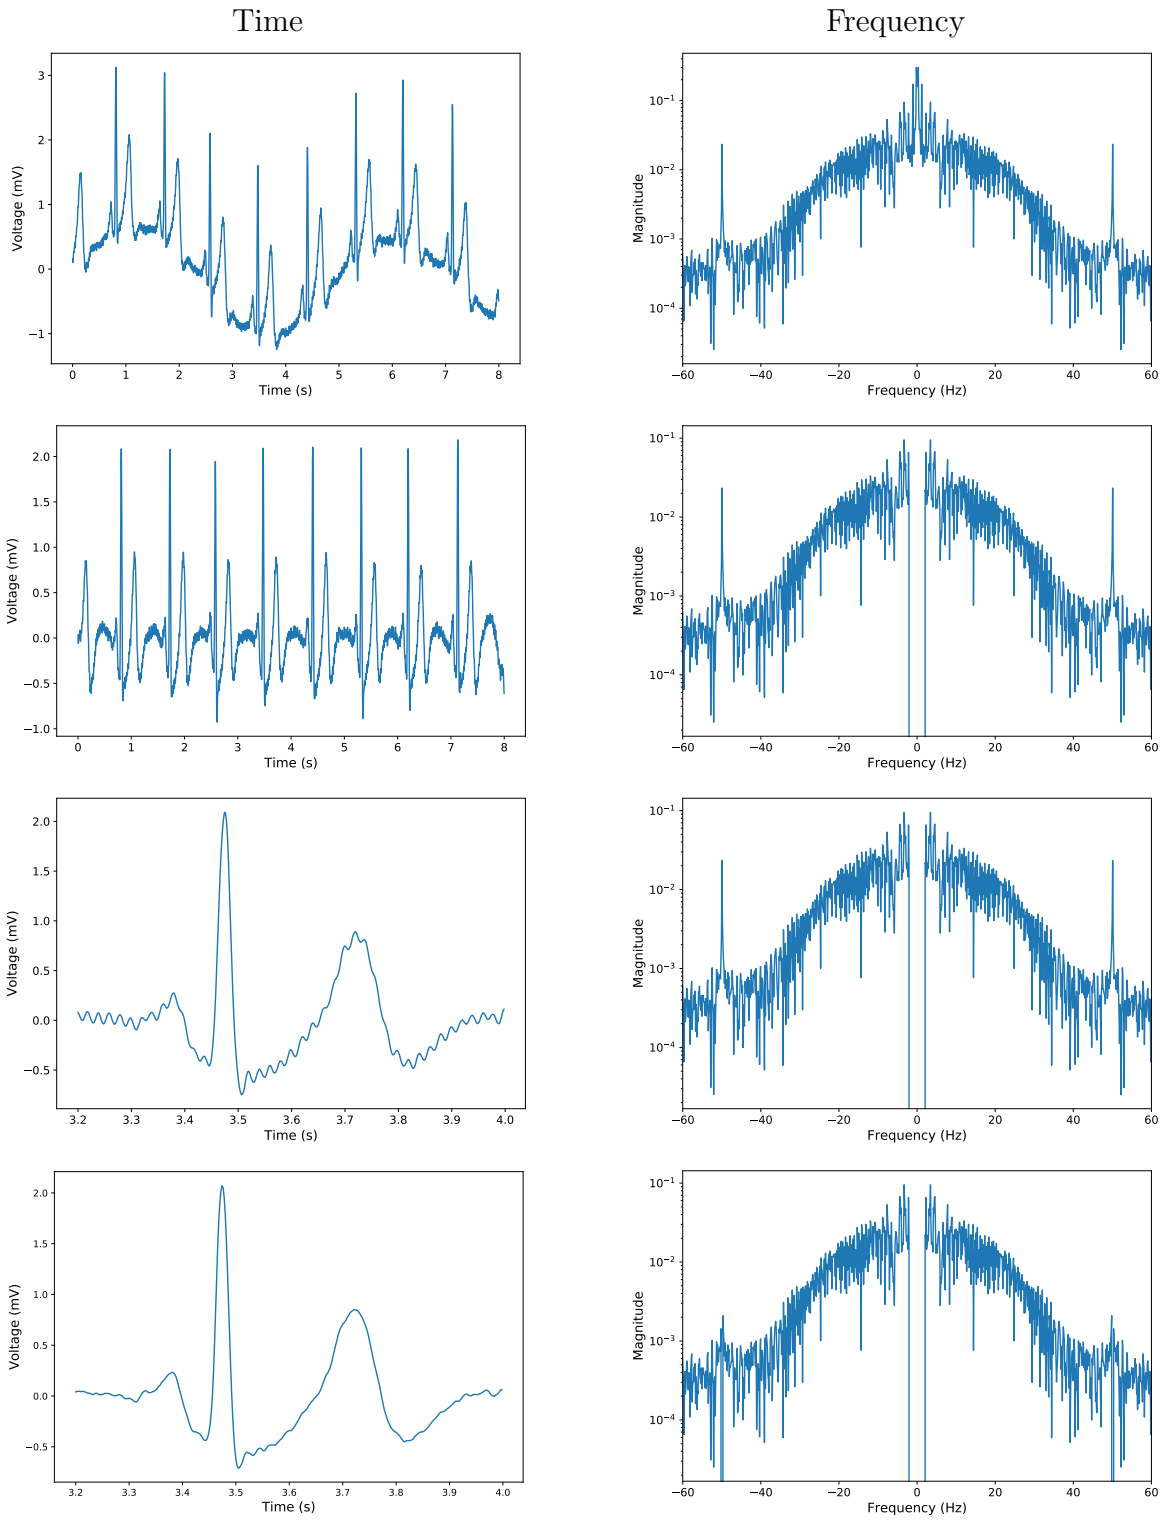


Figure 5: The ECG data from Example 2.5 (top row) is filtered to first remove low-frequency components (second and third rows) and then electric-grid interference at 50 Hz (bottom row). The left column displays the signals in the time domain. The right columns shows the magnitude of the corresponding Fourier series coefficients.

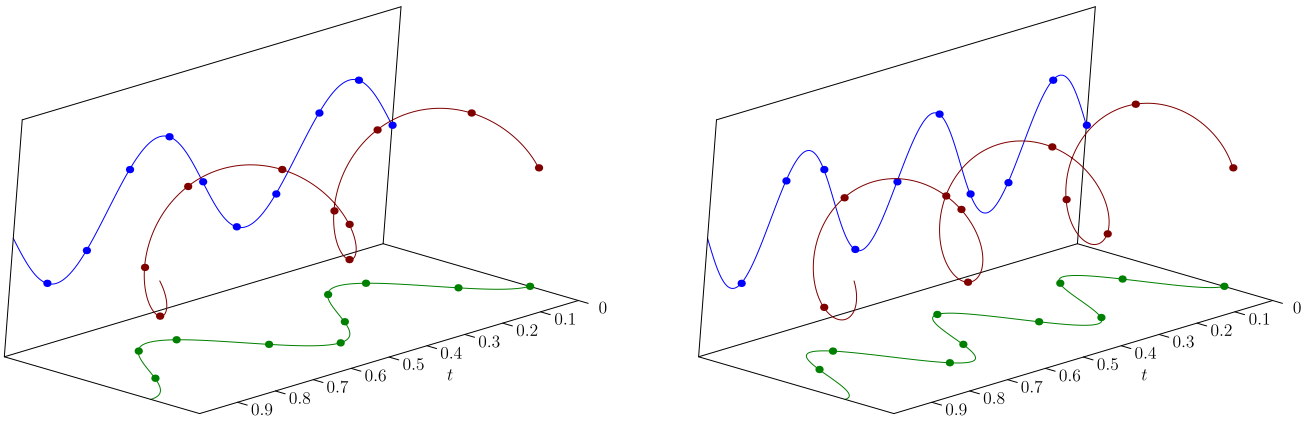


Figure 6: Discrete complex sinusoids $\vec{\phi}_2$ (left) and $\vec{\phi}_3$ (right) for $N = 10$, displayed in red. The real and imaginary parts are shown in green and blue respectively.

This yields the following system of linear equations:

$$\begin{bmatrix} x\left(\frac{0}{N}\right) \\ x\left(\frac{T}{N}\right) \\ \cdots \\ x\left(\frac{jT}{N}\right) \\ \cdots \\ x\left(T-\frac{T}{N}\right) \end{bmatrix} = \begin{bmatrix} 1 & 1 & \cdots & 1 \\ \exp\left(\frac{i2\pi(-k_c)}{N}\right) & \exp\left(\frac{i2\pi(-k_c+1)}{N}\right) & \cdots & \exp\left(\frac{i2\pi k_c}{N}\right) \\ \cdots & \cdots & \cdots & \cdots \\ \exp\left(\frac{i2\pi(-k_c)j}{N}\right) & \exp\left(\frac{i2\pi(-k_c+1)j}{N}\right) & \cdots & \exp\left(\frac{i2\pi k_c j}{N}\right) \\ \cdots & \cdots & \cdots & \cdots \\ \exp\left(\frac{i2\pi(-k_c)(N-1)}{N}\right) & \exp\left(\frac{i2\pi(-k_c+1)(N-1)}{N}\right) & \cdots & \exp\left(\frac{i2\pi k_c(N-1)}{N}\right) \end{bmatrix} \begin{bmatrix} \hat{x}[-k_c] \\ \hat{x}[-k_c+1] \\ \cdots \\ \hat{x}[k_c] \end{bmatrix}$$

More succinctly, we have

$$\vec{x}_{[N]} = \widetilde{F}_{[N]} \hat{x}_{[k_c]}, \tag{22}$$

where $\vec{x}_{[N]}$ is the vector of samples and $\hat{x}_{[k_c]}$ is the vector of Fourier coefficients. If the matrix $\widetilde{F}_{[N]}$ is invertible, we can recover the Fourier coefficients from the samples and completely recover the bandlimited function. The columns of $\widetilde{F}_{[N]}$ are discretized complex sinusoids.

**Definition 2.7** (Discrete complex sinusoids)**.** *The discrete complex sinusoid $\vec{\phi}_k \in \mathbb{C}^N$ with integer frequency $k$ is defined as*

$$\vec{\phi}_k[j] := \exp\left(\frac{i2\pi kj}{N}\right), \qquad 0 \leq j, k \leq N-1. \tag{23}$$

Figure 6 shows $\vec{\phi}_2$ and $\vec{\phi}_3$ for $N := 10$.

Discrete complex sinusoids form an orthonormal basis of $\mathbb{C}^N$ if we scale them by $1/\sqrt{N}$.

**Lemma 2.8** (Orthonormal sinusoidal basis)**.** *The discrete complex exponentials $\frac{1}{\sqrt{N}}\vec{\phi}_0, \ldots, \frac{1}{\sqrt{N}}\vec{\phi}_{N-1}$ form an orthonormal basis of $\mathbb{C}^N$.*

*Proof.* Each vector has $\ell_2$ norm equal to $\sqrt{N}$,

$$\left|\left|\vec{\phi}_k\right|\right|_2^2 = \sum_{j=0}^{N-1} \left|\vec{\phi}_k\left[j\right]\right|^2 \tag{24}$$

$$= \sum_{j=0}^{N-1} 1 \tag{25}$$

$$= N, \tag{26}$$

and

$$\left\langle \vec{\phi}_k, \vec{\phi}_l \right\rangle = \sum_{j=0}^{N-1} \vec{\phi}_k\left[j\right] \overline{\vec{\phi}_l\left[j\right]} \tag{27}$$

$$= \sum_{j=0}^{N-1} \exp\left(\frac{i2\pi\left(k-l\right)j}{N}\right) \tag{28}$$

$$= \frac{1-\exp\left(\frac{i2\pi(k-l)N}{N}\right)}{1-\exp\left(\frac{i2\pi(k-l)}{N}\right)} \tag{29}$$

$$= 0, \tag{30}$$

if $k \neq l$. Since there are $N$ vectors in the set and they are linearly independent, they form a basis of $\mathbb{C}^N$. □

An immediate consequence is that bandlimited signals with cut-off frequency $k_c$ can be recovered from just $2k_c + 1$ samples.

**Theorem 2.9** (Nyquist-Shannon-Kotelnikov sampling theorem)**.** *Any bandlimited signal $x \in \mathcal{L}_2[0,T)$, where $T > 0$, with cut-off frequency $k_c$ can be recovered exactly from $N$ uniformly spaced samples $x\left(0\right)$, $x\left(T/N\right)$, $\ldots$, $x\left(T - T/N\right)$ as long as*

$$N \geq 2k_c + 1, \tag{31}$$

*where $2k_c + 1$ is known as the* Nyquist rate*. The Fourier series coefficients $\hat{x}_{[k_c]}$ can be recovered from the vector of samples $\vec{x}_{[N]}$ as follows*

$$\hat{x}_{[k_c]} = \frac{1}{N}\widetilde{F}^*_{[N]}\vec{x}_{[N]}, \tag{32}$$

*where $\widetilde{F}_{[N]}$ is defined in Eq. (22).*

*Proof.* Note that for $-k_c \leq k \leq -1$ and $0 \leq j \leq N-1$,

$$\exp\left(\frac{i2\pi kj}{N}\right) = \exp\left(\frac{i2\pi\left(N+k\right)j}{N}\right), \tag{33}$$

so the columns of matrix $\widetilde{F}_{[N]}$ in Eq. (22) equal $\vec{\phi}_{N-k_c}$, ..., $\vec{\phi}_{N-1}$, $\vec{\phi}_0$, ..., $\vec{\phi}_{k_c}$. If $N \geq 2k_c + 1$, then no column is repeated. In that case, by Lemma 3.8 the matrix $\widetilde{F}_{[N]}$ has orthogonal columns with norm $\sqrt{N}$, which implies

$$\frac{1}{N}\widetilde{F}^*_{[N]}\widetilde{F}_{[N]} = \left(\frac{1}{\sqrt{N}}\widetilde{F}_{[N]}\right)^* \frac{1}{\sqrt{N}}\widetilde{F}_{[N]} = I. \tag{34}$$

□

The range of frequencies that human beings can hear is from 20 Hz to 20 kHz. The sampling theorem therefore dictates that audio should be sampled at a rate of at least 40 kHz. Typical rates used in practice are 44.1 kHz (CD), 48 kHz, 88.2 kHz, or 96 kHz.

In practice, signals are never exactly bandlimited. It is therefore very important to understand what happens if we sample a signal that is not bandlimited and try to reconstruct it under the assumption that it is. In general, from $N$ samples, we can only hope to estimate the first $2k_c + 1$ Fourier coefficients, where $k_c$ is such that $N := 2k_c+1$. The following lemma establishes that, unless the signal is exactly bandlimited with cut-off frequency $k_c$, the estimated Fourier coefficients are *not* equal to the true coefficients because there is additive interference from coefficients corresponding to higher frequencies. This phenomenon is known as *aliasing*.

**Lemma 2.10** (Aliasing)**.** *Let $x$ be a defined on $[0, T)$, $T > 0$, and let $\vec{x}_{[N]}$ be a vector of $N$ samples of $x$ at 0, $T/N$, $2T/N$, ... $T - T/N$. The vector*

$$\hat{x}^{\text{rec}}[k] := \frac{1}{N}(\widetilde{F}^*_{[N]}\vec{x}_{[N]})[k], \quad -k_{\text{samp}} \leq k \leq k_{\text{samp}}, \tag{35}$$

*denotes the estimate of the $k$th Fourier coefficient of $x$ computed under the assumption that the signal is bandlimited with a cut-off frequency of $k_{\text{samp}}$, where $N = 2k_{\text{samp}} + 1$. We have*

$$\hat{x}^{\text{rec}}[k] = \sum_{\{(m-k) \bmod N = 0\}} \hat{x}[m], \tag{36}$$

*where $\hat{x}$ denotes the true Fourier coefficients of $x$.*

*Proof.* We denote the true cut-off frequency of $x$, which can be arbitrarily large by $k_{\text{true}}$ (the argument can be extended to non-bandlimited funcions). By the definition of $\widetilde{F}_{[N]}$

$$\frac{1}{N}(\widetilde{F}^*_{[N]}\vec{x}_{[N]})[k] = \frac{1}{N}\sum_{j=0}^{N-1} \exp\left(-\frac{i2\pi kj}{N}\right) \sum_{m=-k_{\text{true}}}^{k_{\text{true}}} \hat{x}[m] \exp\left(\frac{i2\pi mj}{N}\right) \tag{37}$$

$$= \frac{1}{N}\sum_{m=-k_{\text{true}}}^{k_{\text{true}}} \hat{x}[m] \sum_{j=0}^{N-1} \exp\left(\frac{i2\pi(m-k)j}{N}\right). \tag{38}$$

By the argument in Eq. (29), the geometric series

$$\sum_{j=0}^{N-1} \exp\left(\frac{i2\pi(m-k)j}{N}\right) \tag{39}$$

equals $N$ if $(m - k) \bmod N = 0$ and zero otherwise. This completes the proof. □

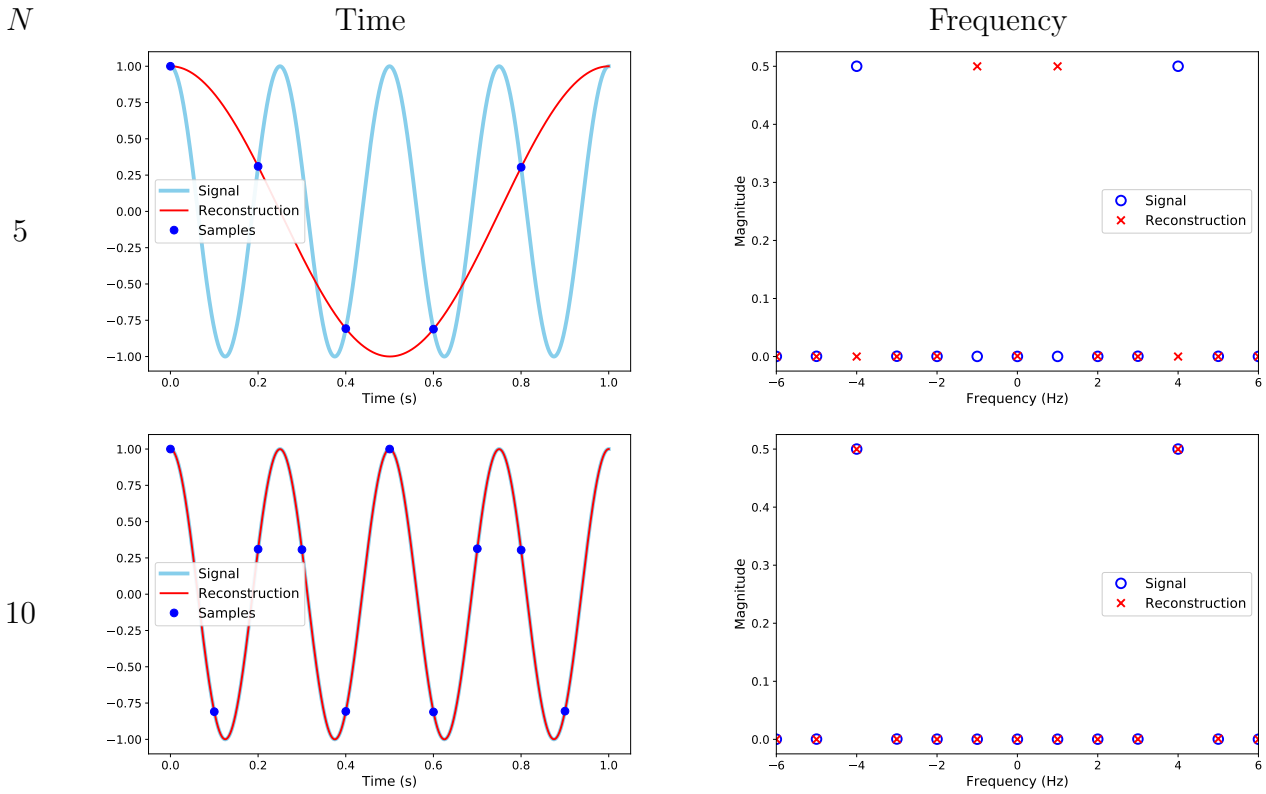


Figure 7: Illustration of aliasing (top row) and perfect recovery (bottom row) of a real sinusoid from equispaced samples, as described in Example 2.11. The left column shows the samples, the original signal and the recovered signal in the time domain. The right column shows the corresponding Fourier coefficients.

The reason for aliasing is that samples from a sinusoid with frequency $m$ are identical to the samples from a sinusoid with frequency $k = m \mod N$, where $N$ is the number of samples. This is illustrated in the following example and Figure 7. The sampling theorem establishes that if the only sinusoidal components present in a signal all have frequencies less than the Nyquist rate, then there is no possible confusion.

**Example 2.11** (Sampling a real sinusoid)**.** Let us consider a real sinusoid with frequency equal to 4 Hz

$$x(t) := \cos(8\pi t) \tag{40}$$

$$= 0.5 \exp(-i2\pi 4t) + 0.5 \exp(i2\pi 4t) \tag{41}$$

measured over one second, i.e. $T = 1$ s. The Fourier coefficients of the original signal are all zero except for $\hat{x}[-4] = 0.5$ and $\hat{x}[4] = 0.5$, so the signal is bandlimited with cut-off frequency $k_c := 4$. If the number of samples $N$ is larger than $2k_c + 1 = 9$ then by Theorem 3.6 the Fourier coefficients, and hence the signal, can be reconstructed perfectly, as shown in Figure 7 for $N = 10$.

To illustrate what happens if $N < 2k_c + 1$ let us assume $N = 5$. In that case we can only hope to recover 5 Fourier coefficients: $k \in \{-2, -1, 0, 1, 2\}$. Applying Lemma 2.10, the Fourier coefficients corresponding to $k \in \{-2, 0, 2\}$ will be estimated to equal zero because $(4 - k) \bmod 5 \neq 0$ and $(-4-k) \bmod 5 \neq 0$ for $k = -2, 0, 2$. For $k = -1$ we have $(4-(-1)) \bmod N = 0$ so the corresponding estimated coefficient $\hat{x}^{\text{rec}}[-1]$ equals $\hat{x}_4 = 0.5$. Similarly, $(-4 - 1) \bmod N = 0$ so $\hat{x}^{\text{rec}}[1]$ equals $\hat{x}[-4] = 0.5$. The samples are consistent with a real sinusoid with cut-off frequency 1 Hz, as shown in Figure 7. △

**Example 2.12** (Sampling an electrocardiogram)**.** The ECG signal from Example 2.5 is only approximately bandlimited. Its approximate cut-off frequency is roughly above 50 Hz, where there is a significant contribution from electric-grid interference. In this example, we have $T = 8$ s, so 50 Hz corresponds to $k_c = 50/(1/T) = 400$. By Theorem 3.6 the number of samples should be at least 801 to avoid significant aliasing. Figure 8 shows that for $N = 1,000$ the aliasing is indeed almost imperceptible. In contrast, if $N = 625$ the aliasing is evident in both the time and the frequency domain. In particular the frequency component at $k = 400$ (corresponding to the electric-grid interference at 50 Hz) shows up at $k = \pm 225$, which corresponds to 28.125 Hz. This follows from Lemma 2.10 because $(400 - (-225)) \bmod 625 = 0$ and $(-400 - 225) \bmod 625 = 0$. △

## 2.3 The discrete Fourier transform

The discrete Fourier transform is a change of basis that expresses a finite-dimensional vector in terms of discrete complex sinusoids. It is the discrete counterpart of the Fourier series.

**Definition 2.13** (Discrete Fourier transform)**.** *The discrete Fourier transform (DFT) of a vector*

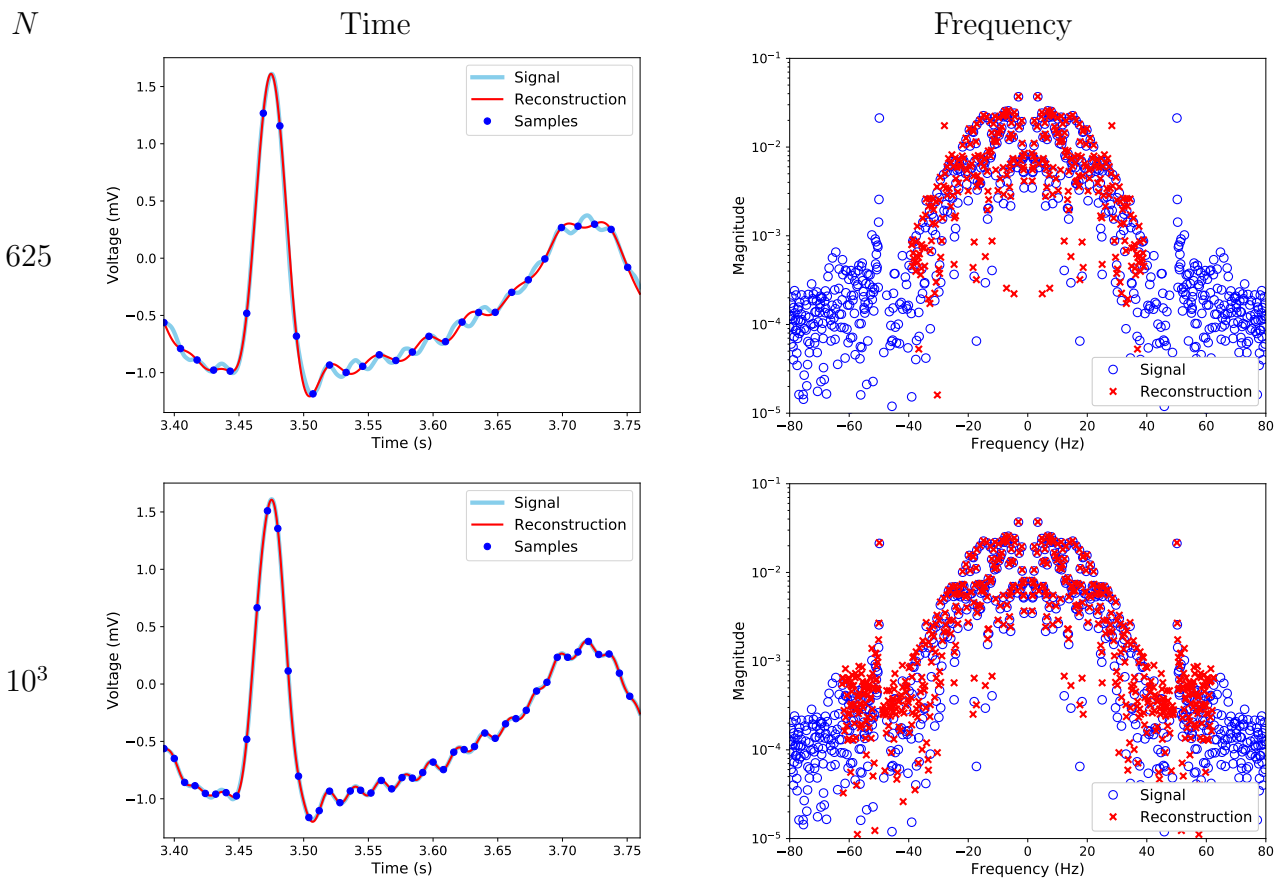


Figure 8: Effect of different sampling rates on the recovery of the ECG data from Example 2.5. The left column shows the samples, the original signal and the recovered signal in the time domain. The right column shows the corresponding Fourier coefficients. As explained in Example 2.12, if the sampling rate is not sufficiently high, there is significant aliasing (top row). This can be mitigated by increasing the rate (bottom row).

$\vec{x} \in \mathbb{C}^N$ *is given by*

$$\hat{x} := \begin{bmatrix} 1 & 1 & 1 & \cdots & 1 \\ 1 & \exp\left(-\frac{i2\pi}{N}\right) & \exp\left(-\frac{i2\pi 2}{N}\right) & \cdots & \exp\left(-\frac{i2\pi(N-1)}{N}\right) \\ 1 & \exp\left(-\frac{i2\pi 2}{N}\right) & \exp\left(-\frac{i2\pi 4}{N}\right) & \cdots & \exp\left(-\frac{i2\pi 2(N-1)}{N}\right) \\ \cdots & \cdots & \cdots & \cdots & \cdots \\ 1 & \exp\left(-\frac{i2\pi(N-1)}{N}\right) & \exp\left(-\frac{i2\pi 2(N-1)}{N}\right) & \cdots & \exp\left(-\frac{i2\pi(N-1)^2}{N}\right) \end{bmatrix} \vec{x} = F_{[N]}\vec{x}. \tag{42}$$

*In terms of the discrete complex sinusoids in Definition 2.7,*

$$\hat{x}\left[k\right] = \left\langle \vec{x}, \vec{\phi}_k \right\rangle, \qquad 0 \leq k \leq N-1. \tag{43}$$

*The inverse DFT of a vector* $\hat{y} \in \mathbb{C}^N$ *equals*

$$\vec{y} = \frac{1}{N} F_{[N]}^* \hat{y}. \tag{44}$$

The rows of the DFT matrix $F_{[N]}$ are exactly the same as the rows of the matrix $\widetilde{F}_{[N]}^*$ in Eq. (22), only in a different order (since for any $k$ and any integer $l$ $\phi_k = \phi_{k+lN}$). This observation combined with Theorem 3.6 provides a interesting interpretation of the DFT: if $\vec{x} \in \mathbb{C}^N$ contains equispaced samples from a bandlimited signal $x$ with cut-off frequency $k_c$– where $N \geq 2k_c + 1$– then the DFT of $\vec{x}$ contains the nonzero Fourier series coefficients of $x$. In addition, by Lemma 3.8 the rows (and columns) of $F_{[N]}$ are orthogonal and have norm $\sqrt{N}$, which justifies the definition of the inverse DFT.

**Corollary 2.14.** *The inverse DFT inverts the DFT.*

In general, the time complexity of multiplying an $N \times N$ matrix with an $N$-dimensional vector is $N^2$. However one can compute the DFT much faster. The fast-Fourier transform (FFT) algorithm exploits the structure of the DFT matrix to compute the DFT with complexity $\mathcal{O}\left(N \log N\right)$. It is difficult to overstate the importance of the FFT. Gilbert Strang has described it as *the most important numerical algorithm of our lifetime*. The main insight underlying the FFT is that the $N$-order DFT matrix can be expressed in terms of $N/2$-order DFT submatrices. To simplify the exposition we assume that $N$ is even, but similar decompositions are possible for odd $N$.

**Lemma 2.15.** *Let* $F_{[N]}$ *denote the* $N \times N$ *DFT matrix, where* $N$ *is even. Then for* $k = 0, 1, \ldots, N/2 - 1$*, and any vector* $\vec{x} \in \mathbb{C}^N$

$$F_{[N]}\vec{x}\left[k\right] = F_{[N/2]}\vec{x}_{\text{even}}\left[k\right] + \exp\left(-\frac{i2\pi k}{N}\right) F_{[N/2]}\vec{x}_{\text{odd}}\left[k\right], \tag{45}$$

$$F_{[N]}\vec{x}\left[k + N/2\right] = F_{[N/2]}\vec{x}_{\text{even}}\left[k\right] - \exp\left(-\frac{i2\pi k}{N}\right) F_{[N/2]}\vec{x}_{\text{odd}}\left[k\right], \tag{46}$$

*where* $\vec{x}_{\text{even}}$ *and* $\vec{x}_{\text{odd}}$ *contain the even and odd entries of* $\vec{x}$ *respectively.*

*Proof.* The proof is illustrated in Figure 9. □

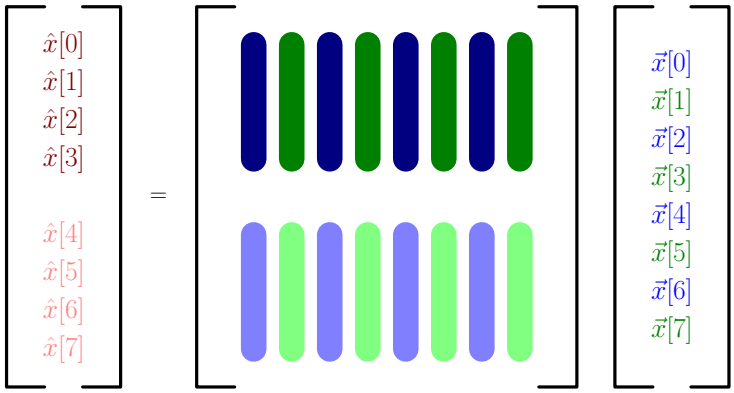


(a) Separation in even/odd columns and top/bottom rows.

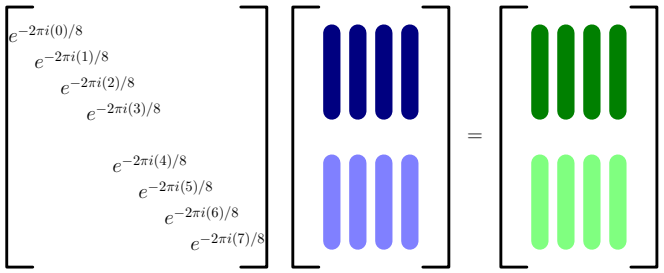


(b) Even columns can be scaled to yield odd columns.

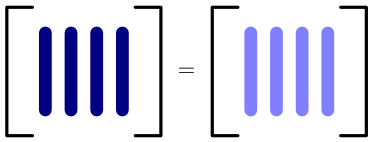


(c) Top even submatrix and bottom even submatrix are both an $N/2$-order DFT matrix.

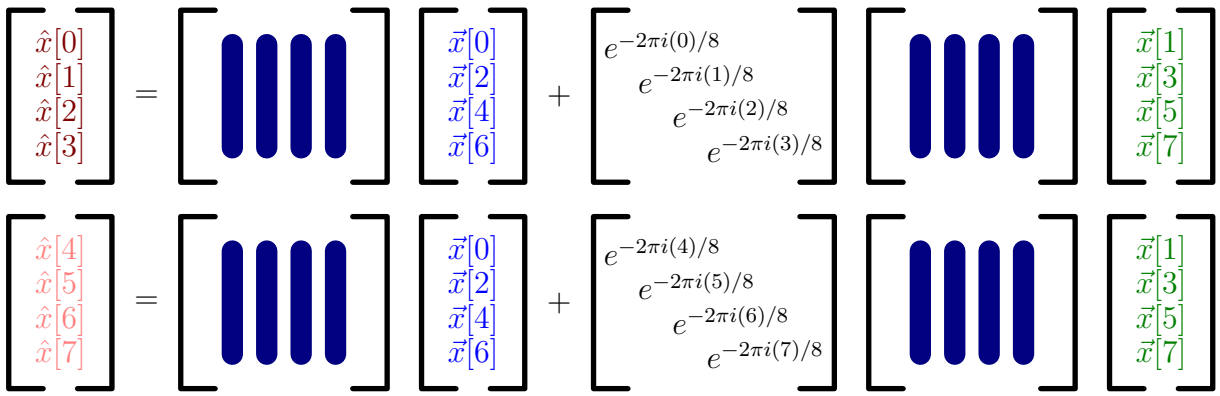


(d) Application of $N$-order DFT matrix to an $N$-dimensional vector expressed in terms of two $N/2$-order DFT matrices applied to two $N/2$-dimensional vectors.

Figure 9: Illustration of the proof of Lemma 2.15 for $N := 8$.

The Cooley-Tukey FFT algorithm exploits Lemma 2.15 to compute the DFT recursively.

**Algorithm 2.16** (Cooley-Tukey Fast Fourier transform)**.** *If $N = 1$, output $F_{[1]}\vec{x} := \vec{x}$. Otherwise apply the following steps:*

1. *Compute $F_{[N/2]}\vec{x}_{\text{even}}$.*
2. *Compute $F_{[N/2]}\vec{x}_{\text{odd}}$.*
3. *For $k = 0, 1, \ldots, N/2 - 1$ set*

$$F_{[N]}\vec{x}\,[k] := F_{[N/2]}\vec{x}_{\text{even}}\,[k] + \exp\left(-\frac{i2\pi k}{N}\right) F_{[N/2]}\vec{x}_{\text{odd}}\,[k]\,, \tag{47}$$

$$F_{[N]}\vec{x}\,[k + N/2] := F_{[N/2]}\vec{x}_{\text{even}}\,[k] - \exp\left(-\frac{i2\pi k}{N}\right) F_{[N/2]}\vec{x}_{\text{odd}}\,[k]\,. \tag{48}$$

**Lemma 2.17** (Complexity of the FFT)**.** *The FFT algorithm has complexity $O(N \log N)$.*

*Proof.* For simplicity, we assume that $N = 2^L$ for some positive integer $L$; the argument can be adapted to the general case. The algorithm recursively decomposes the task of computing the DFT into a tree with $L$ levels depicted in Figure 10. Let us number the levels starting at 0, which represents the root. At level $l \in \{1, \ldots, L\}$ there are $2^l$ nodes. At each node we need to compute Eq. (47), which amounts to scaling a vector of length $2^{L-l}$ and adding it to another vector of the same length. As a result at each level the complexity is of order $2^l 2^{L-l} = 2^L = N$. Since there are exactly $\log_2 N$ levels, the proof is complete. □

Figure 11 shows a numerical comparison of the running time of the FFT and a naive matrix-based implementation of the DFT. The difference is dramatic, as expected from Lemma 2.17.

# 3 Frequency representations in multiple dimensions

## 3.1 Fourier series of multidimensional signals

In this section we discuss the extension of frequency representations to multidimensional spaces. We consider the Hilbert space of complex-valued square-integrable functions defined on a hyper-rectangle $\mathcal{I} := [a_1, b_1] \times \ldots \times [a_p, b_p] \subset \mathbb{R}^p$ with the standard inner product,

$$\langle x, y \rangle := \int_{\mathcal{I}} x\left(\vec{t}\right) \overline{y\left(\vec{t}\right)}\, \mathrm{d}\vec{t}. \tag{49}$$

A natural multidimensional generalization of the sinusoid in Eq. (5) is

$$a \cos\left(2\pi \langle \vec{f}, \vec{t}\,\rangle + \theta\right). \tag{50}$$

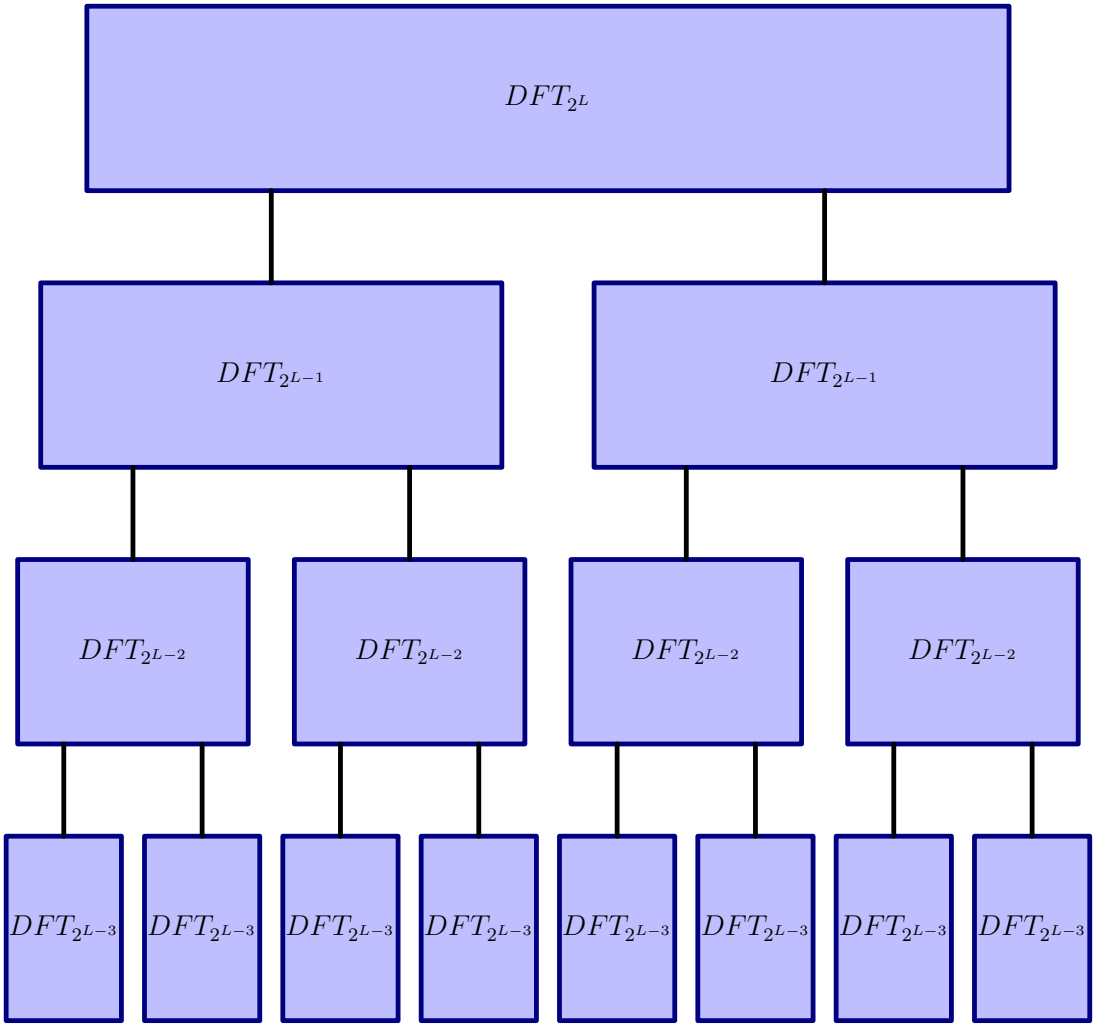


Figure 10: Recursive structure of the FFT algorithm for $N := 8$.

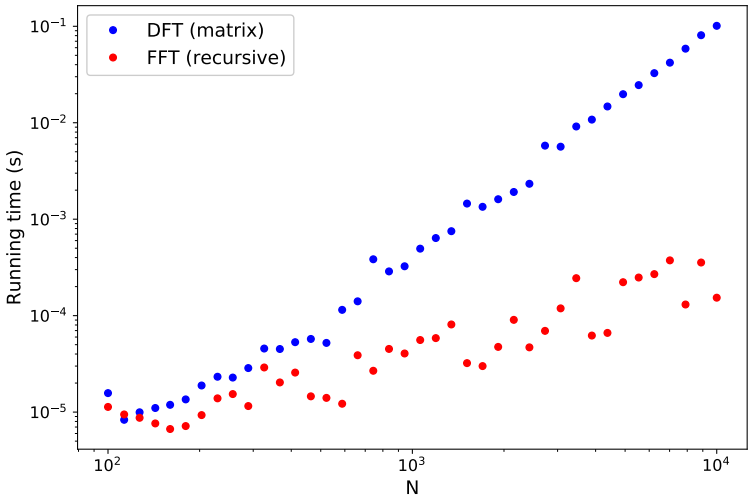


Figure 11: Numerical comparison of the running times of a naive matrix-based implementation of the DFT (assuming the matrix is already built) and the FFT algorithm. Both approaches were applied to 20 random complex vectors for each value of $N$.

Here the frequency and the time variables are $d$-dimensional vectors. The multidimensional sinusoid is periodic with period $1/\left|\left|f\right|\right|_2$ in the direction of the unit vector $\vec{v} := \vec{f}/\left|\left|f\right|\right|_2$. Indeed,

$$a\cos\left(i2\pi\left\langle \vec{f},\vec{t}+\frac{m}{\left|\left|f\right|\right|_2}\frac{\vec{f}}{\left|\left|f\right|\right|_2}\right\rangle+\theta\right)=a\cos\left(i2\pi\langle\vec{f},\vec{t}\rangle+i2\pi m+\theta\right) \tag{51}$$

$$=a\cos\left(i2\pi\langle\vec{f},\vec{t}\,\rangle+\theta\right), \tag{52}$$

for any integer $m$. Multidimensional sinusoids are planar waves, constant along any hyperplane orthogonal to $\vec{f}$. In the direction of $\vec{f}$ they oscillate with a frequency equal to the $\ell_2$ norm of $\vec{f}$. Figure 12 shows some examples. As in one dimension, working with complex sinusoids makes it easier to define frequency representations, because the phase can be encoded linearly as a complex multiplicative coefficient.

**Definition 3.1** (Multidimensional complex sinusoids)**.** *The complex sinusoid with frequency* $\vec{f}\in R^d$ *is given by*

$$\exp(i2\pi\langle\vec{f},\vec{t}\,\rangle):=\cos(2\pi\langle\vec{f},\vec{t}\,\rangle)+i\sin(2\pi\langle\vec{f},\vec{t}\,\rangle), \tag{53}$$

*where* $\vec{t}\in\mathbb{R}^d$.

Any multidimensional sinusoid with frequency $\vec{f}$ can be expressed as the sum of two complex sinusoids with frequencies $\vec{f}$ and $-\vec{f}$ respectively:

$$\cos\left(i2\pi\langle\vec{f},\vec{t}\,\rangle+\theta\right)=\frac{\exp(i\theta)}{2}\exp(i2\pi\langle\vec{f},\vec{t}\,\rangle)+\frac{\exp(-i\theta)}{2}\exp(-i2\pi\langle\vec{f},\vec{t}\,\rangle). \tag{54}$$

Multidimensional complex sinusoids can be expressed as the product of complex sinusoids aligned with each axis.

$$\exp(i2\pi\langle \vec{f}, \vec{t}\,\rangle) := \exp\left(i2\pi \sum_{j=1}^{d} \vec{f}[j]\vec{t}[j]\right) \tag{55}$$

$$= \prod_{j=1}^{d} \exp(i2\pi \vec{f}[j]\vec{t}[j]). \tag{56}$$

For simplicity, and because we are mostly interested in applications related to images, from now on we focus on the two-dimensional case. However, most of the results generalize to higher dimensions. To alleviate notation, we denote each coordinate by $t_1 := \vec{t}[1]$ and $t_2 := \vec{t}[2]$. If turns out that for any fixed positive $T \in \mathbb{R}$, complex sinusoids with frequency coordinates equal to $k_1/T$ and $k_2/T$ for $k_1, k_2 \in \mathbb{Z}$ are all orthogonal. Figure 12 shows real-valued 2D sinusoids corresponding to different values of $k_1$ and $k_2$.

**Lemma 3.2** (Orthogonality of multidimensional complex sinusoids)**.** *The family of complex sinusoids with integer frequencies*

$$\phi^{\text{2D}}_{k_1,k_2}(t_1, t_2) := \exp\left(\frac{i2\pi k_1 t_1}{T}\right) \exp\left(\frac{i2\pi k_2 t_2}{T}\right), \qquad k_1, k_2 \in \mathbb{Z}, \tag{57}$$

*is an orthogonal set of functions on any interval of the form $[a, a+T]\times[b, b+T]$, where $a, b, T \in \mathbb{R}$ and $T > 0$.*

*Proof.* The result follows from Lemma 2.2. By Eq. (56)

$$\phi^{\text{2D}}_{k_1,k_2}(t_1, t_2) = \phi_{k_1}(t_1)\,\phi_{k_2}(t_2), \tag{58}$$

so that

$$\left\langle \phi^{\text{2D}}_{k_1,k_2}, \phi^{\text{2D}}_{j_1,j_2} \right\rangle = \int_{t_1=a}^{a+T} \int_{t_2=b}^{b+T} \phi_{k_1}(t_1)\,\phi_{k_2}(t_2)\,\overline{\phi_{j_1}(t_1)\,\phi_{j_2}(t_2)}\,\mathrm{d}t_1\,\mathrm{d}t_2 \tag{59}$$

$$= \langle \phi_{k_1}, \phi_{j_1} \rangle \langle \phi_{k_2}, \phi_{j_2} \rangle \tag{60}$$

$$= 0 \tag{61}$$

as long as $j_1 \neq k_1$ or $j_2 \neq k_2$. □

The 2D Fourier series represents 2D functions in terms of complex sinusoids.

**Definition 3.3** (Fourier series)**.** *The Fourier series coefficients of a function $x \in \mathcal{L}_2[a, a+T]$ for any $a, T \in \mathbb{R}$, $T > 0$, are given by*

$$\hat{x}[k_1, k_2] := \left\langle x, \phi^{\text{2D}}_{k_1,k_2} \right\rangle = \int_{t_1=a}^{a+T} \int_{t_2=b}^{b+T} x(t_1, t_2) \exp\left(-\frac{i2\pi k_1 t_1}{T}\right) \exp\left(-\frac{i2\pi k_2 t_2}{T}\right) dt_1\, dt_2. \tag{62}$$

*The Fourier series of order $k_{c,1}$, $k_{c,2}$ is defined as*

$$\mathcal{F}_{k_{c,1},k_{c,2}}\{x\} := \frac{1}{T} \sum_{k_1=-k_{c,1}}^{k_{c,1}} \sum_{k_2=-k_{c,2}}^{k_{c,2}} \hat{x}[k_1, k_2]\phi^{\text{2D}}_{k_1,k_2}. \tag{63}$$

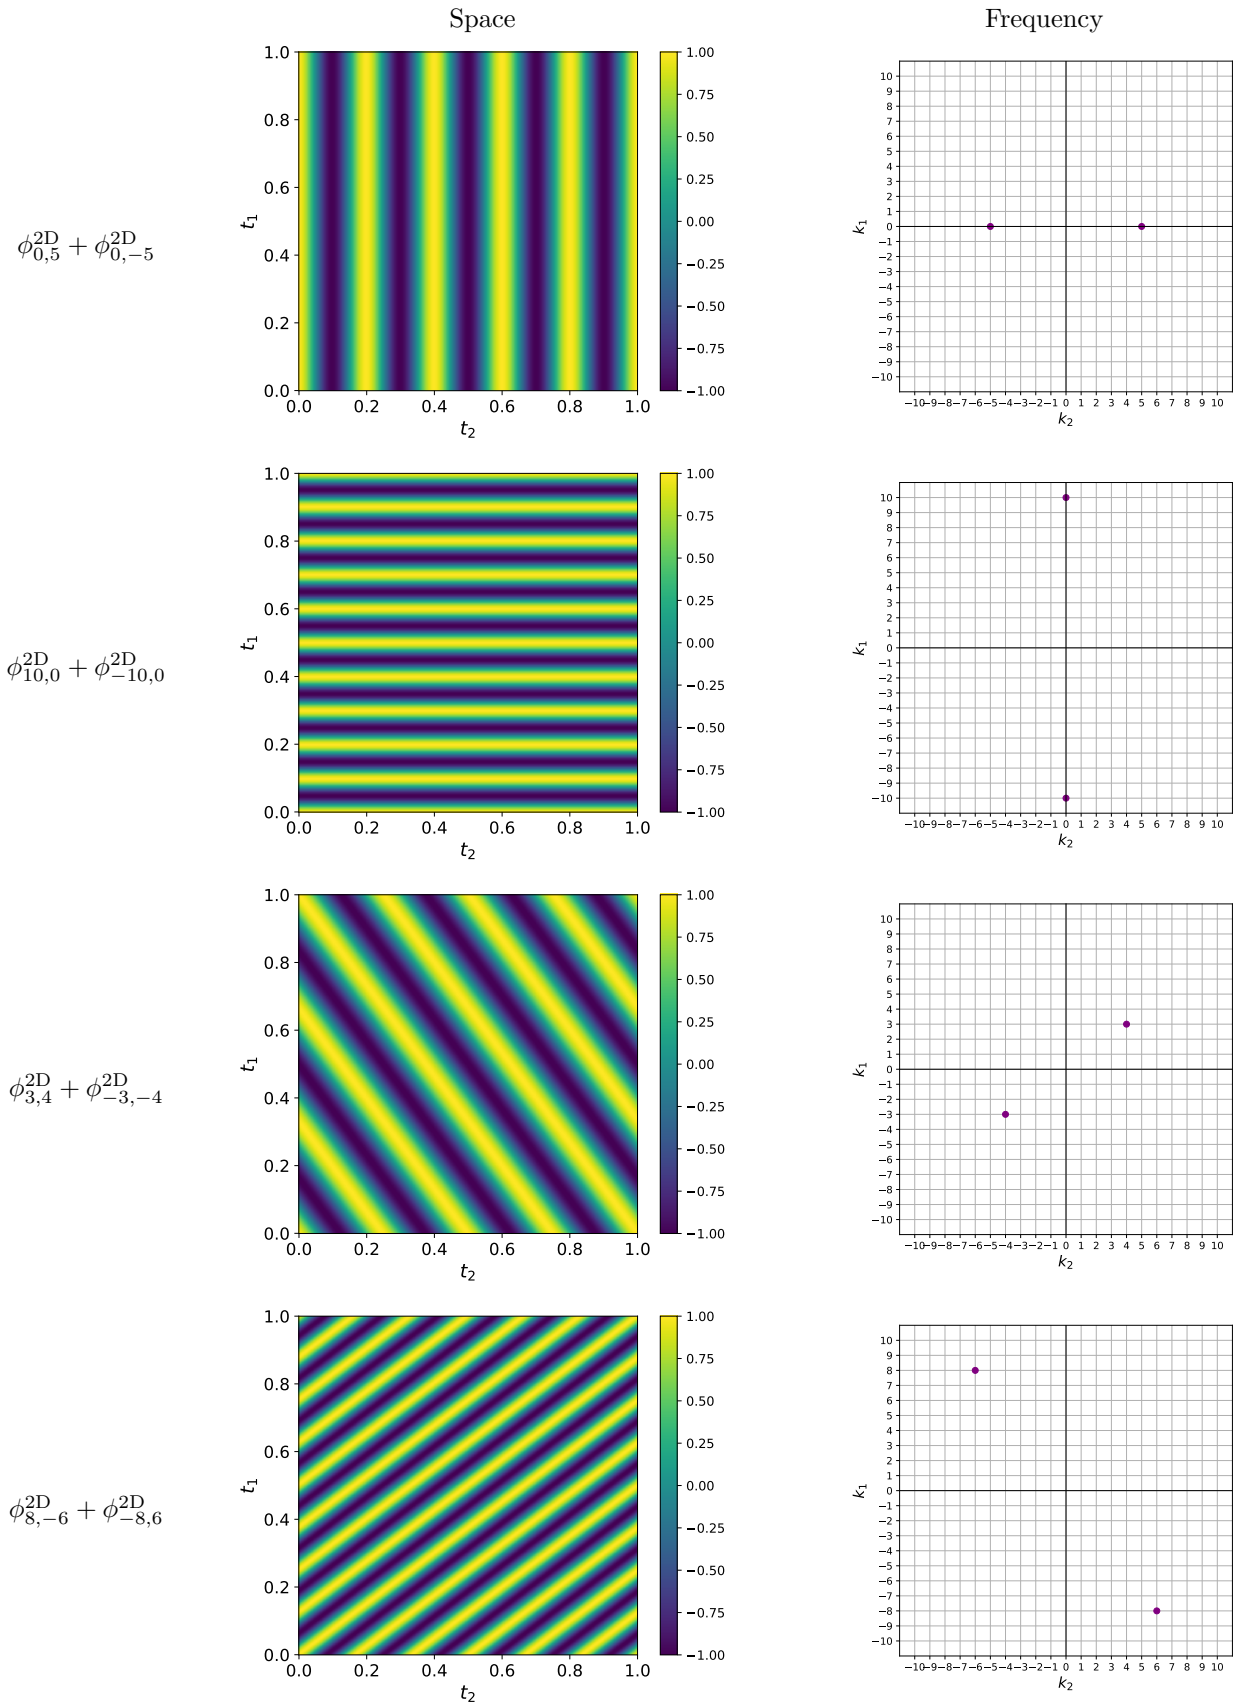


Figure 12: Spatial (left) and frequency (right) representation of real-valued 2D sinusoids with different 2D frequencies.

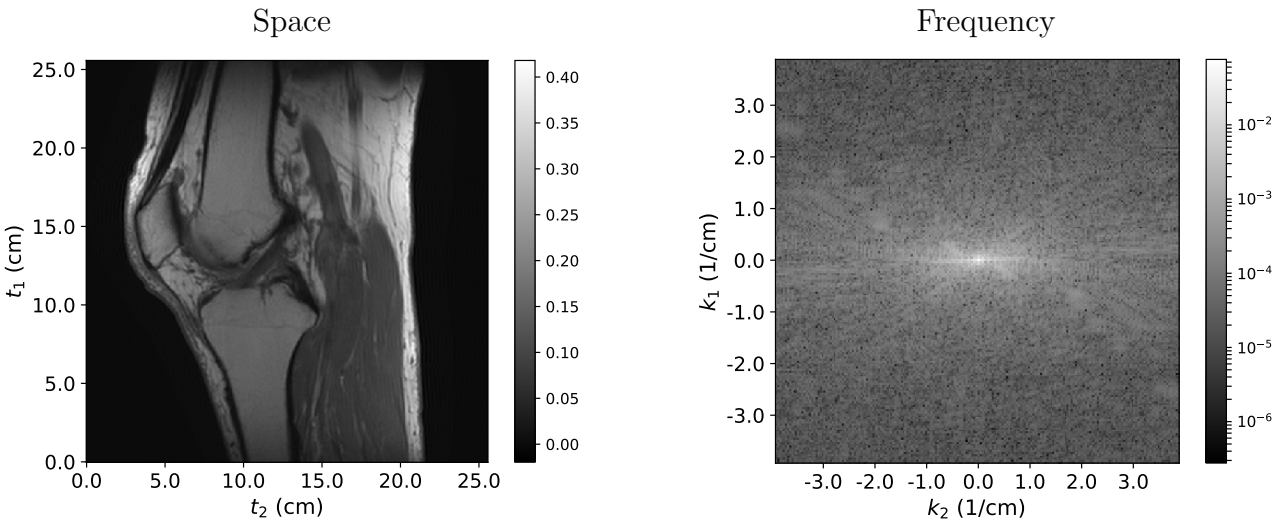


Figure 13: Sagittal section of a human knee measured by magnetic-resonance imaging. The image on the left shows the proton density of the section, whereas the image on the right shows the magnitude of the corresponding Fourier series coefficients.

As in the one-dimensional case, the Fourier series of a square-integrable function in multiple dimensions converges to the function as the order tends to infinity.

**Example 3.4** (Magnetic resonance imaging)**.** Magnetic resonance imaging (MRI) is a medical-imaging technique that measures the response of the atomic nuclei in biological tissues to high-frequency radio waves when placed in a strong magnetic field. Remarkably, the radio waves can be adjusted so that each measurement corresponds to a 2D Fourier coefficient of the proton density of the hydrogen atoms in a region of interest. In MRI the 2D frequency representation is called $k$-space. In order to reconstruct an MRI image, we compute the Fourier series corresponding to the $k$-space measurements. Figure 13 shows an example. In this case, $T := 25$ cm, so that the basic frequency in the Fourier series equals 0.04 $\text{cm}^{-1}$. Figure 14 shows the Fourier series for different values of $k_{c,1}$ and $k_{c,2}$. These values govern the vertical and horizontal resolution of the image. △

## 3.2 Sampling theorem in 2D

As in one dimension, bandlimited signals are signals that are well represented by a finite number of Fourier-series coefficient, like the image in Figure 13. To simplify the definition, we consider Fourier series that have the same order in the vertical and horizontal indices.

**Definition 3.5** (Bandlimited signal)**.** *A signal defined on the 2D rectangle* $[a, a+T] \times [b, b+T]$*, where* $a, b, T \in \mathbb{R}$ *and* $T > 0$ *is bandlimited with a cut-off frequency* $k_c$ *if it is equal to its Fourier series representation of order* $k_c$*, i.e.*

$$x(t_1, t_2) = \sum_{k_1=-k_c}^{k_c} \sum_{k_2=-k_c}^{k_c} \hat{x}[k_1, k_2] \exp\left(\frac{i2\pi k_1 t_1}{T}\right) \exp\left(\frac{i2\pi k_2 t_2}{T}\right). \tag{64}$$

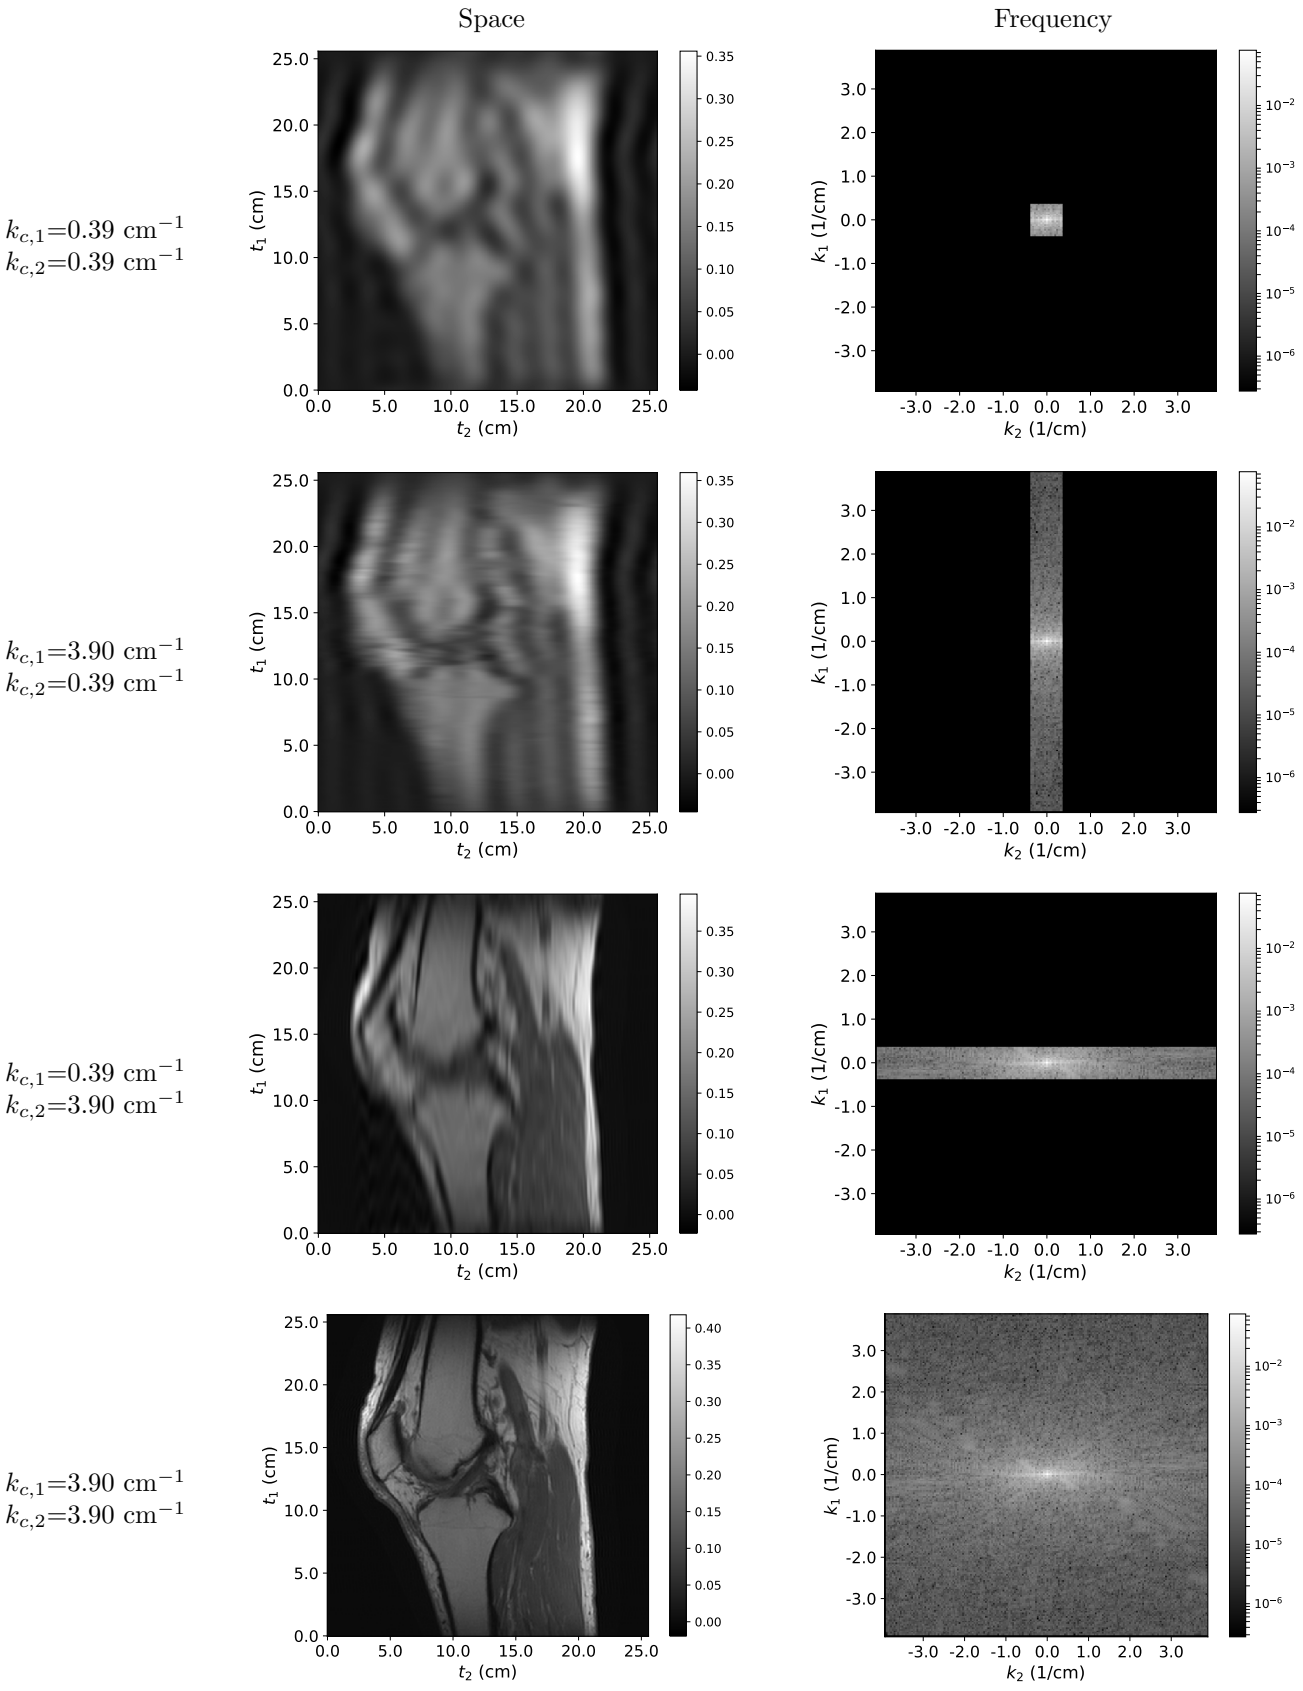


Figure 14: On the right, Fourier series of the image in Figure 13 for different values of $k_{c,1}$ and $k_{c,2}$. The magnitude of the corresponding Fourier coefficients is shown on the right.

Bandlimited functions can be exactly recovered from a finite number of samples. To simplify notation we consider functions supported on the square $[0,T)^2$ for some real $T>0$, without loss of generality. We will focus on equispaced sampling patterns of the form:

$$X_{[N]} := \begin{bmatrix} x\left(\frac{0}{N},\frac{0}{N}\right) & x\left(\frac{0}{N},\frac{T}{N}\right) & \cdots & x\left(\frac{0}{N},T-\frac{T}{N}\right) \\ x\left(\frac{T}{N},\frac{0}{N}\right) & x\left(\frac{T}{N},\frac{T}{N}\right) & \cdots & x\left(\frac{T}{N},T-\frac{T}{N}\right) \\ \cdots & \cdots & \cdots & \cdots \\ x\left(T-\frac{T}{N},\frac{0}{N}\right) & x\left(T-\frac{T}{N},\frac{T}{N}\right) & \cdots & x\left(T-\frac{T}{N},T-\frac{T}{N}\right) \end{bmatrix}. \tag{65}$$

The sampling theorem, generalized to 2D, provides a condition on $N$ to ensure exact recovery of the bandlimited signal from such measurements.

**Theorem 3.6** (Nyquist-Shannon-Kotelnikov sampling theorem)**.** *Any bandlimited signal $x \in \mathcal{L}_2[0,T)^2$, where $T>0$, with cut-off frequency $k_c$ can be recovered exactly from $N^2$ uniformly spaced samples as long as*

$$N \geq 2k_c+1, \tag{66}$$

*where $2k_c+1$ is known as the* Nyquist rate*. Recovery can be carried out by obtaining the matrix of Fourier series coefficients*

$$\widehat{X}_{[k_c]} := \begin{bmatrix} \hat{x}_{-k_c,-k_c} & \hat{x}_{-k_c,-k_c+1} & \cdots & \hat{x}_{-k_c,k_c} \\ \hat{x}_{-k_c+1,-k_c} & \hat{x}_{-k_c+1,-k_c+1} & \cdots & \hat{x}_{-k_c+1,k_c} \\ \cdots & \cdots & \cdots & \cdots \\ \hat{x}_{k_c,-k_c} & \hat{x}_{k_c,-k_c+1} & \cdots & \hat{x}_{k_c,k_c} \end{bmatrix} \tag{67}$$

*from the vector of samples $X_{[N]}$ as follows*

$$\widehat{X}_{[k_c]} = \frac{1}{N^2}\widetilde{F}^*_{[N]} X_{[N]} \left(\widetilde{F}^*_{[N]}\right)^T, \tag{68}$$

*where $\widetilde{F}_{[N]}$ is defined in Eq. (22).*

*Proof.* The result follows from the identity

$$X_{[N]} = \widetilde{F}_{[N]} \widehat{X}_{[k_c]} \widetilde{F}^T_{[N]}, \tag{69}$$

which can be verified by computing the matrix product, and from the fact that $\frac{1}{\sqrt{N}}\widetilde{F}_{[N]}$ is an orthogonal matrix by Lemma 3.8. □

## 3.3 Discrete Fourier transform in 2D

In this section we extend the definition of discrete Fourier transform to two dimensions. We represent discrete signals in 2D as matrices belonging to the vector space of $\mathbb{C}^{N\times N}$ matrices endowed with the standard inner product

$$\langle A,B\rangle := \operatorname{tr}\left(B^*A\right), \quad A,B\in\mathbb{C}^{N\times N}. \tag{70}$$

This is equivalent to representing the signals as vectors in $\mathbb{C}^{N^2}$ with the usual dot product. Discrete complex sinusoids form a basis of the space.

**Definition 3.7** (Discrete complex sinusoids)**.** *The discrete complex sinusoid $\Phi_{k_1,k_2} \in \mathbb{C}^{N\times N}$ with integer frequencies $k_1$ and $k_2$ is defined as*

$$\Phi_{k_1,k_2}[j_1,j_2] := \exp\left(\frac{i2\pi k_1 j_1}{N}\right)\exp\left(\frac{i2\pi k_2 j_2}{N}\right), \qquad 0 \le j_1, j_2 \le N-1, \tag{71}$$

*or, in terms of 1D discrete sinusoids*

$$\Phi_{k_1,k_2} = \vec{\phi}_{k_1}\vec{\phi}_{k_2}^{\,T}. \tag{72}$$

**Lemma 3.8** (Orthonormal sinusoidal basis)**.** *The discrete complex exponentials $\frac{1}{N}\Phi_{k_1,k_2}$, $0 \le k_1, k_2 \le N-1$, form an orthonormal basis of $\mathbb{C}^{N\times N}$.*

*Proof.* By Eq. (72), we can reformulate the inner product of two 2D discrete sinusoids in terms of 1D discrete sinusoids

$$\begin{aligned}\langle \Phi_{k_1,k_2}, \Phi_{l_1,l_2}\rangle &= \operatorname{tr}\left(\left(\Phi_{l_1,l_2}\right)^* \Phi_{k_1,k_2}\right) &(73)\\ &= (\vec{\phi}_{k_1})^*\vec{\phi}_{l_1}(\vec{\phi}_{k_2})^*\vec{\phi}_{l_2}. &(74)\end{aligned}$$

Lemma 3.8 then implies that the inner product equals $N$ if $k_1 = l_1$ and $k_2 = l_2$, and zero otherwise. Since there are $N^2$ vectors in the set and they are linearly independent, they form a basis of $\mathbb{C}^{N^2}$. □

The 2D discrete Fourier transform is just a change of basis that expresses a 2D vector in terms of 2D discrete complex sinusoids.

**Definition 3.9** (2D discrete Fourier transform)**.** *The discrete Fourier transform (DFT) of a 2D array $X \in \mathbb{C}^{N\times N}$ is given by*

$$\widehat{X}[k_1,k_2] := \langle X, \Phi_{k_1,k_2}\rangle, \qquad 0 \le k_1, k_2 \le N-1, \tag{75}$$

*or equivalently by*

$$\widehat{X} := F_{[N]} X F_{[N]}, \tag{76}$$

*where $F_{[N]}$ is the 1D DFT matrix defined in Eq. (42). The inverse DFT of a 2D array $\widehat{Y} \in \mathbb{C}^{N\times N}$ equals*

$$Y = \frac{1}{N^2} F_{[N]}^* \widehat{Y} F_{[N]}^*. \tag{77}$$

Recall that by Lemma 3.8 the rows (and columns) of $F_{[N]}$ are orthogonal and have norm $\sqrt{N}$, which justifies the definition of the inverse DFT.

**Corollary 3.10.** *The inverse 2D DFT inverts the 2D DFT.*

As in one dimension, if we interpret the entries of a matrix $X$ as equispaced samples from a bandlimited signal measured at the Nyquist rate, then the 2D DFT of $X$ is exactly equal to the 2D Fourier series of $x$ (recall that the rows of $F_{[N]}$ are the same as the rows of $\widetilde{F}_{[N]}$, only in a different order). By Lemma (2.17), the complexity of computing the 2D DFT using the FFT algorithm is $O(N^2 \log N)$. In contrast, a naive implementation using matrix-vector multiplications would have complexity $O(N^3)$.

**Example 3.11** (Undersampling in the frequency domain)**.** An important goal in MRI is to reduce the time that the patient stays in the scanner. Scanner time is very expensive, and patients tend to move during long scans, which decreases the image quality. Scan time is roughly proportional to the frequency samples measured by the MR scanner. As shown in Figure 14, limiting the number of $k$-space samples by reducing the cut-off frequency results in a loss of resolution. An alternative could be to undersample regularly up to the cut-off frequency, i.e. for example measure the Fourier coefficients $\hat{x}[k_1, k_2]$ for $k_1 = -k_c, -k_c + 2, \ldots, k_c - 2, k_c$ instead of for all integer values of $k_1$ between $-k_c$ and $k_c$.

Unfortunately undersampling in the frequency domain produces aliasing artifacts in the image domain, just like sampling below the Nyquist rate produces aliasing in the frequency domain (recall Lemma 2.10). To gain some intuition, let us sketch a simple proof when the undersampling factor equals 2 in the horizontal axis. Consider an image $X \in \mathbb{C}^{N \times N}$, where $N$ is even to simplify matters. By Lemma 2.15, the submatrix of even columns of $F_{[N]}$ equals

$$(F_{[N]})_{\text{even}} := \begin{bmatrix} F_{[N/2]} \\ F_{[N/2]} \end{bmatrix}. \tag{78}$$

Let us separate the image into its left $N/2$ columns and its right $N/2$ columns

$$X := \begin{bmatrix} X_{\text{left}} & X_{\text{right}} \end{bmatrix}. \tag{79}$$

If we undersample by a factor of 2 in the $k_2$ axis, the data equal

$$\widehat{X}_{even} = F_{[N]} X (F_{[N]})_{\text{even}} \tag{80}$$

$$= F_{[N]} \begin{bmatrix} X_{\text{left}} & X_{\text{right}} \end{bmatrix} \begin{bmatrix} F_{[N/2]} \\ F_{[N/2]} \end{bmatrix} \tag{81}$$

$$= F_{[N]} \left( X_{\text{left}} + X_{\text{right}} \right) F_{[N/2]}. \tag{82}$$

The left and right halves are scrambled in the measurements. Let $\widehat{Y} \in \mathbb{C}^{N \times N}$ be a matrix whose even columns equal $\widehat{X}_{even}$ and whose odd columns are zero. Recovering from $\widehat{Y}$ as if we had measured all Fourier coefficients yields

$$Y = \frac{1}{N^2} F^*_{[N]} \widehat{Y} F^*_{[N]} \tag{83}$$

$$= \frac{1}{N^2} F^*_{[N]} \widehat{X}_{\text{even}} (F_{[N]})^*_{\text{even}} \tag{84}$$

$$= \frac{1}{N} F^*_{[N]} F_{[N]} \left( X_{\text{left}} + X_{\text{right}} \right) \frac{1}{N} F_{[N/2]} \begin{bmatrix} F^*_{[N/2]} & F^*_{[N/2]} \end{bmatrix} \tag{85}$$

$$= \frac{1}{2} \begin{bmatrix} X_{\text{left}} + X_{\text{right}} & X_{\text{left}} + X_{\text{right}} \end{bmatrix}. \tag{86}$$

The aliased image is equal to the sum of the original image and a circular shift. Figure 15 illustrates this phenomenon. △

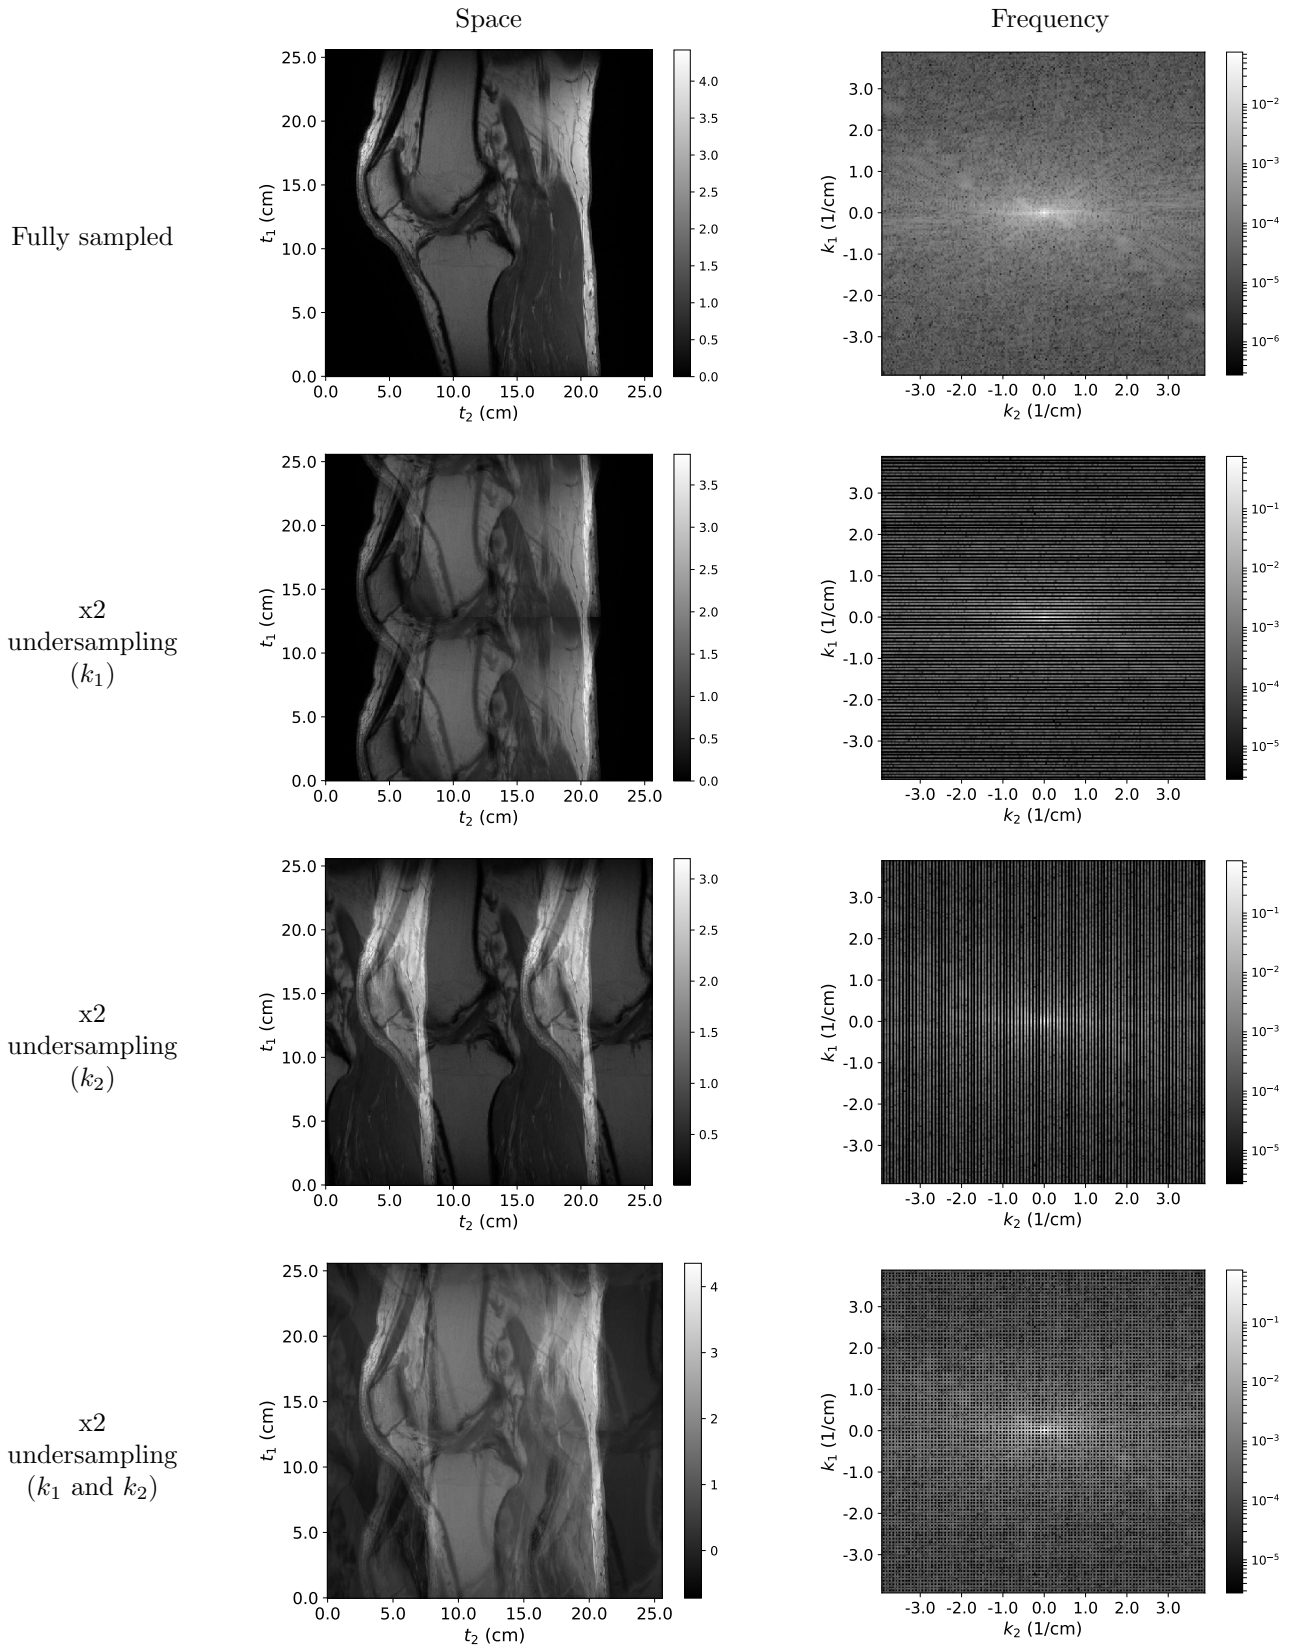


Figure 15: Aliasing in the spatial domain (left) caused by different undersampling patterns in the frequency domain (right) for the MR image in Figure 13.

# 4 Translation Invariance

As discussed in Section 1.3, linear models that map signals of dimension $N$ to signals of have $N^2$ parameters. This makes them intractable for signals like images, audio, video, etc. In order to design more parsimonious models, we consider classes of signals that are *translation invariant*. If a signal $x$ belongs to a set of signals that is translation invariant, then any shift of $x$ also belongs to the set. Audio, images and videos are all translation invariant; a shifted picture of a dog is still a picture of a dog.

## 4.1 Translation

In practice, when considering translations of finite-dimensional signals one needs to be careful about border effects. We focus on circular translations, which assume that signals lie on a circle in 1D or a torus in 2D.

**Definition 4.1** (Circular translation)**.** *We denote by $\vec{x}^{\downarrow s}$ the $s$th circular translation or shift of a vector $\vec{x} \in \mathbb{C}^N$. For all $1 \leq j \leq N$,*

$$\vec{x}^{\downarrow s}[j] = \begin{cases} \vec{x}\,[j-s] & \text{if } 1 \leq j-s \leq N, \\ \vec{x}\,[N+j-s] & \text{if } j-s < 1. \end{cases} \tag{87}$$

*For an $N \times N$ signal $X \in \mathbb{C}^{N \times N}$, circular translation by $(s_1, s_2) \in$ is denoted by $X^{\downarrow (s_1,s_2)}$. For all $1 \leq j_1, j_2 \leq N$,*

$$X^{\downarrow (s_1,s_2)}[j_1, j_2] = \begin{cases} X[j_1 - s_1, j_2 - s_2] & \text{if } 1 \leq j_1 - s_1, j_2 - s_2 \leq N, \\ X[N + j_1 - s_1, j_2 - s_2] & \text{if } j_1 - s_1 < 1, \\ X[j_1 - s_1, N + j_2 - s_2] & \text{if } j_2 - s_2 < 1, \\ X[N + j_1 - s_1, N + j_2 - s_2] & \text{if } j_1 - s_1 < 1,\ j_2 - s_2 < 1. \end{cases} \tag{88}$$

Recall that we use complex exponentials to represent sinusoids when computing frequency representations, because the phase can be encoded as a linear coefficient. Shifting a sinusoid amounts to modifying its phase. As a result, shifting a complex sinusoid is equivalent to multiplying it by a complex constant.

**Lemma 4.2.** *For any discrete complex sinusoid $\vec{\phi}_k \in \mathbb{C}^N$ with integer frequency $k$ and any integer shift $s$*

$$\vec{\phi}_k^{\downarrow s} = \exp\left(-\frac{i2\pi ks}{N}\right) \vec{\phi}_k. \tag{89}$$

*Proof.* By the definition of a complex sinusoid,

$$\vec{\phi}_k^{\downarrow s}[l] = \exp\left(\frac{i2\pi k(l-s)}{N}\right) \tag{90}$$

$$= \exp\left(-\frac{i2\pi ks}{N}\right) \vec{\phi}_k[l]. \tag{91}$$

□

**Corollary 4.3.** *The discrete complex sinusoid $\Phi_{k_1,k_2} \in \mathbb{C}^{N\times N}$ with integer frequencies $k_1$ and $k_2$*

$$\Phi_{k_1,k_2}^{\downarrow(s_1,s_2)} = \exp\left(-\frac{i2\pi k_1 s_1}{N}\right) \exp\left(-\frac{i2\pi k_2 s_2}{N}\right) \Phi_{k_1,k_2}. \tag{92}$$

*Proof.* The result follows immediately from Lemma 4.2 because $\Phi_{k_1,k_2}^{\downarrow(s_1,s_2)} = \vec{\phi}_{k_1}^{\downarrow s_1} (\vec{\phi}_{k_2}^{\downarrow s_2})^T$. □

One can think of a vector as a sum of discrete sinusoids each scaled by the corresponding DFT coefficient. If the vector is shifted, then Lemma 4.2 implies that each discrete sinusoid is just scaled by a complex constant. This yields a simple expression for the Fourier coefficients of shifted signals.

**Theorem 4.4.** *Let $\vec{x} \in \mathbb{C}^N$ with DFT $\hat{x}$ and $y := x^{\downarrow s}$, $0 \leq s \leq N-1$. The DFT coefficients of $y$ are given by*

$$\hat{y}\,[k] := \exp\left(-\frac{i2\pi ks}{N}\right) \hat{x}\,[k]\,, \quad 1 \leq k \leq N. \tag{93}$$

*Let $X \in \mathbb{C}^{N\times N}$ with DFT $\widehat{X}$ and $Y := X^{\downarrow(s_1,s_2)}$, $0 \leq s_1, s_2 \leq N-1$. The DFT coefficients of $Y$ are given by*

$$\widehat{Y}\,[k_1,k_2] := \exp\left(-\frac{i2\pi ks_1}{N}\right) \exp\left(-\frac{i2\pi ks_2}{N}\right) \widehat{X}\,[k_1,k_2]\,, \quad 1 \leq k_1, k_2 \leq N. \tag{94}$$

*Proof.* We only prove the one-dimensional case, the 2D case follows by the same argument:

$$\begin{aligned}
\hat{y}\,[k] &:= \langle \vec{x}^{\downarrow s}, \vec{\phi}_k \rangle && (95)\\
&= \langle \vec{x}, \vec{\phi}_k^{\downarrow -s} \rangle && (96)\\
&= \left\langle \vec{x}, \exp\left(\frac{i2\pi ks}{N}\right) \vec{\phi}_k \right\rangle \quad \text{by Lemma 4.2} && (97)\\
&= \exp\left(-\frac{i2\pi ks}{N}\right) \langle \vec{x}, \vec{\phi}_k \rangle && (98)\\
&= \exp\left(-\frac{i2\pi ks}{N}\right) \hat{x}\,[k]\,. && (99)
\end{aligned}$$

□

## 4.2 Linear translation-invariant maps

Let us consider a linear map $\mathcal{D} : \mathbb{C}^N \to \mathbb{C}^N$ that performs denoising on a class of translation-invariant signals. For a fixed noisy data vector $\vec{y}$ the denoised signal equals $\mathcal{D}(\vec{y})$. How should the output change if we shift the input $\vec{y}^{\downarrow s}$? If the statistics of the noise are translation invariant (i.e. the noise level does not depend on the location), the output should just be a shifted copy of $\mathcal{D}(\vec{y})$: $\mathcal{D}(\vec{y})^{\downarrow s}$. Linear maps with this property are said to be translation or shift invariant.

**Definition 4.5** (Linear translation-invariant map)**.** *A map $\mathcal{L}$ from $\mathbb{C}^N$ to $\mathbb{C}^N$ is linear if for any vectors $\vec{x}, \vec{y} \in \mathbb{C}^N$ and any constant $\alpha \in \mathbb{C}$*

$$\mathcal{L}(\vec{x} + \vec{y}) = \mathcal{L}(\vec{x}) + \mathcal{L}(\vec{y}), \tag{100}$$

$$\mathcal{L}(\alpha \vec{x}) = \alpha \mathcal{L}(\vec{x}), \tag{101}$$

*and translation invariant if for any shift $0 \leq s \leq N-1$*

$$\mathcal{L}(\vec{x}^{\downarrow s}) = \mathcal{L}(\vec{x})^{\downarrow s}. \tag{102}$$

*A map $\mathcal{L}$ from $\mathbb{C}^{N \times N}$ to $\mathbb{C}^{N \times N}$ is linear if for any $X, Y \in \mathbb{C}^{N \times N}$ and any constant $\alpha \in \mathbb{C}$*

$$\mathcal{L}(X + Y) = \mathcal{L}(X) + \mathcal{L}(Y), \tag{103}$$

$$\mathcal{L}(\alpha X) = \alpha \mathcal{L}(X), \tag{104}$$

*and translation invariant if for any $0 \leq s_1, s_2 \leq N-1$*

$$\mathcal{L}(X^{\downarrow (s_1, s_2)}) = \mathcal{L}(X)^{\downarrow (s_1, s_2)}. \tag{105}$$

A linear translation-invariant (LTI) map can be parametrized by its action on the first standard-basis vector. Standard-basis vectors are often called *impulses* in the signal-processing literature, so the corresponding output of the map is called an impulse response. Impulse responses are also often called *filters*.

**Definition 4.6** (Impulse response)**.** *Let $\vec{e}_0 \in \mathbb{C}^N$ be the first standard-basis vector if we index the entries from 0 to $N-1$ ($\vec{e}_0[0] = 1$ and $\vec{e}_0[j] = 0$ for $1 \leq j \leq N-1$). The vector $\vec{h}_{\mathcal{L}} \in \mathbb{C}^n$ obtained by applying a map $\mathcal{L} : \mathbb{C}^n \to \mathbb{C}^n$ to $\vec{e}_0$ is called the impulse response of the map,*

$$\vec{h}_{\mathcal{L}} := \mathcal{L}(\vec{e}_0). \tag{106}$$

*Let $E_0 \in \mathbb{C}^{N \times N}$ be the first standard-basis vector in $\mathbb{C}^{N \times N}$ if we index each dimension from 0 to $N-1$ ($E[0,0] = 1$ and $E[j_1, j_2] = 0$ for $1 \leq j_1, j_2 \leq N-1$). The vector $H_{\mathcal{L}} \in \mathbb{C}^{N \times N}$ obtained by applying a map $\mathcal{L} : \mathbb{C}^{N \times N} \to \mathbb{C}^{N \times N}$ to $E_0$ is the impulse response of the map,*

$$H_{\mathcal{L}} := \mathcal{L}(E_0). \tag{107}$$

In order to express LTI systems in terms of their impulse response, we first need to introduce an operation between two vectors known as convolution. Note that we index the entries of the vectors from 0 to $N-1$, but this is without loss of generality, the definition is equivalent if we consider any $N$ contiguous integers.

**Definition 4.7** (Circular convolution)**.** *The circular convolution between two vectors $\vec{x}, \vec{y} \in \mathbb{C}^N$ is defined as*

$$\vec{x} * \vec{y}\,[j] := \sum_{s=0}^{N-1} \vec{x}\,[s]\, \vec{y}^{\downarrow s}\,[j], \quad 0 \leq j \leq N-1. \tag{108}$$

*The 2D circular convolution between $X \in \mathbb{C}^{N \times N}$ and $Y \in \mathbb{C}^{N \times N}$ is defined as*

$$X * Y\,[j_1, j_2] := \sum_{s_1=0}^{N-1} \sum_{s_2=0}^{N-1} X\,[s_1, s_2]\, Y^{\downarrow (s_1, s_2)}\,[j_1, j_2], \quad 0 \leq j_1, j_2 \leq N-1. \tag{109}$$

A general linear map from $\mathbb{C}^N$ to $\mathbb{C}^N$ can be parametrized by a matrix with $N^2$ entries. In contrast, the next theorem establishes that an LTI map can be parametrized by its impulse response, which has dimension $N$. This is a dramatic reduction in the number of parameters.

**Theorem 4.8.** *The action of a LTI map $\mathcal{L} : \mathbb{C}^N \to \mathbb{C}^N$ on a vector $\vec{x} \in \mathbb{C}^N$ is equal to the circular convolution of $\vec{x}$ and the impulse response $\vec{h}_{\mathcal{L}}$ of $\mathcal{L}$,*

$$\mathcal{L}(\vec{x}) = \vec{x} * \vec{h}_{\mathcal{L}}. \tag{110}$$

*The action of a 2D LTI map $\mathcal{L} : \mathbb{C}^{N \times N} \to \mathbb{C}^{N \times N}$ on $X \in \mathbb{C}^{N \times N}$ is equal to the circular convolution of $X$ and the impulse response $H_{\mathcal{L}}$ of $\mathcal{L}$,*

$$\mathcal{L}(X) = X * H_{\mathcal{L}}. \tag{111}$$

*Proof.* We prove the 1D case; the 2D case follows by the same argument. To simplify the argument we index the entries of the vector from 0 to $N-1$. Expressed in terms of the basis vectors $\vec{x} = \sum_{j=0}^{N-1} \vec{x}[j]\vec{e}_j$. By linearity and translation invariance,

$$\begin{aligned}
\mathcal{L}(\vec{x}) &= \mathcal{L}\left(\sum_{j=0}^{N-1} \vec{x}[j]\vec{e}_j\right) && (112)\\
&= \mathcal{L}\left(\sum_{j=0}^{N-1} \vec{x}[j]\vec{e}_0^{\downarrow j}\right) && (113)\\
&= \sum_{j=0}^{N-1} \vec{x}[j]\mathcal{L}\left(\vec{e}_0^{\downarrow j}\right) && (114)\\
&= \sum_{j=0}^{N-1} \vec{x}[j]\mathcal{L}(\vec{e}_0)^{\downarrow j} && (115)\\
&= \sum_{j=0}^{N-1} \vec{x}[j]\vec{h}_{\mathcal{L}}^{\downarrow j}. && (116)
\end{aligned}$$

□

The convolution between two vectors can be efficiently computed by multiplying their Fourier coefficients. This is illustrated in Figures 16 and 17.

**Theorem 4.9** (Convolution in time is multiplication in frequency)**.** *Let $\vec{y} := \vec{x}_1 * \vec{x}_2$ for $\vec{x}_1, \vec{x}_2 \in \mathbb{C}^N$. Then the DFT of $\vec{y}$ equals*

$$\hat{y}[k] = \hat{x}_1[k]\,\hat{x}_2[k], \quad 0 \le k \le N-1, \tag{117}$$

*$\hat{x}_1$ and $\hat{x}_2$ are the DFTs of $\vec{x}_1$ and $\vec{x}_2$ respectively.*

*Let $Y := X_1 * X_2$ for $X_1, X_2 \in \mathbb{C}^{N \times N}$. Then the 2D DFT of $Y$ is given by*

$$\widehat{Y}[k_1, k_2] = \widehat{X}_1[k_1, k_2]\,\widehat{X}_2[k_1, k_2]. \tag{118}$$

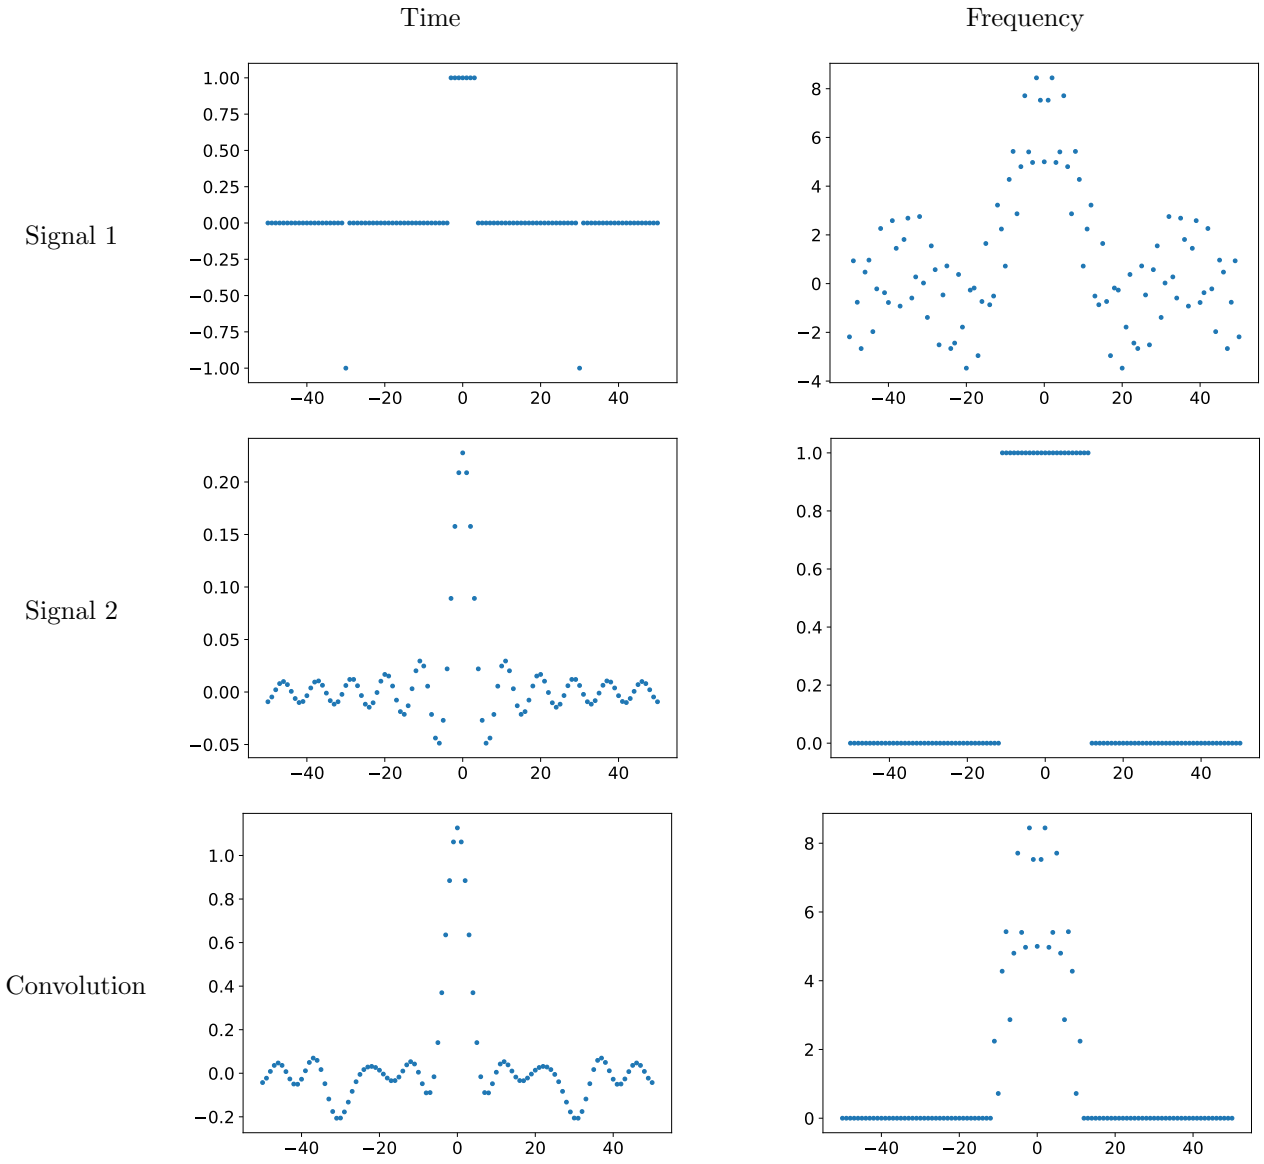


Figure 16: Convolution between two one-dimensional signals. The frequency representation of both signals and their convolution is shown on the right. One of them is a low-pass filter, which only has nonzero coefficients in a low-frequency band. As a result, these coefficients are also zero in the convolved signal.

*Proof.* We have

$$\hat{y}\left[k\right] := \left\langle \vec{x}_1 * \vec{x}_2, \vec{\phi}_k \right\rangle \tag{119}$$

$$= \left\langle \sum_{s=0}^{N-1} \vec{x}_1\left[s\right] \vec{x}_2^{\downarrow s}, \vec{\phi}_k \right\rangle \tag{120}$$

$$= \left\langle \sum_{s=0}^{N-1} \vec{x}_1\left[s\right] \frac{1}{N} \sum_{j=0}^{N-1} \exp\left(-\frac{i2\pi js}{N}\right) \hat{x}_2[j] \vec{\phi}_j, \vec{\phi}_k \right\rangle \quad \text{by Theorem 4.4} \tag{121}$$

$$= \sum_{j=0}^{N-1} \sum_{s=0}^{N-1} \vec{x}_1\left[s\right] \exp\left(-\frac{i2\pi js}{N}\right) \frac{1}{N} \left\langle \hat{x}_2[j] \vec{\phi}_j, \vec{\phi}_k \right\rangle \tag{122}$$

$$= \sum_{j=0}^{N-1} \hat{x}_1[j] \hat{x}_2[j] \frac{1}{N} \left\langle \vec{\phi}_j, \vec{\phi}_k \right\rangle \tag{123}$$

$$= \hat{x}_1[k] \hat{x}_2[k]. \tag{124}$$

The 2D result follows from the same argument. □

A useful consequence of this theorem is that we can compute convolutions by computing the DFT of the two vectors via the FFT, multiplying their DFT coefficients, and then taking the inverse DFT again via the FFT. The complexity for $N$-dimensional vectors is of order $N \log N$.

The DFT of the impulse response of an LTI map is known as its transfer function. By Theorems 4.8 and 4.9 applying an LTI map to a vector just scales each of its Fourier coefficients by the corresponding entry of the transfer function.

**Corollary 4.10.** *Let $\hat{h}_{\mathcal{L}}$ be the DFT of the impulse response of an LTI map $\mathcal{L}$. Then for any $\vec{x} \in C^N$*

$$\mathcal{L}(\vec{x}) = \sum_{k=0}^{N-1} \hat{h}_{\mathcal{L}}[k] \hat{x}[k] \vec{\phi}_k. \tag{125}$$

*Let $\widehat{H}_{\mathcal{L}}$ be the 2D DFT of the impulse response of a 2D LTI map $\mathcal{L}$. Then for any $X \in C^{N \times N}$*

$$\mathcal{L}(X) = \sum_{k_1=0}^{N-1} \sum_{k_2=0}^{N-1} \hat{H}_{\mathcal{L}}[k_1, k_2] \widehat{X}[k_1, k_2] \Phi_{k_1,k_2}. \tag{126}$$

## 4.3 Wiener filtering

LTI maps can be applied to estimate signals from data. This is equivalent to solving a regression problem where the function is constrained to be linear and translation invariant. The assumption of translation invariance results in a dramatic reduction of parameters with respect to a general linear transformation. The optimal LTI map with respect to the $\ell_2$-norm estimation error is known as the Wiener filter. The following theorem derives the expression of the transfer function of the Wiener filter for an arbitrary training dataset of examples.

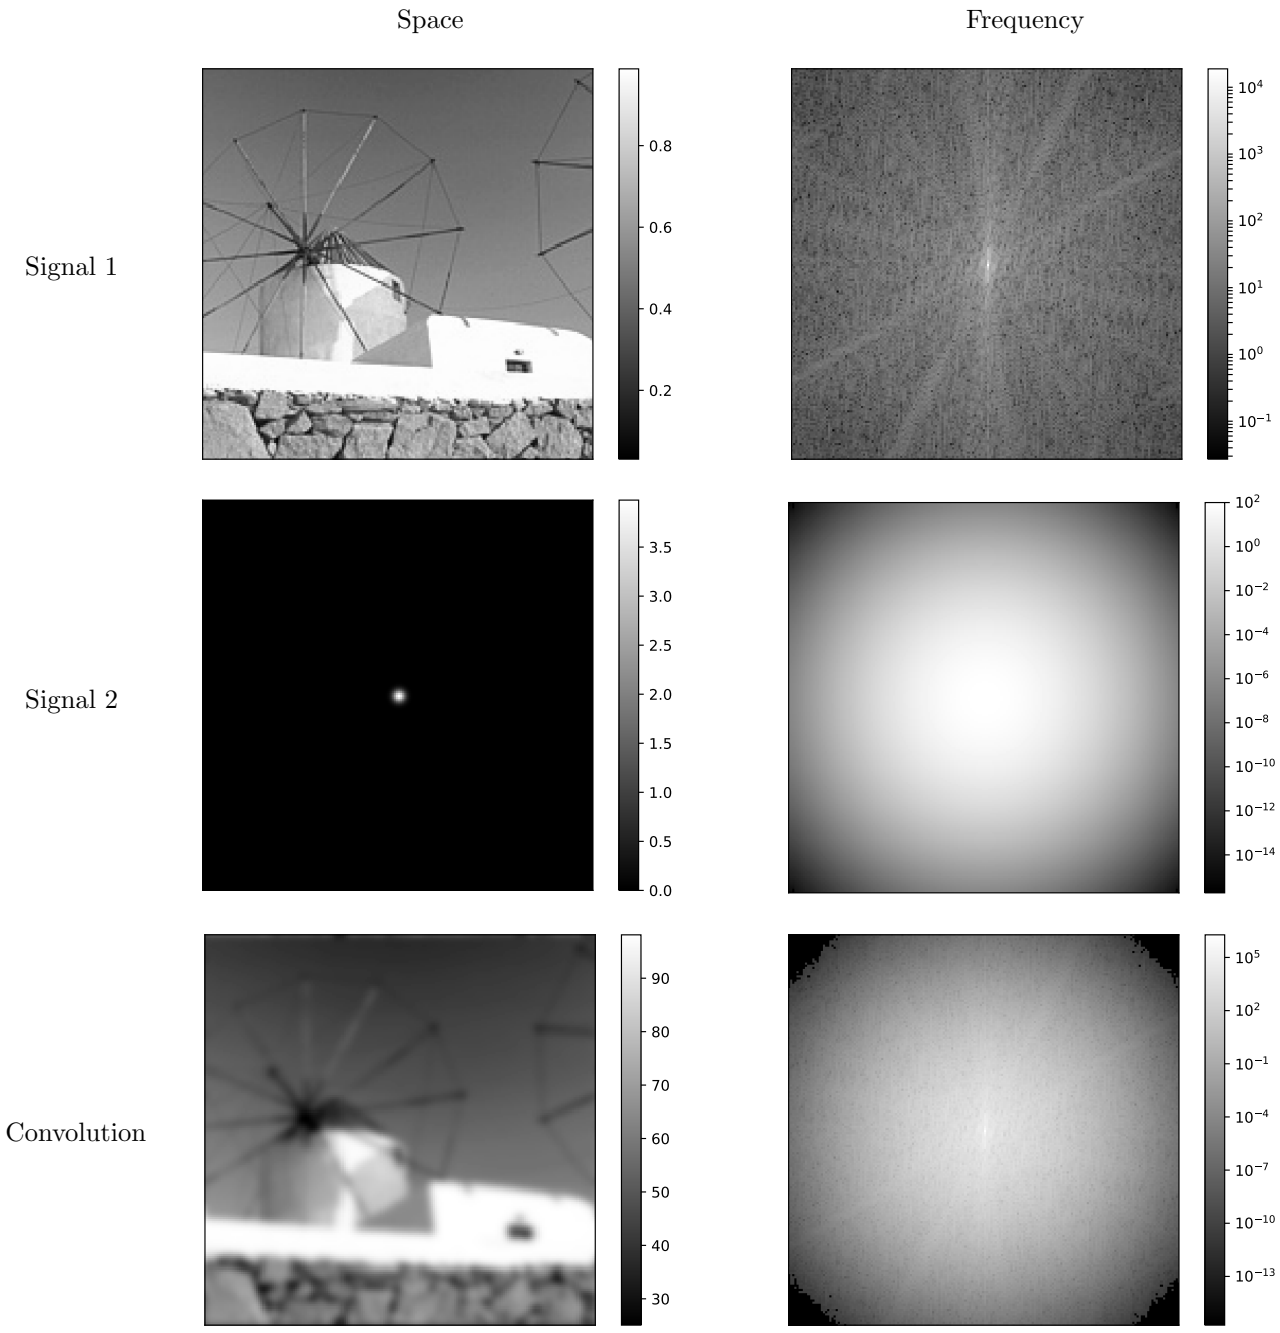


Figure 17: Image blurring caused by diffraction in optical systems is often modeled as a convolution between a high-resolution image and a blurring kernel or point-spread function. The result is a low-resolution image due to the local averaging resulting from convolving with the kernel, which in this example is Gaussian. In the frequency domain, the high-frequency components of the image are suppressed because the corresponding frequency components of the kernel are very small.

**Theorem 4.11** (Wiener filter)**.** *Let* $(\vec{x}^{[1]}, \vec{y}^{[1]})$, ..., $(\vec{x}^{[n]}, \vec{y}^{[n]})$ *be a training set of* $n$ *pairs of clean and noisy signals of dimension* $N$*. The solution to the* $\ell_2$*-norm loss optimization problem*

$$\min_{\vec{w}\in\mathbb{C}^N} \sum_{j=1}^{n} \left|\left|\vec{x}^{[j]} - \vec{w} * \vec{y}^{[j]}\right|\right|_2^2 \tag{127}$$

*is an LTI map with a transfer function given by*

$$\hat{w}_{\text{opt}}[k] = \frac{\sum_{j=1}^{n} \hat{x}^{[j]}[k]\overline{\hat{y}^{[j]}[k]}}{\sum_{j=1}^{n} \left|\hat{y}^{[j]}[k]\right|^2}, \quad 1 \leq k \leq N. \tag{128}$$

*If the signals are two dimensional:* $(X^{[1]}, Y^{[1]})$, ..., $(X^{[n]}, Y^{[n]})$ *then the transfer function of the optimal filter equals*

$$\widehat{W}_{\text{opt}}[k_1, k_2] = \frac{\sum_{j=1}^{n} \widehat{X}^{[j]}[k_1, k_2]\overline{\widehat{Y}^{[j]}[k_1, k_2]}}{\sum_{j=1}^{n} \left|\widehat{Y}^{[j]}[k_1, k_2]\right|^2}, \quad 1 \leq k_1, k_2 \leq N. \tag{129}$$

*Proof.* The cost function can be rewritten in terms of the Fourier coefficients of the signals in the training set by Theorem 4.9,

$$C\left(w\right) := \frac{1}{2} \sum_{j=1}^{n} \left|\left|\vec{x}^{[j]} - \vec{w} * \vec{y}^{[j]}\right|\right|_2^2 \tag{130}$$

$$= \frac{1}{2} \sum_{j=1}^{n} \left|\left|\frac{1}{N} F_{[N]}^{*}(\hat{x}^{[j]} - \hat{w} \circ \hat{y}^{[j]})\right|\right|_2^2 \tag{131}$$

$$= \frac{1}{2N} \sum_{j=1}^{n} \left|\left|\hat{x}^{[j]} - \hat{w} \circ \hat{y}^{[j]}\right|\right|_2^2 \tag{132}$$

$$= \frac{1}{2N} \sum_{j=1}^{n} \sum_{k=1}^{N} \left|\hat{x}^{[j]}[k] - \hat{w}[k]\hat{y}^{[j]}[k]\right|^2 := \frac{1}{N} \sum_{k=1}^{N} \widetilde{C}_k\left(\hat{w}[k]\right). \tag{133}$$

This decouples the cost function into $N$ terms, each of which only depends on one Fourier coefficient.

$$\widetilde{C}_k\left(\alpha\right) := \frac{1}{2} \sum_{j=1}^{n} \left|\hat{x}^{[j]}[k] - \alpha\, \hat{y}^{[j]}[k]\right|^2 \tag{134}$$

$$= \frac{1}{2} \sum_{j=1}^{n} \left|\hat{x}^{[j]}[k]\right|^2 - 2\,\text{Re}\left\{\hat{x}^{[j]}[k]\overline{\hat{y}^{[j]}[k]\alpha}\right\} + \left|\hat{y}^{[j]}[k]\right|^2 |\alpha|^2 \tag{135}$$

$$= \frac{1}{2} \sum_{j=1}^{n} \left|\hat{x}^{[j]}[k]\right|^2 - 2\,\text{Re}\left\{\hat{x}^{[j]}[k]\overline{\hat{y}^{[j]}[k]}\right\}\alpha_R - 2\,\text{Im}\left\{\hat{x}^{[j]}[k]\overline{\hat{y}^{[j]}[k]}\right\}\alpha_I + \left|\hat{y}^{[j]}[k]\right|^2 \left(\alpha_R^2 + \alpha_I^2\right),$$

where $\alpha_R$ and $\alpha_I$ denote the real and imaginary part of $\alpha$ respectively. Setting the derivatives with respect to these parameters to zero, we obtain

$$\arg\min \widetilde{C}_k\left(\alpha\right) = \frac{\sum_{j=1}^{n} \text{Re}\left\{\hat{x}^{[j]}[k]\overline{\hat{y}^{[j]}[k]}\right\}}{\sum_{j=1}^{n} \left|\hat{y}^{[j]}[k]\right|^2} + i\frac{\sum_{j=1}^{n} \text{Im}\left\{\hat{x}^{[j]}[k]\overline{\hat{y}^{[j]}[k]}\right\}}{\sum_{j=1}^{n} \left|\hat{y}^{[j]}[k]\right|^2}. \tag{136}$$

The proof for the 2D result is analogous. □

Optimal linear translation-invariant filtering is equivalent to scaling each Fourier coefficient of the noisy signal by a fixed constant. In the case of additive iid Gaussian noise, the scaling has a very simple expression.

**Example 4.12** (Denoising via Wiener filtering)**.** Let us consider a denoising problem where the data are modeled probabilistically as $n$ samples from a random vector

$$\vec{\mathbf{y}} = \vec{\mathbf{x}} + \vec{\mathbf{z}}, \tag{137}$$

where $\vec{\mathbf{z}}$ is zero-mean Gaussian noise with variance $\sigma^2$, which is independent of the signal $\vec{\mathbf{x}}$. The goal is to estimate the signal from the noisy measurements.

As we will see in a future lecture, a linear transformation $A\vec{\mathbf{z}}$ of a Gaussian vector with mean $\vec{\mu}$ and covariance matrix $\Sigma$ is still Gaussian with mean $A\vec{\mu}$ and covariance matrix $A\Sigma A^*$. This implies that the DFT of $\vec{\mathbf{z}}$ is a Gaussian vector with covariance matrix $F_{[N]}\sigma^2 I F_{[N]}^* = N\sigma^2 I$ and mean zero: in other words, the Fourier coefficients of the noise are also iid Gaussian with zero mean and variance $N\sigma^2$. Here we are considering Gaussian random variables, which are complex-valued random values whose real and imaginary parts have a Gaussian distribution.

The numerator in Eq. (144) can be interpreted as a sample crosscorrelation of the Fourier coefficients of the data and the signal. By the independence assumption we have

$$\frac{1}{n}\sum_{j=1}^{n} \hat{x}^{[j]}[k]\overline{\hat{y}^{[j]}[k]} \approx \mathrm{E}\left(\hat{\boldsymbol{x}}[k]\overline{\hat{\boldsymbol{y}}[k]}\right) \tag{138}$$

$$= \mathrm{E}\left(\left|\hat{\boldsymbol{x}}[k]\right|^2\right) + \mathrm{E}\left(\hat{\boldsymbol{x}}[k]\overline{\hat{\boldsymbol{z}}[k]}\right) \tag{139}$$

$$= \mathrm{E}\left(\left|\hat{\boldsymbol{x}}[k]\right|^2\right) + \mathrm{E}\left(\hat{\boldsymbol{x}}[k]\right)\mathrm{E}\left(\overline{\hat{\boldsymbol{z}}[k]}\right) \tag{140}$$

$$= \mathrm{E}\left(\left|\hat{\boldsymbol{x}}[k]\right|^2\right). \tag{141}$$

Similarly, the denominator is equal to the sample mean square of the data. Again by the independence assumption,

$$\frac{1}{n}\sum_{j=1}^{n} \left|\hat{y}^{[j]}[k]\right|^2 \approx \mathrm{E}\left(\left|\hat{\boldsymbol{y}}[k]\right|^2\right) \tag{142}$$

$$= \mathrm{E}\left(\left|\hat{\boldsymbol{x}}[k]\right|^2\right) + \mathrm{E}\left(\left|\hat{\boldsymbol{z}}[k]\right|^2\right), \tag{143}$$

where the cross term is zero because $\hat{\boldsymbol{z}}$ is zero mean.

We conclude that the transfer function of the optimal denoising Wiener filter is given by

$$\hat{w}_{\mathrm{opt}}[k] = \frac{\mathrm{E}\left(\left|\hat{\boldsymbol{x}}[k]\right|^2\right)}{\mathrm{E}\left(\left|\hat{\boldsymbol{x}}[k]\right|^2\right) + N\sigma^2}, \quad 1 \leq k \leq N. \tag{144}$$

In 2D, by the same argument

$$\widehat{W}_{\mathrm{opt}}[k_1, k_2] = \frac{\mathrm{E}\left(\left|\widehat{\boldsymbol{X}}[k_1, k_2]\right|^2\right)}{\mathrm{E}\left(\left|\widehat{\boldsymbol{X}}[k_1, k_2]\right|^2\right) + N^2\sigma^2}, \quad 1 \leq k_1, k_2 \leq N. \tag{145}$$

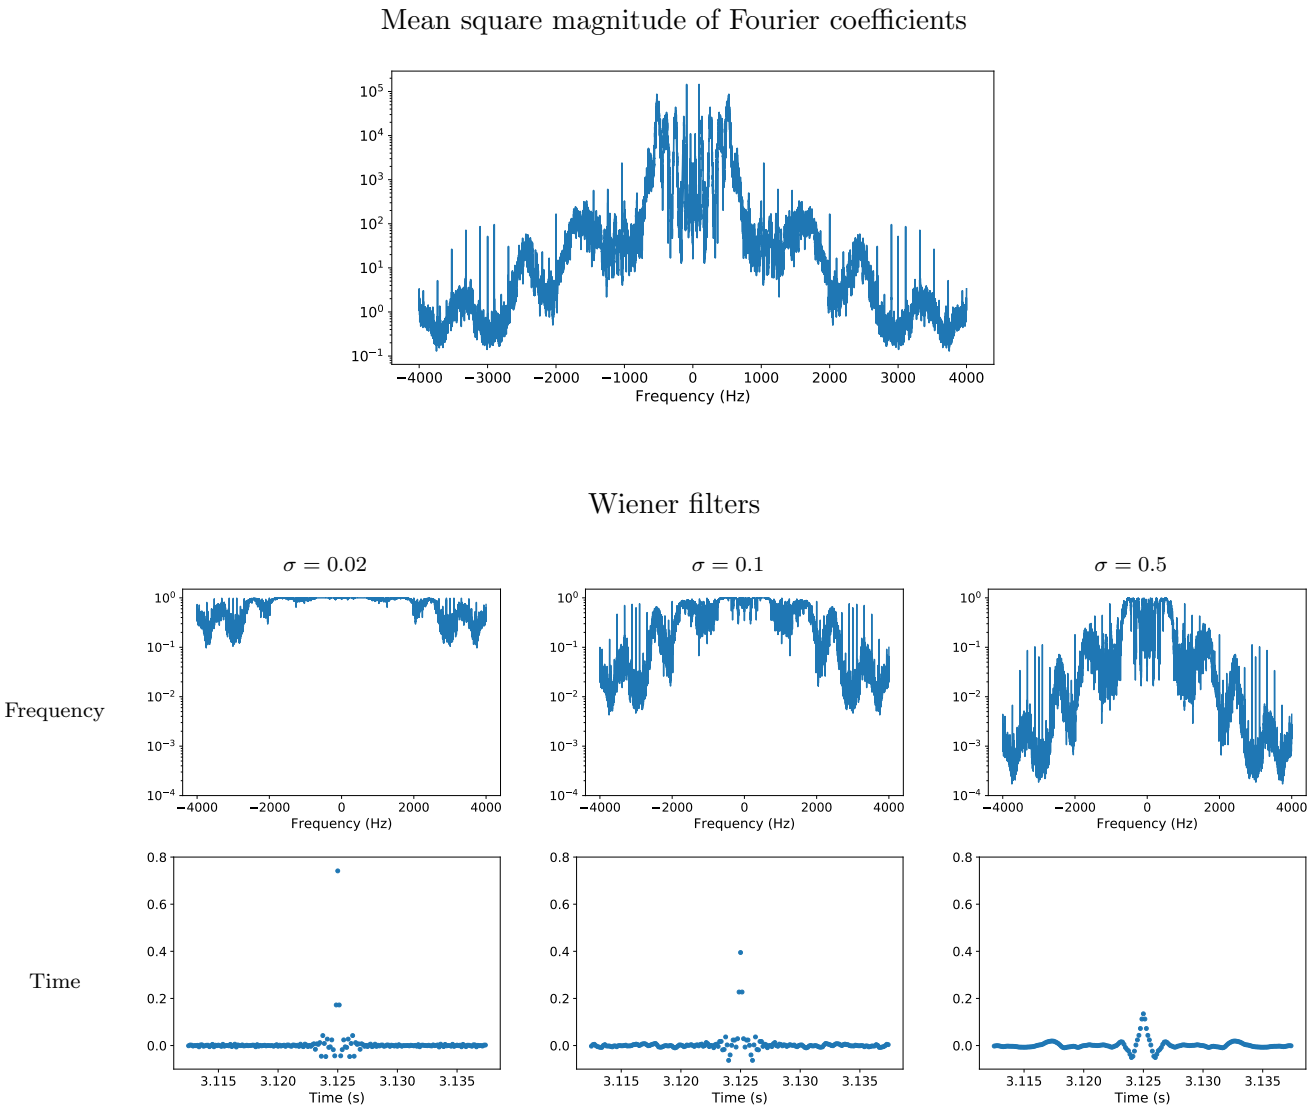


Figure 18: The plot at the top shows the mean square magnitude of the Fourier coefficients of a set of audio signals, in which people say *yes* and *no* alternatively in Hebrew. The corresponding Wiener filters for different noise levels are shown below in the frequency and time domains.

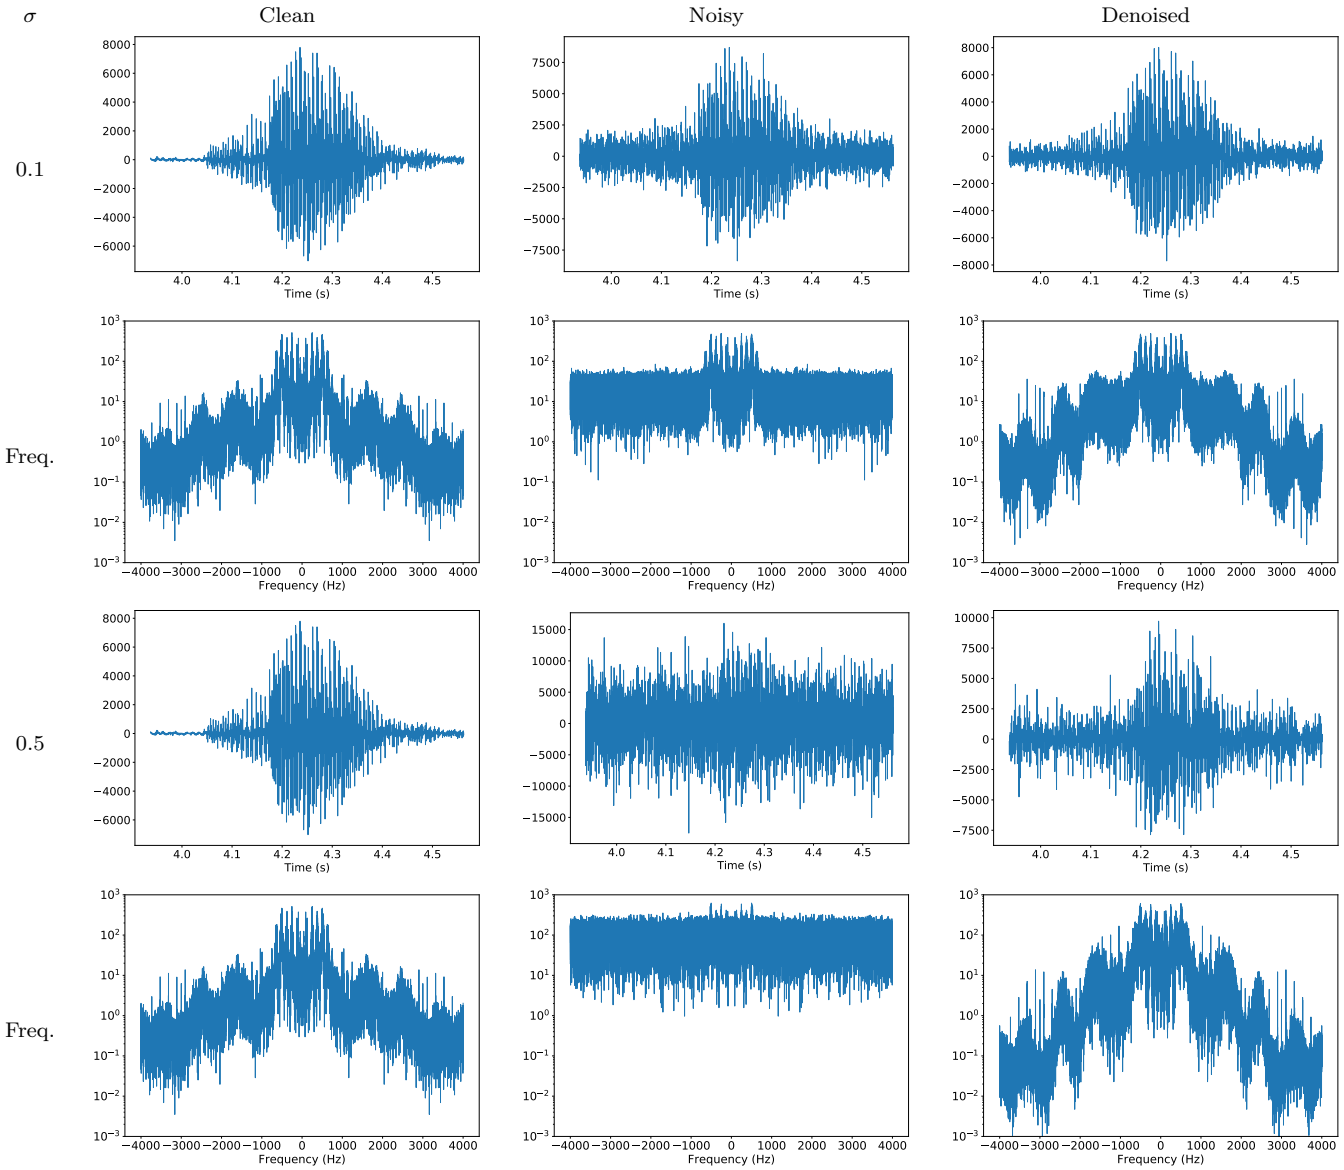


Figure 19: Result of applying the filters in Figure 18 to two test audio signals with different noise levels. Click on these links to listen to the audio: clean signal, noisy ($\sigma = 0.1$), denoised ($\sigma = 0.1$), noisy ($\sigma = 0.5$), denoised ($\sigma = 0.5$).

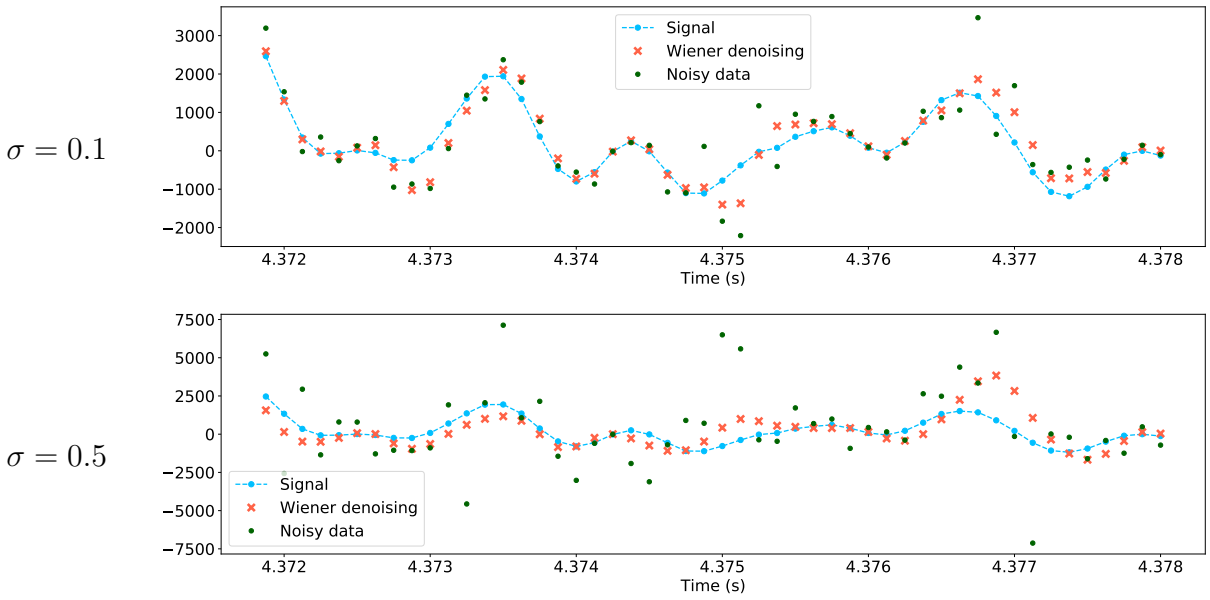


Figure 20: Zoomed plots of the denoising results shown in Figure 19.

Figures 18 and 21 show the mean square magnitude of the Fourier coefficients of a set of signals, and the corresponding Wiener filters for different noise levels. When the noise level is low, the filter is concentrated at the origin, which means that each entry of the estimated signal is almost equal to the noisy measurement. As the noise rises, the filter suppresses the frequencies at which the signals in the training set have less energy. For audio and natural images, this usually results in filtering out high frequencies. In the time domain, this results in a filter that averages over more neighboring entries to produce the estimate. This is apparent in the denoising examples shown in Figures 19, 20 and 19. △

## 4.4 PCA of translation-invariant signals

In this section, we consider performing PCA on translation-invariant signals. A simple example, are piecewise-constant one-dimensional signals. Figure 23 shows the results of applying PCA to a set of such signals generated at random. It turns out that the principal directions are sinusoidal! To understand why, we need to study the covariance matrices of translation-invariant signals.

Up to now, we have not made mathematically precise what we mean by translation-invariant signals. A popular definition in signal processing is that such signals are *stationary*, meaning that their statistics are translation-invariant.

**Definition 4.13.** *An $N$-dimensional random vector $\vec{\mathbf{x}}$ is wide-sense or weak-sense stationary if it satisfies two conditions. First, it has a constant mean, i.e. there exists a constant $\mu$ such that*

$$\mathrm{E}\left(\vec{\mathbf{x}}[j]\right) = \mu, \quad 1 \leq j \leq N. \tag{146}$$

*Second, it has a translation-invariant covariance; there exists a function $\kappa$ such that*

$$\mathrm{E}\left(\vec{\mathbf{x}}[j_1]\vec{\mathbf{x}}[j_2]\right) = \kappa(|j_2 - j_1|), \quad 1 \leq j_1, j_2 \leq N. \tag{147}$$

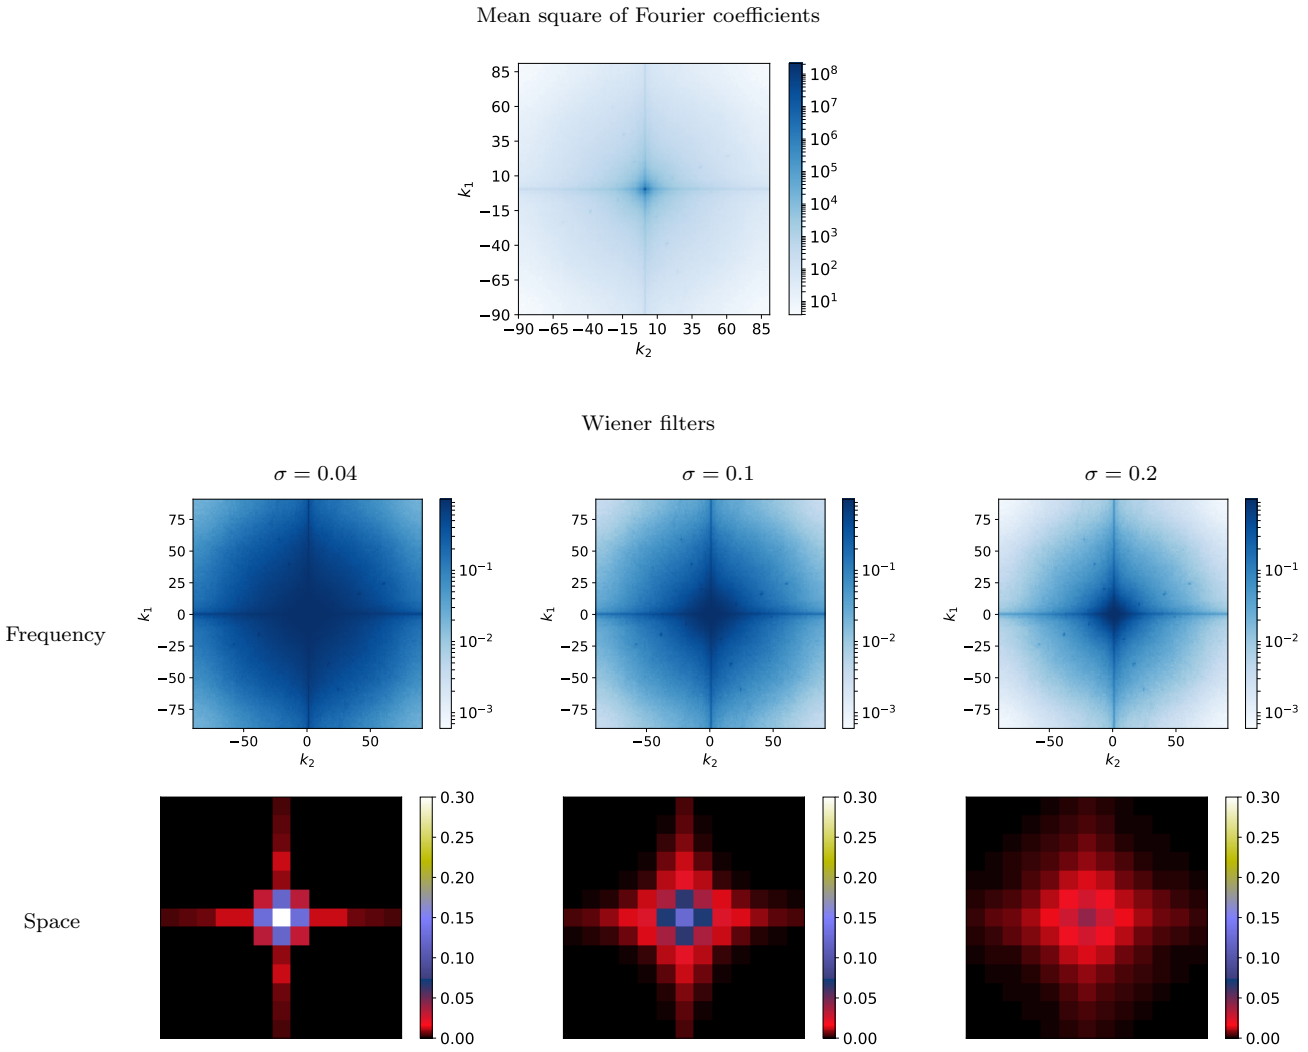


Figure 21: The heatmap at the top shows the mean square magnitude of the Fourier coefficients of a set of natural images. The corresponding Wiener filters for different noise levels are shown below in the frequency and spatial domains.

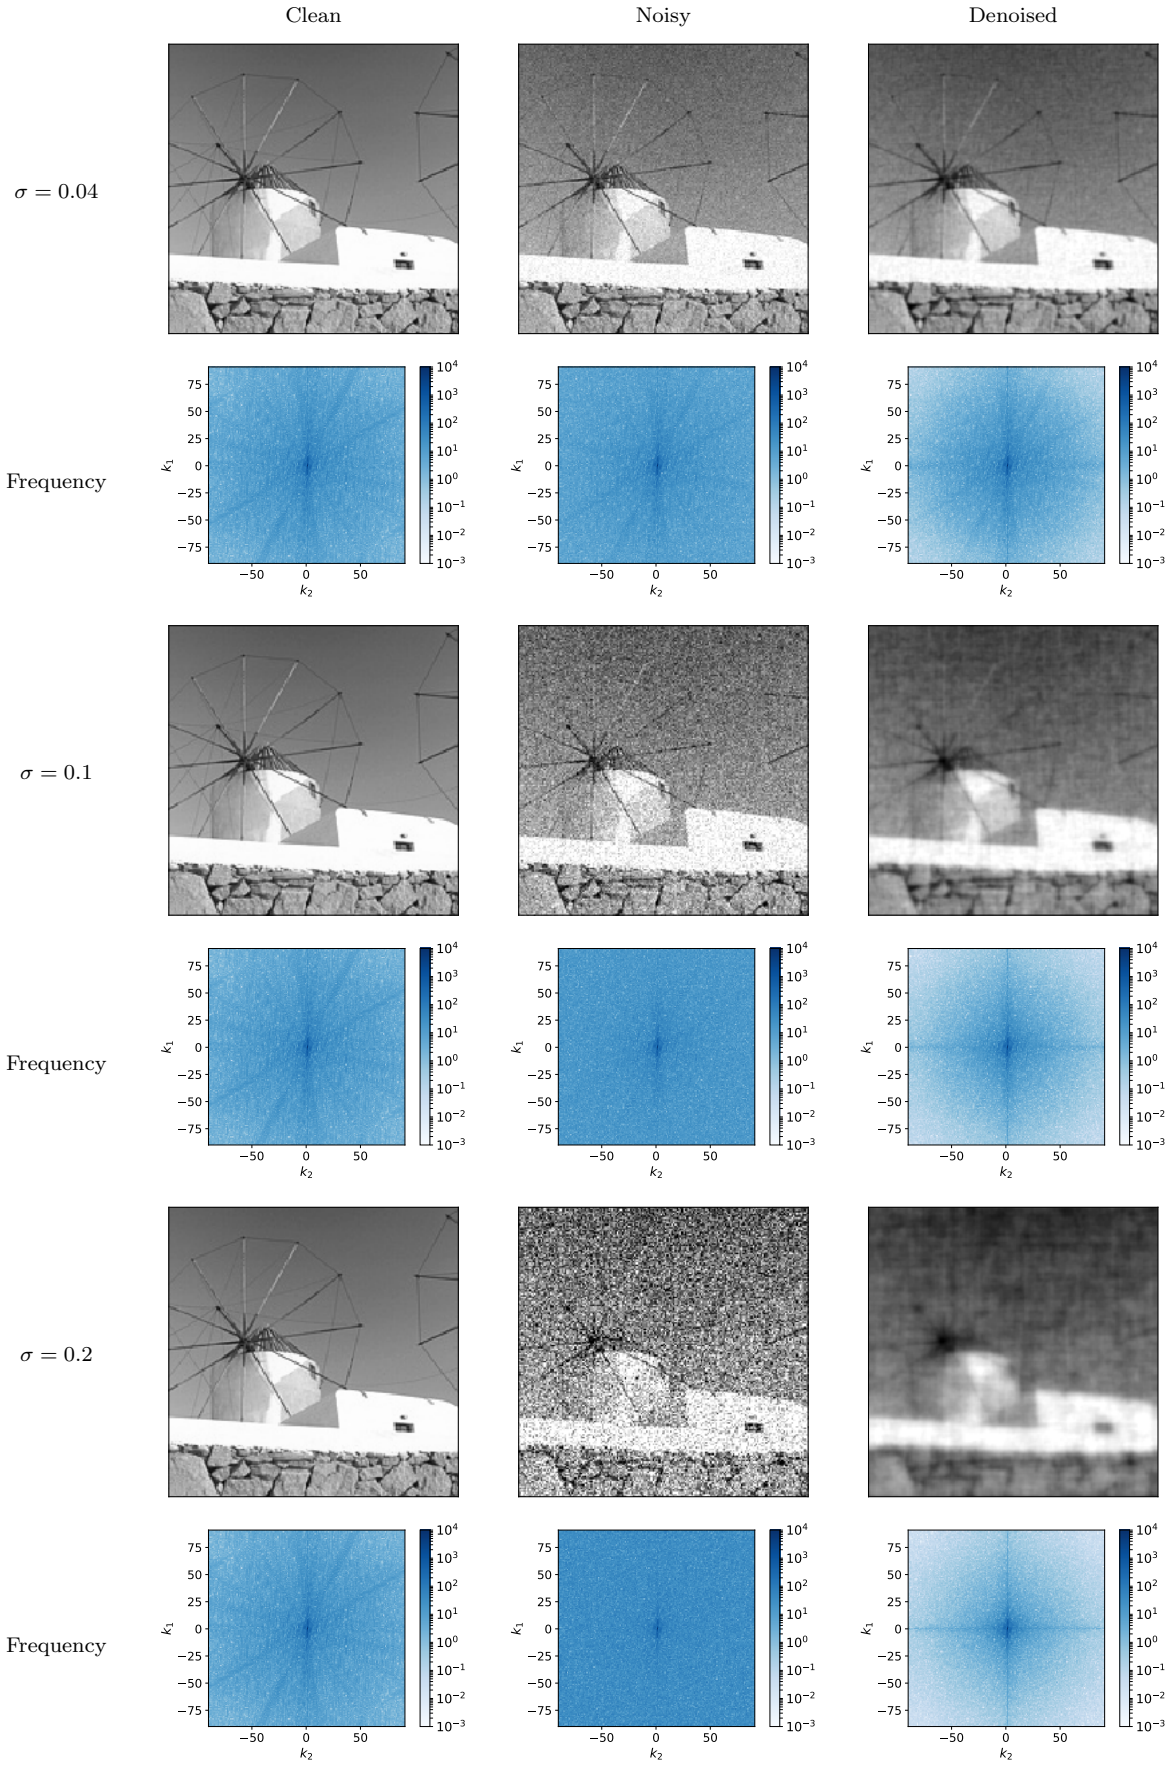


Figure 22: Result of applying the filters in Figure 21 to a test image with different noise levels.

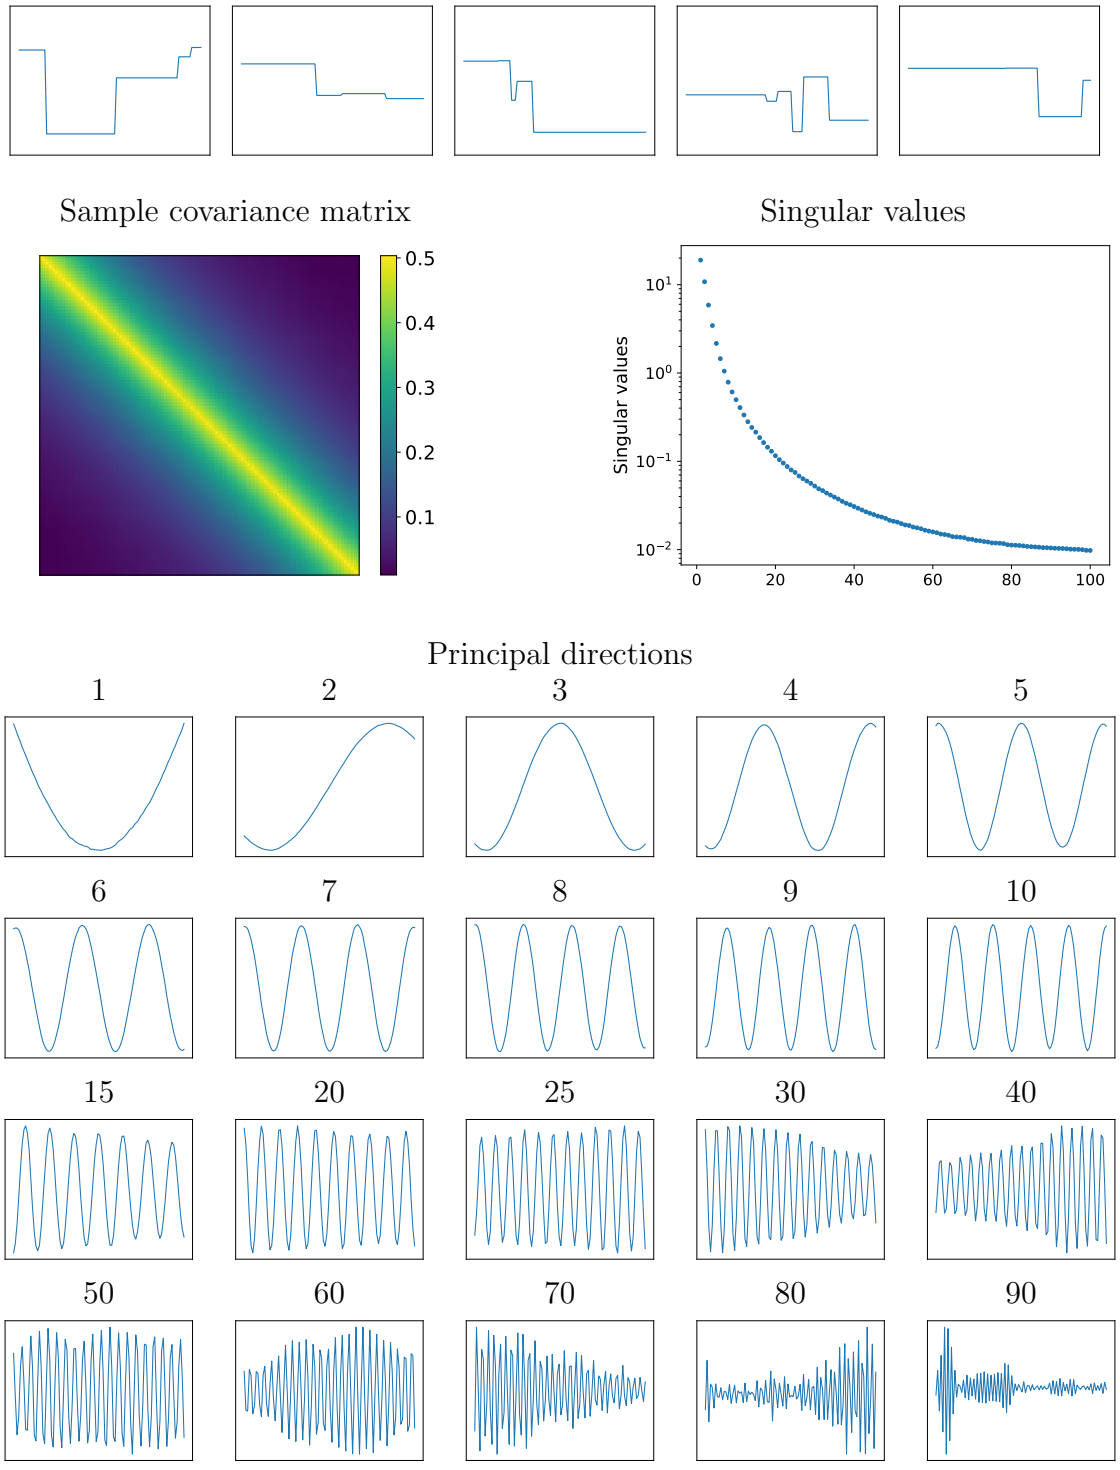


Figure 23: The top row shows examples of piecewise constant signals. The heatmap below displays the sample covariance matrix of a set of such signals. The graph to the right of the heatmap shows the singular values of the sample covariance matrix. Finally, the images at the bottom show the singular vectors, which represent the principal directions of the data.

*The function $\kappa$ is known as the autocovariance function of the random vector.*

The covariance matrix of a wide-sense stationary random vector is Toeplitz symmetric. For instance, for $N = 4$ it is of the form.

$$\Sigma = \begin{bmatrix} \kappa(0) & \kappa(1) & \kappa(2) & \kappa(3) \\ \kappa(1) & \kappa(0) & \kappa(1) & \kappa(2) \\ \kappa(2) & \kappa(1) & \kappa(0) & \kappa(1) \\ \kappa(3) & \kappa(2) & \kappa(1) & \kappa(0) \end{bmatrix}. \tag{148}$$

In real signals, the autocovariance function $\kappa$ often decays at large separations because the entries become uncorrelated (think for example of pixels in an image, which are far apart). In that case, the covariance matrix can be well approximated by a circulant matrix.

**Definition 4.14** (Circulant matrix)**.** *A circulant matrix is a Toeplitz matrix where each row vector is obtained by applying a unit circular shift to the previous row. For example,*

$$\begin{bmatrix} a & b & c & d \\ d & a & b & c \\ c & d & a & b \\ b & c & d & a \end{bmatrix}, \tag{149}$$

*is a circulant matrix.*

To illustrate the similarity of Toeplitz symmetric and circulant matrices when the rows decay, let $N = 5$ and assume that $\kappa(j)$ is negligeable for $j > 1$. The covariance matrix

$$\Sigma = \begin{bmatrix} \kappa(0) & \kappa(1) & 0 & 0 & 0 \\ \kappa(1) & \kappa(0) & \kappa(1) & 0 & 0 \\ 0 & \kappa(1) & \kappa(0) & \kappa(1) & 0 \\ 0 & 0 & \kappa(1) & \kappa(0) & \kappa(1) \\ 0 & 0 & 0 & \kappa(1) & \kappa(0) \end{bmatrix}. \tag{150}$$

is quite close to the circulant matrix

$$\begin{bmatrix} \kappa(0) & \kappa(1) & 0 & 0 & \kappa(1) \\ \kappa(1) & \kappa(0) & \kappa(1) & 0 & 0 \\ 0 & \kappa(1) & \kappa(0) & \kappa(1) & 0 \\ 0 & 0 & \kappa(1) & \kappa(0) & \kappa(1) \\ \kappa(1) & 0 & 0 & \kappa(1) & \kappa(0) \end{bmatrix}. \tag{151}$$

The following theorem establishes that the eigenvectors of circulant matrices are sinusoidal.

**Theorem 4.15.** *Any circulant matrix $C \in \mathbb{C}^{n \times n}$ can be written as*

$$C := \frac{1}{N} F_{[N]}^* \Lambda F_{[N]}, \tag{152}$$

*where $F_{[N]}$ is the DFT matrix and $\Lambda$ is a diagonal matrix.*

*Proof.* Let $\vec{c}$ be equal to the first column of the circulant matrix. By the definition of circular convolution, for any vector $\vec{x} \in \mathbb{C}^n$

$$C\vec{x} = \vec{c} * \vec{x} \tag{153}$$

$$= \frac{1}{N} F_{[N]}^* \operatorname{diag}(\hat{c}) F_{[N]} \vec{x}, \tag{154}$$

where the last step follows from Theorem 4.9 and $\operatorname{diag}(\hat{c})$ is a diagonal matrix that contains the DFT coefficients of $\vec{c}$. This implies $C = \frac{1}{N} F_{[N]}^* \operatorname{diag}(\hat{c}) F_{[N]}$. □

In words, the eigenvalues of circulant matrices are equal to the DFT coefficients of their rows, and the eigenvectors equal complex sinusoids of corresponding frequencies. When the circulant matrix is a covariance matrix, the eigenvalues are nonnegative, as we proved in the lecture notes on the SVD. As a result the SVD of the matrix equals $\frac{1}{\sqrt{N}} F_{[N]}^* \operatorname{diag}(\hat{c}) \frac{1}{\sqrt{N}} F_{[N]}$. In particular, if the Fourier coefficients of $\vec{c}$ have different magnitudes, the principal directions are sinusoids! This explains why covariance matrices that are *almost* circulant, like the one in Figure 23 tend to have sinusoidal singular vectors. We refer to the seminal review by Gray [1] for a more in-depth discussion.

The same reasoning can be used to analyze covariance matrices obtained from translation-invariant 2D signals, which are vectorized. In that case, the rows of the matrix have additional structure, they can be reshaped into a 2D array that corresponds to the shape of the signals. Figure 24 shows an example obtained using 20x20 patches extracted from natural images. If the covariance between pixels decays when the pixels are far from each other, the covariance matrix of such signals is close to being equivalent to a 2D convolution matrix, which has 2D sinusoids as its eigenvectors. This suggests that applying PCA to patches of natural images should yields sinusoidal principal directions. Figure 25 confirms this prediction for the data in Figure 24.

An important consequence is that sinusoids are guaranteed to be close to optimal in terms of dimensionality reduction for natural images. This is the reason that the JPEG compression standard uses the discrete cosine transform (DCT) to compress images. The DCT is a variant of the discrete Fourier transform for real signals, which assumes that the signal is symmetric (but only one half is observed). In that case the DFT components correspond to discrete cosines (the coefficients corresponding to the sines are zero due to symmetry). Figure 26 shows the basis vectors of the $8 \times 8$ 2D DCT[2].

Figure 27 shows the result of dividing a natural image into 64-pixel patches and projecting it onto the span of the low-frequency DCT basis functions. Ignoring some of the high-frequency components is almost imperceptible. This is the main insight behind the JPEG method for lossy compression of digital images. JPEG divides the image in $8 \times 8$ patches and then quantizes the coefficients of each DCT band differently. More bits are used for lower-frequency bands because errors in these bands are much more apparent to the human eye.

# References

[1] R. M. Gray. Toeplitz and circulant matrices: A review. *Foundations and Trends in Communications and Information Theory*, 2(3):155–239, 2006.

[2]The image is borrowed from https://upload.wikimedia.org/wikipedia/commons/2/24/DCT-8x8.png

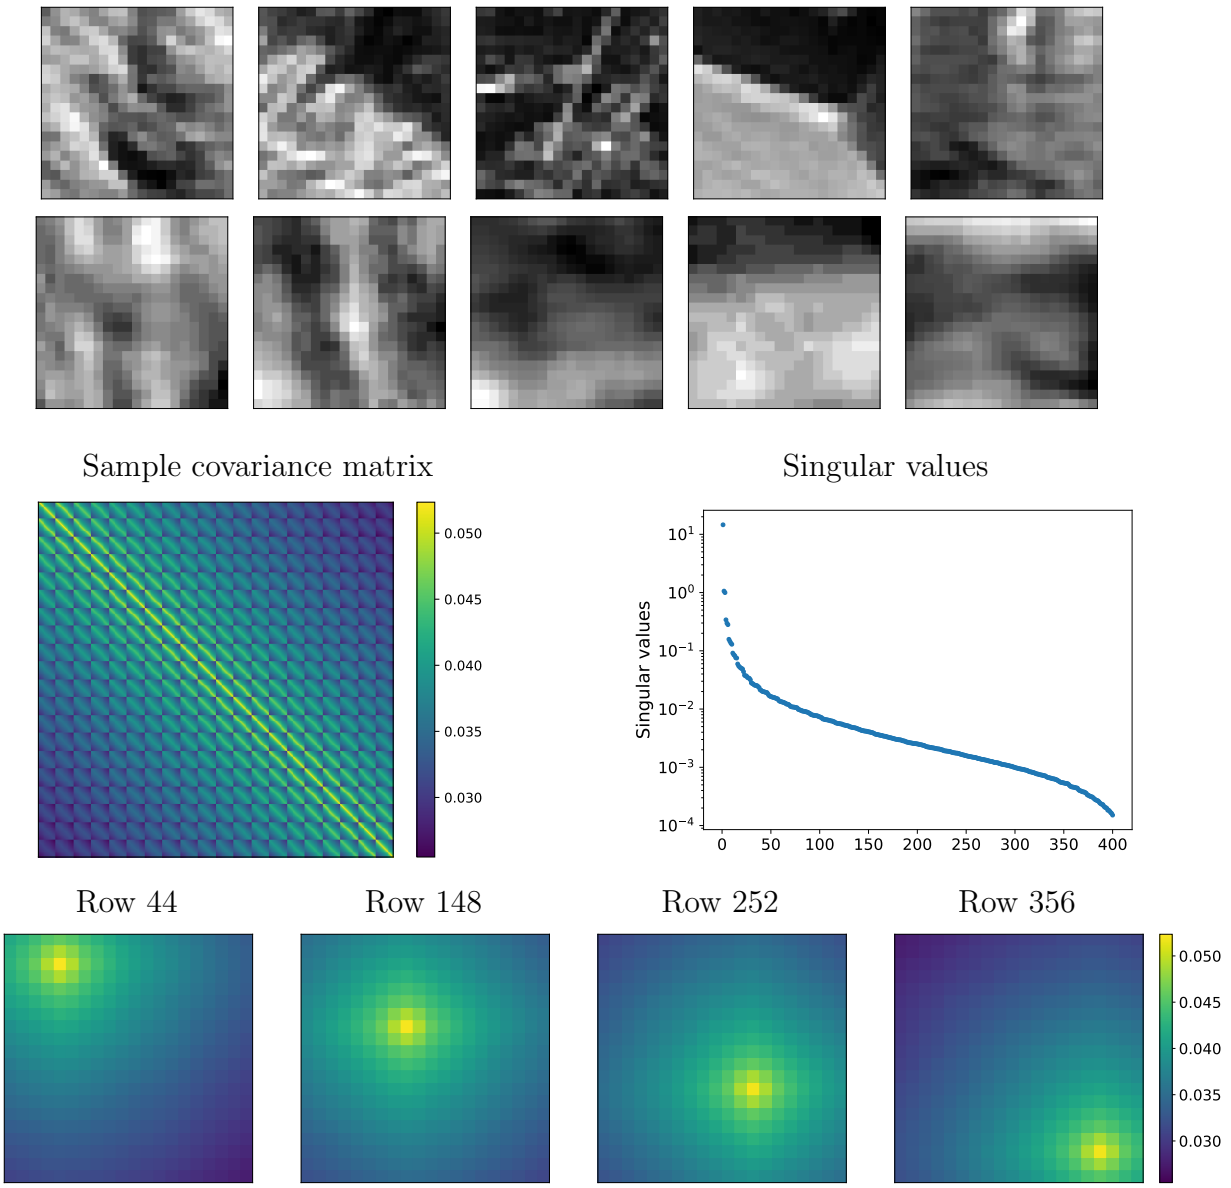


Figure 24: The top row shows examples of image patches. The heatmap below displays the sample covariance matrix of patches sampled from a set of natural images. The graph to the right of the heatmap shows the singular values of the sample covariance matrix. The bottom row shows rows of the covariance matrix reshaped to show the values corresponding to each pixel in a patch.

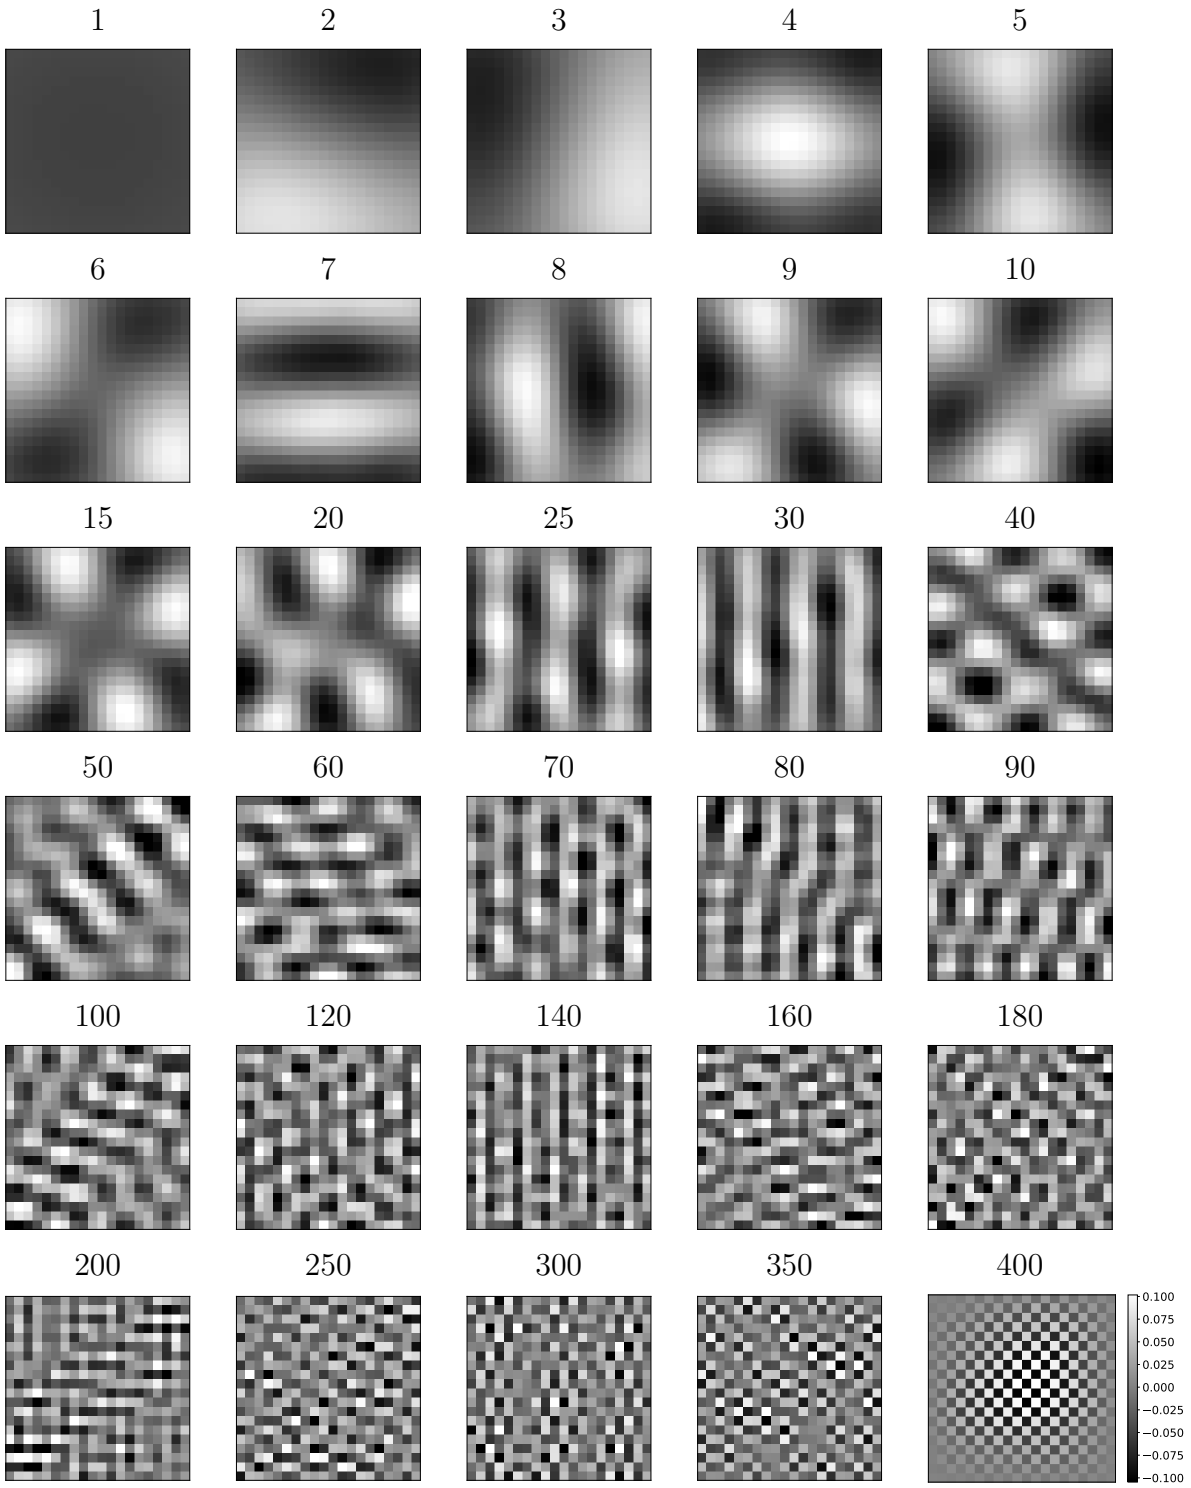


Figure 25: Singular vectors of the sample covariance matrix in Figure 24 reshaped to show the values corresponding to each pixel. These are the principal directions of the dataset of image patches.

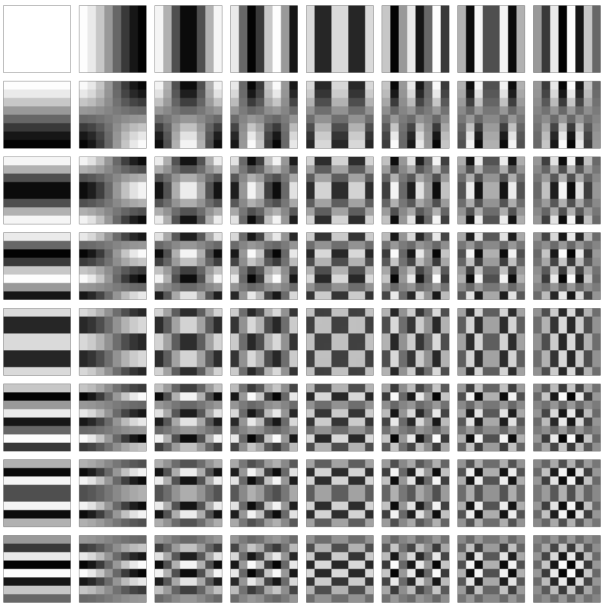

Figure 26: $8 \times 8$ DCT basis vectors.

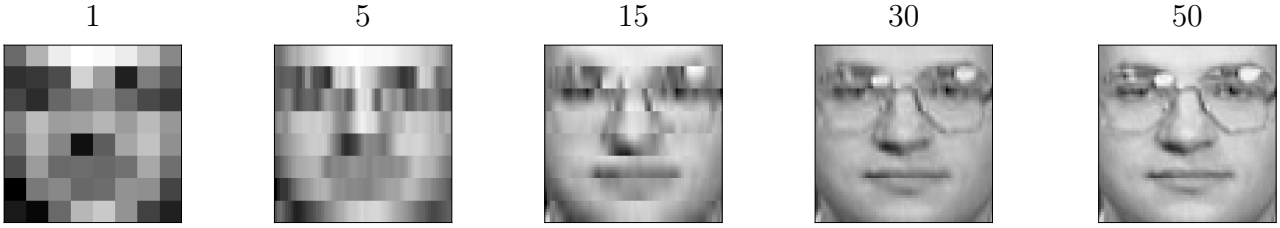


Figure 27: Result of projecting each 64-pixel patch from the image of a face onto the lowest 1, 5, 15, 30 and 50 2D DCT basis functions.