

2D Fourier Transforms

In 2D, for signals $h(n, m)$ with N columns and M rows, the idea is exactly the same:

$$\hat{h}(k, l) = \sum_{n=0}^{N-1} \sum_{m=0}^{M-1} e^{-i(\omega_k n + \omega_l m)} h(n, m)$$
$$h(n, m) = \frac{1}{NM} \sum_{k=0}^{N-1} \sum_{l=0}^{M-1} e^{i(\omega_k n + \omega_l m)} \hat{h}(k, l)$$



Often it is convenient to express frequency in vector notation with $\vec{k} = (k, l)^t$, $\vec{n} = (n, m)^t$, $\vec{\omega}_{kl} = (\omega_k, \omega_l)^t$ and $\vec{\omega}^t \vec{n} = \omega_k n + \omega_l m$.

2D Fourier Basis Functions: Sinusoidal waveforms of different wavelengths (scales) and orientations. Sinusoids on $N \times M$ images with 2D frequency $\vec{\omega}_{kl} = (\omega_k, \omega_l) = 2\pi(k/N, l/M)$ are given by:


$$e^{i(\vec{\omega}^t \vec{n})} = e^{i\omega_k n} e^{i\omega_l m} = \cos(\vec{\omega}^t \vec{n}) + i \sin(\vec{\omega}^t \vec{n})$$

Separability: If $h(\vec{n})$ is separable, e.g., $h(n, m) = f(n)g(m)$, then, because complex exponentials are also separable, so is the Fourier spectrum, $\hat{h}(k, l) = \hat{f}(k) \hat{g}(l)$.

2D Fourier Basis Functions



Blocks image and its amplitude spectrum



Properties of the Fourier Transform

Some key properties of the Fourier transform, $\hat{f}(\vec{\omega}) = \mathcal{F}[f(\vec{x})]$.

Symmetries:

For $s(x) \in \mathcal{R}$, the Fourier transform is symmetric, i.e., $\hat{s}(\omega) = \hat{s}^*(-\omega)$.

For $s(x) = s(-x)$ the transform is real-valued, i.e., $\hat{s}(\omega) \in \mathcal{R}$.

For $s(x) = -s(-x)$ the transform is imaginary, i.e., $i\hat{s}(\omega) \in \mathcal{R}$.

Shift Property:

$$\mathcal{F}[f(\vec{x} - \vec{x}_0)] = \exp(-i\vec{\omega}^t \vec{x}_0) \hat{f}(\vec{\omega}) \quad (1)$$

The amplitude spectrum is invariant to translation. The phase spectrum is not. In particular, note that $\mathcal{F}[\delta(\vec{x} - \vec{x}_0)] = \exp(-i\vec{\omega}^t \vec{x}_0)$.

Proof: substitution and change of variables.

Differentiation:

$$\mathcal{F}\left[\frac{\partial^n f(\vec{x})}{\partial x_j^n}\right] = (i\omega_j)^n \hat{f}(\vec{\omega}) \quad (2)$$

For intuition, remember that $\frac{\partial e^{i\omega x}}{\partial x} = i\omega e^{i\omega x}$ and $\frac{\partial \sin(\omega x)}{\partial x} = \omega \cos(\omega x)$.

Linear Scaling: Scaling the signal domain causes scaling of the Fourier domain; i.e., given $a \in \mathcal{R}$, $\mathcal{F}[s(ax)] = \frac{1}{a}\hat{s}(\omega/a)$.

Parseval's Theorem: Sum of squared Fourier coefficients is a constant multiple of the sum of squared signal values.

Convolution Theorem

The Fourier transform of the convolution of two signals is equal to the product of their Fourier transforms:

$$\mathcal{F}[f * g] = \mathcal{F}[f] \mathcal{F}[g] \equiv \hat{f}(\omega) \hat{g}(\omega) . \quad (3)$$

Proof in the discrete 1D case:






$$\begin{aligned} \mathcal{F}[f * g] &= \sum_n f * g e^{-i\omega n} = \sum_n \sum_m f(m) g(n - m) e^{-i\omega n} \\ &= \sum_m f(m) \sum_n g(n - m) e^{-i\omega n} \\ &= \sum_m f(m) \hat{g}(\omega) e^{-i\omega m} \quad (\text{shift property}) \\ &= \hat{g}(\omega) \hat{f}(\omega) . \end{aligned}$$

Remarks:

- This theorem means that one can apply filters efficiently in the Fourier domain, with multiplication instead of convolution.
- Fourier spectra help characterize how different filters behave, by expressing both the impulse response and the signal in the Fourier domain (e.g, with the DTFT). The filter's amplitude spectrum tells us how each signal frequency will be attenuated. The filter's phase spectrum tells us how each sinusoidal signal component will be phase shifted in the response.
- Convolution theorem also helps prove properties. E.g. prove:

$$\frac{\partial}{\partial x} (h * g) = \frac{\partial h}{\partial x} * g = h * \frac{\partial g}{\partial x}$$



Common Filters and their Spectra





Top Row: Image of Al and a **low-pass** (blurred) version of it. The low-pass kernel was separable, composed of 5-tap 1D impulse responses $\frac{1}{16}(1, 4, 6, 4, 1)$ in the x and y directions.

Bottom Row: From left to right are the amplitude spectrum of Al, the amplitude spectrum of the impulse response, and the product of the two amplitude spectra, which is the amplitude spectrum of the blurred version of Al. (Brightness in the left and right images is proportional to log amplitude.)

Common Filters and their Spectra (cont)







From left to right is the original AI, a **high-pass** filtered version of AI, and the amplitude spectrum of the filter. This impulse response is defined by $\delta(n) - h(n, m)$ where $h[n, m]$ is the separable blurring kernel used in the previous figure.



From left to right is the original AI, a **band-pass** filtered version of AI, and the amplitude spectrum of the filter. This impulse response is defined by the difference of two low-pass filters.

Common Filters and their Spectra (cont)



Top Row: Convolution of AI with a horizontal derivative filter, along with the filter's Fourier spectrum. The 2D separable filter is composed of a vertical smoothing filter (i.e., $\frac{1}{4}(1, 2, 1)$) and a first-order central difference (i.e., $\frac{1}{2}(-1, 0, 1)$) horizontally.

Bottom Row: Convolution of AI with a vertical derivative filter, and the filter's Fourier spectrum. The filter is composed of a horizontal smoothing filter and a vertical first-order central difference.

Nyquist Sampling Theorem


Theorem: Let $f(x)$ be a band-limited signal such that

$$\hat{f}(\omega) = 0 \quad \text{for } |\omega| > \omega_0$$

for some ω_0 . Then $f(x)$ is uniquely determined by its samples $g(m) = f(m n_s)$ when

$$\frac{2\pi}{n_s} > 2\omega_0 \quad \text{or equivalently} \quad n_s < \frac{\lambda_0}{2}$$

where $\lambda_0 = 2\pi/\omega_0$. In words, the distance between samples must be smaller than half a wavelength of the highest frequency in the signal.




Here the replicas can be isolated by an ideal low-pass filter (the dotted pass-band), so the original signal can be perfectly reconstructed.

Corollary: Let $f(x)$ be a single-sided band-pass signal with bandwidth $2\omega_0$. Then $f(x)$ is uniquely determined if sampled at a rate such that $n_s < \frac{\lambda_0}{2}$.

Aliasing



Aliasing occurs when replicas overlap:



Consider a perspective image of an infinite checkerboard. The signal is dominated by high frequencies in the image near the horizon. Properly designed cameras blur the signal before sampling, using

- the point spread function due to diffraction,
- imperfect focus,
- averaging the signal over each CCD element.

These operations attenuate high frequency components in the signal. Without this (physical) preprocessing, the sampled image can be severely aliased (corrupted):



Dimensionality

A guiding principal throughout signal transforms, sampling, and aliasing is the underlying dimension of the signal, that is, the number of linearly independent degrees of freedom (dof). This helps clarify many issues that might otherwise appear mysterious.

- Real-valued signals with N samples have N dof. We need a basis of dimension N to represent them uniquely.
- Why did the DFT of a signal of length N use N sinusoids? Because N sinusoids are linearly independent, providing a minimal spanning set for signals of length N . We need no more than N .
- But wait: Fourier coefficients are complex-valued, and therefore have $2N$ dofs. This matches the dof needed for complex signals of length N but not real-valued signals. For real signals the Fourier spectra are symmetric, so we keep half of the coefficients.
- When we down-sample a signal by a factor of two we are moving to a basis with $N/2$ dimensions. The Nyquist theorem says that the original signal should lie in an $N/2$ dimensional space before you down-sample. Otherwise information is corrupted (i.e. signal structure in multiple dimensions of the original N -D space appear the same in the $N/2$ -D space).
- The Nyquist theorem is not primarily about highest frequencies and bandwidth. The issue is really one of having a model for the signal; that is, how many non-zero frequency components are in the signal (i.e., the dofs), and which frequencies are they.