


A gentle introduction to Fourier analysis




[Image source](#)

Many slides borrowed from S. Seitz, A. Efros, D. Hoiem, B. Freeman, A. Zisserman

Mystery 1


- Why can downsampling sometimes lead to aliasing?



Mystery 2

- Why does filtering with a Gaussian give a nice smooth image, but filtering with a box filter gives artifacts?

Gaussian




Box filter



Mystery 3

- How do hybrid images work?

“Low frequencies”




“High frequencies”

A. Oliva et al. [Hybrid Images](#). SIGGRAPH 2006

Fourier analysis

- To understand such phenomena, we need a representation of images that allows us to tease apart slow and fast changes



Outline

- 1D Fourier transform
 - Definition and properties
 - Discrete Fourier transform
- 2D Fourier transform
 - Definition
 - Examples and properties
- Convolution theorem
- Understanding the behavior of filters
- Understanding the sampling theorem
- More applications of Fourier analysis

Fourier analysis


- Any “reasonable” univariate function can be expressed as a weighted sum of sinusoids of different frequencies (1807)



Jean-Baptiste Joseph Fourier (1768-1830)

Example: series for a square wave

$$\sum_{k=1,3,5,\dots}^{\infty} \frac{1}{k} \sin(kt)$$



Fourier analysis


- Our building block:

$$A \sin(\omega t + \phi)$$

Diagram illustrating the components of the sine wave equation $A \sin(\omega t + \phi)$:

- A : Amplitude (indicated by a red arrow pointing to A)
- ω : Frequency (indicated by a red arrow pointing to ω)
- ϕ : Phase (indicated by a red arrow pointing to ϕ)

- Add enough of these to get any signal you want!



1D Fourier transform

- Let's define basis functions parameterized by $u \in (-\infty, \infty)$:

$$\psi_u(t) = e^{i2\pi ut}$$



- What the...?

1D Fourier transform

- Let's define basis functions parameterized by $u \in (-\infty, \infty)$:

$$\psi_u(t) = e^{i2\pi ut}$$

- Euler's formula: $e^{i\theta} = \cos \theta + i \sin \theta$



[Image source](#)

1D Fourier transform

- Let's define basis functions parameterized by $u \in (-\infty, \infty)$:

$$\psi_u(t) = e^{i2\pi ut}$$

- Euler's formula: $e^{i\theta} = \cos \theta + i \sin \theta$
- So $\psi_u(t)$ is just a cosine-sine pair at frequency $2\pi u$ (or period $1/u$)

1D Fourier transform

- Let's define basis functions parameterized by $u \in (-\infty, \infty)$:

$$\psi_u(t) = e^{i2\pi ut}$$

- Inner product for complex functions is given by:

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

Complex conjugate:
real part stays the same,
imaginary part is flipped

1D Fourier transform

- Let's define basis functions parameterized by $u \in (-\infty, \infty)$:

$$\psi_u(t) = e^{i2\pi ut}$$

- Inner product for complex functions is given by:

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

- Orthonormality:

$$\langle \psi_{u_1}, \psi_{u_2} \rangle = \begin{cases} 1 & \text{if } u_1 = u_2 \\ 0 & \text{otherwise} \end{cases}$$

1D Fourier transform

- Given a signal $f(t)$, we want to represent it as a weighted combination of the basis functions $\psi_u(t) = e^{i2\pi ut}$ with weights $F(u)$:

$$f(t) = \int_{-\infty}^{\infty} F(u) e^{i2\pi ut} du$$

- Each weight $F(u)$ is given by the inner product of f and ψ_u :

$$F(u) = \langle f, \psi_u \rangle = \int_{-\infty}^{\infty} f(t) e^{-i2\pi ut} dt$$

1D Fourier transform

- Forward transform: $f(t) \xrightarrow{\mathcal{F}} F(u)$

$$F(u) = \int_{-\infty}^{\infty} f(t) e^{-i2\pi ut} dt$$

- Note: for the FT to exist, the *energy* $\int_{-\infty}^{\infty} |f(t)|^2 dt$ has to be finite

1D Fourier transform

- Forward transform: $f(t) \xrightarrow{\mathcal{F}} F(u)$

$$F(u) = \int_{-\infty}^{\infty} f(t) e^{-i2\pi ut} dt$$

- For each u , $F(u)$ is a *complex number* that encodes both the *amplitude* A and *phase* ϕ of the sinusoid $A \sin(2\pi ut + \phi)$ in the decomposition of $f(t)$:

$$F(u) = \operatorname{Re}(F(u)) + i \operatorname{Im}(F(u)),$$

$$A = \sqrt{\operatorname{Re}(F(u))^2 + \operatorname{Im}(F(u))^2}, \quad \phi = \tan^{-1} \frac{\operatorname{Im}(F(u))}{\operatorname{Re}(F(u))}$$

- If $f(t)$ is real, then $\operatorname{Re}(F(u)) = \operatorname{Re}(F(-u))$,
 $\operatorname{Im}(F(u)) = -\operatorname{Im}(F(-u))$

1D Fourier transform

- Forward transform: $f(t) \xrightarrow{\mathcal{F}} F(u)$

$$F(u) = \int_{-\infty}^{\infty} f(t) e^{-i2\pi ut} dt$$

- Important properties:

- Energy preservation:

$$\int_{-\infty}^{\infty} |f(t)|^2 dt = \int_{-\infty}^{\infty} |F(u)|^2 du$$

- Linearity: $\mathcal{F}\{af_1 + bf_2\} = a\mathcal{F}\{f_1\} + b\mathcal{F}\{f_2\}$

1D Fourier transform


- Forward transform: $f(t) \xrightarrow{\mathcal{F}} F(u)$

$$F(u) = \int_{-\infty}^{\infty} f(t) e^{-i2\pi ut} dt$$



- Inverse transform: $F(u) \xrightarrow{\mathcal{F}^{-1}} f(t)$

$$f(t) = \int_{-\infty}^{\infty} F(u) e^{i2\pi ut} du$$





- Duality: if $f(t) \xrightarrow{\mathcal{F}} F(u)$, then $F(t) \xrightarrow{\mathcal{F}} f(-u)$
 - Thus, we can talk about *Fourier transform pairs* $f(t) \leftrightarrow F(u)$




Important Fourier transform pairs



Important Fourier transform pairs



Important Fourier transform pairs





*The last one is formal since these functions don't meet the mathematical requirements for FT

Outline

- 1D Fourier transform
 - Definition and properties
 - Discrete Fourier transform

Discrete Fourier transform (DFT)

- Now suppose our signal consists of N samples $f(n)$,
 $n = 0, \dots, N - 1$
- We can also discretize frequencies to k/N , $k = 0, \dots, N - 1$
(k cycles per N samples)



[Image source](#)

Discrete Fourier transform (DFT)


- Now suppose our signal consists of N samples $f(n)$,
 $n = 0, \dots, N - 1$
- We can also discretize frequencies to k/N , $k = 0, \dots, N - 1$
(k cycles per N samples)
- DFT formula:

$$F(k) = \sum_{n=0}^{N-1} f(n) \exp\left(-i \frac{2\pi k}{N} n\right)$$

- We can pack the values $\exp\left(-i \frac{2\pi k}{N} n\right)$, $k, n = 0, \dots, N - 1$ into an $N \times N$ matrix U , and DFT becomes just a matrix-vector multiplication!
- Fast Fourier transform: only $N \log N$ complexity!

DFT: Just a change of basis!

$$U f = F$$



Inverse DFT

- Forward DFT:

$$F(k) = \sum_{n=0}^{N-1} f(n) \exp\left(-i \frac{2\pi}{N} kn\right) \quad \text{or } F = Uf$$

- Inverse DFT:

$$f(n) = \frac{1}{N} \sum_{k=0}^{N-1} F(k) \exp\left(i \frac{2\pi}{N} kn\right) \quad \text{or } f = \frac{1}{N} U^{-1} F$$

where U^{-1} is the transpose of the *complex conjugate* of U

Periodicity of DFT and inverse DFT

- The result of DFT is periodic: because $F(k)$ is obtained as a sum of complex exponentials with a common period of N samples, $F(k + aN) = F(k)$ for any integer a :

$$\begin{aligned} F(k + aN) &= \sum_{n=0}^{N-1} f(n) \exp\left(-i \frac{2\pi}{N} n(k + aN)\right) \\ &= \sum_{n=0}^{N-1} f(n) \exp\left(-i \frac{2\pi n}{N} k\right) \exp(-i2\pi an) = F(k) \end{aligned}$$

- Likewise, the result of the inverse DFT is a periodic signal: $f(t + aN) = f(t)$ for any integer a

Outline

- 1D Fourier transform
 - Definition and properties
 - Discrete Fourier transform
- 2D Fourier transform

2D Fourier transform


- To represent 2D signals $f(x, y)$, we need to extend our 1D basis functions $\psi_u(t) = e^{i2\pi ut}$ to two variables:

$$\begin{aligned}\psi_{u,v}(x, y) &= e^{i2\pi ux} e^{i2\pi vy} = e^{i2\pi(ux+vy)} \\ &= \cos 2\pi(ux + vy) + i \sin 2\pi(ux + vy)\end{aligned}$$

- What does this look like?


2D Fourier transform

- 2D basis functions are oriented sinusoidal “gratings”:



Basis function examples

(u, v)




Real
component



Basis function examples

(u, v)




Real
component



Basis function examples

(u, v)




Real
component



Linear combination of basis functions

(u, v)



Real
component



2D Fourier transform

$$F(u, v) = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} f(x, y) e^{-i2\pi(ux+vy)} dx dy$$


- Output is 2D and complex-valued:

$$F(u, v) = \text{Re}(F(u, v)) + i \text{Im}(F(u, v))$$

- Magnitude spectrum: $|F(u, v)| = \sqrt{\text{Re}(F(u, v))^2 + \text{Im}(F(u, v))^2}$
- Phase angle spectrum: $\tan^{-1} \frac{\text{Im}(F(u, v))}{\text{Re}(F(u, v))}$
- Symmetry: the Fourier transform of a real-valued image has coefficients that come in pairs, with $F(u, v)$ being the *complex conjugate* of $F(-u, -v)$
 - This means that the magnitude spectrum is symmetric about the origin

2D discrete Fourier transform

$$F(u, v) = \sum_{n=0}^{N-1} \sum_{m=0}^{M-1} f(n, m) \exp \left(-i2\pi \left(\frac{un}{N} + \frac{vm}{M} \right) \right)$$




Source: [B. Freeman](#)

Real image examples

intensity image



log fft magnitude




Real image examples



intensity image



log fft magnitude



Which image goes with which spectrum?




Phase vs. magnitude

- Which has more information, the phase or the magnitude?
- Let's take the phase from one image and combine it with the magnitude from another image



Magnitude



Phase





Phase




Magnitude





Images with periodic patterns

- The magnitude image has peaks corresponding to the frequencies of repetition


Image



Application: Removing periodic patterns



Magnitude
image



Remove
peaks




Application: Removing periodic patterns


Lunar orbital image
(1966)



Magnitude
image



Remove
peaks





Join lines
removed



Source: [A. Zisserman](#)

Image transformations


- How does the FT change when the image is scaled?



Scaled by the
inverse factor!

Image transformations


- How does the FT change when the image is rotated?



Rotates the same way!

Image transformations


- How does the FT change when the image is rotated?



Caution: in real images this is not always the case because of edge artifacts (recall that DFT treats images as periodically tiled)

Image transformations


- How does the FT change when the image is translated?



Magnitude spectrum
doesn't change,
phase gets
modulated

Last time: Fourier analysis


- 1D Fourier transform
 - Definition and properties
 - Discrete Fourier transform
- 2D Fourier transform
 - Definition
 - Examples and properties



Review: 2D Fourier transform

- 2D basis functions are oriented sinusoidal “gratings”:

$$\psi_{u,v}(x, y) = e^{i2\pi(ux+vy)} = \cos 2\pi(ux + vy) + i \sin 2\pi(ux + vy)$$




- (u, v) is the direction *normal* to the grating
- The period is $1/\sqrt{u^2 + v^2}$

[Fourier image decomposition demo](#)

Aside: Discrete Cosine Transform

- Real-valued
- Fewer basis elements are needed to get a good approximation of the image
 - more useful for compression

DCT basis functions



https://en.wikipedia.org/wiki/Discrete_cosine_transform

Review: 2D Fourier transform

$$F(u, v) = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} f(x, y) e^{-i2\pi(ux+vy)} dx dy$$

- Output is 2D and complex-valued:

$$F(u, v) = \text{Re}(F(u, v)) + i \text{Im}(F(u, v))$$

- Magnitude spectrum: $|F(u, v)| = \sqrt{\text{Re}(F(u, v))^2 + \text{Im}(F(u, v))^2}$
- Phase angle spectrum: $\tan^{-1} \frac{\text{Im}(F(u, v))}{\text{Re}(F(u, v))}$

Review: 2D Fourier transform


$$F(u, v) = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} f(x, y) e^{-i2\pi(ux+vy)} dx dy$$

- Inverse transform:

$$f(x, y) = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} F(u, v) e^{i2\pi(ux+vy)} du dv$$

Review: 2D discrete Fourier transform




$$F(u, v) = \sum_{n=0}^{N-1} \sum_{m=0}^{M-1} f(n, m) \exp \left(-i2\pi \left(\frac{un}{N} + \frac{vm}{M} \right) \right)$$



Source: [B. Freeman](#)

Review: Which image goes with which spectrum?

Swapped magnitude and phase



[Source](#)

Outline

- 1D Fourier transform
 - Definition and properties
 - Discrete Fourier transform
- 2D Fourier transform
 - Definition
 - Examples and properties
- Convolution theorem
- Understanding the behavior of filters
- Understanding the sampling theorem
- More applications of Fourier analysis

Convolution theorem


- **Convolution** in the spatial domain translates to **multiplication** in the frequency domain (and vice versa)
- The Fourier transform of the convolution of two functions is the product of their Fourier transforms:

$$\mathcal{F}\{f * g\} = \mathcal{F}\{f\} \mathcal{F}\{g\}$$

- The inverse Fourier transform of the product of two Fourier transforms is the convolution of the two inverse Fourier transforms:

$$\mathcal{F}^{-1}\{FG\} = \mathcal{F}^{-1}\{F\} * \mathcal{F}^{-1}\{G\}$$

2D convolution theorem example



Convolution theorem

- Suppose f and g both consist of N pixels
- What is the complexity of computing $f * g$ in the spatial domain?
 - $O(N^2)$
- And what is the complexity of computing $\mathcal{F}^{-1}\{\mathcal{F}\{f\}\mathcal{F}\{g\}\}$?
 - $O(N \log N)$ using FFT
- Thus, convolution of an image with a large filter can be more efficiently done in the frequency domain


Outline

- 1D Fourier transform
 - Definition and properties
 - Discrete Fourier transform
- 2D Fourier transform
 - Definition
 - Examples and properties
- Convolution theorem
- Understanding the behavior of filters

Understanding the behavior of filters

- Why does filtering with a Gaussian give a nice smooth image, but filtering with a box filter gives artifacts?



Gaussian



Box filter



Recall: Fourier transform pairs




Filtering with a Gaussian




intensity image



filter: gaussian



filtered image




Filtering with a box filter




intensity image



filter: box




filtered image



Low-pass filtering

Image



[Source \(demo\)](#)



Low-pass filtering

- What can we say about the low-pass filtering result?

Image




Low-pass filtered



[Source \(demo\)](#)

Low-pass filtering

- What can we say about the low-pass filtering result?
- It has ringing artifacts – why?
 - “Hard” low-pass filtering is equivalent to convolving with a sinc function in the spatial domain






- This is why Gaussian filtering is preferred

Image



Low-pass filtered





[Source \(demo\)](#)

High-pass filtering

Image



Low-pass filtered



[Demo](#)

Source: [A. Zisserman](#)

High-pass filtering





Image



Low-pass filtered




High-pass filtered



[Demo](#)

Source: [A. Zisserman](#)


Hybrid images revisited



A. Oliva et al. [Hybrid Images](#). SIGGRAPH 2006

Hybrid images in the frequency domain

Low-pass filter




High-pass filter

A. Oliva et al. [Hybrid Images](#). SIGGRAPH 2006

Hybrid images in the frequency domain




Having a gap in the frequency content is important for the illusion to be effective




A. Oliva et al. [Hybrid Images](#). SIGGRAPH 2006

Psychophysical explanation




[Source](#)

Psychophysical explanation




[Source](#)

Psychophysical explanation



Psychophysical explanation



Hybrid painting

Salvador Dali


*"Gala Contemplating the Mediterranean Sea,
which at 30 meters becomes the portrait
of Abraham Lincoln", 1976*



Hybrid painting

Salvador Dali


*"Gala Contemplating the Mediterranean Sea,
which at 30 meters becomes the portrait
of Abraham Lincoln", 1976*



Hybrid painting

Salvador Dali

*"Gala Contemplating the Mediterranean Sea,
which at 30 meters becomes the portrait
of Abraham Lincoln", 1976*



Outline

- 1D Fourier transform
 - Definition and properties
 - Discrete Fourier transform
- 2D Fourier transform
 - Definition
 - Examples and properties
- Convolution theorem
- Understanding the behavior of filters
- Understanding the sampling theorem

Recall: Sampling and aliasing


256x256

128x128

64x64


32x32

16x16



Recall: Nyquist-Shannon sampling theorem

- When sampling a signal at discrete intervals, the sampling frequency must be at least *twice* the maximum frequency of the input signal to allow us to reconstruct the original perfectly from the sampled version




https://en.wikipedia.org/wiki/Nyquist-Shannon_sampling_theorem

Understanding the sampling theorem


- Suppose we have a continuous function $f(t)$ and we want to sample it at discrete intervals with a spacing of T
- This can be accomplished by multiplying it by the *comb function* or *impulse train*:

$$\text{comb}(t) = \sum_{n=-\infty}^{\infty} \delta(t - nT)$$



Understanding the sampling theorem

- Let's (formally) take the Fourier transform:




$$f_s(t) = f(t) \odot \text{comb}(t; T)$$

*Officially, the FT of the comb function doesn't exist since it's periodic, and since δ is a weird function

Understanding the sampling theorem


- Let's (formally) take the Fourier transform:



*Officially, the FT of the comb function doesn't exist since it's periodic, and since δ is a weird function

Understanding the sampling theorem


- Let's (formally) take the Fourier transform:



$$f_s(t) = f(t) \odot \text{comb}(t; T)$$

$$F_s(u) = F(u) * \frac{1}{T} \text{comb}\left(t; \frac{1}{T}\right)$$



*Officially, the FT of the comb function doesn't exist since it's periodic, and since δ is a weird function



Replicated copies of $F(u)$!


Understanding the sampling theorem

- How do we reconstruct $f(t)$?
- Let's apply a box filter in the frequency domain
 - Equivalent to convolving with a sinc function in the original domain
- When will this work perfectly?
 - When the sampling frequency $1/T$ exceeds twice the greatest frequency contained in $F(u)$!




Understanding the sampling theorem

- If the sampling frequency is too small, frequencies above the Nyquist limit are “folded back” onto smaller frequencies, resulting in aliasing




Sampling theorem in 2D

- If the Fourier transform of a continuous function $f(x, y)$ is zero for all frequencies beyond u_b and v_b , i.e., if the Fourier transform is *band-limited*, then $f(x, y)$ can be completely reconstructed from its samples as long as the sampling distances w and h along the x and y directions are such that $w \leq \frac{1}{2u_b}$ and $h \leq \frac{1}{2v_b}$




Aside: Analyzing interpolation methods

- Perfect reconstruction of the subsampled signal requires convolution with a sinc filter in the spatial domain, which is bad because sinc has infinite support
- Instead, faster and cheaper reconstruction (interpolation) methods are typically used




Aside: Analyzing different interpolation methods


- Linear reconstruction can be done by convolving the sampled signal with a *triangle filter*:



- However, the Fourier transform of the triangle filter is the sinc^2 function, so multiplying the signal's spectrum by it introduces high-frequency artifacts



Bilinear interpolation closeup




[Image source](#)

Outline

- 1D Fourier transform
 - Definition and properties
 - Discrete Fourier transform
- 2D Fourier transform
 - Definition
 - Examples and properties
- Convolution theorem
- Understanding the behavior of filters
- Understanding the sampling theorem
- More applications of Fourier analysis

Accelerating convolutional networks



M. Mathieu, M. Henaff, and Y. LeCun. [Fast Training of Convolutional Networks through FFTs](#). ICLR 2014

Understanding neural network shortcuts




Figure 1. More empirical studies found that humans rely on more phase components to achieve robust recognition. However, CNN without effective training restrictions tends to converge at the local optimum related to the amplitude spectrum of the image, leading to generalization behaviors counter-intuitive to humans.

G. Chen et al. [Amplitude-Phase Recombination: Rethinking Robustness of Convolutional Neural Networks in Frequency Domain](#). ICCV 2021




Figure 1: Images of 'container ship' and 'siamese cat' and their DFM-filtered versions with only top-5% dominant frequencies (the white dots in the central figures) retained can both be recognized correctly by NNs.

S. Wang et al. [What do neural networks learn in image classification? A frequency shortcut perspective](#). ICCV 2023

Detecting fake images




Figure 7: **Frequency analysis on each dataset.** We show the average spectra of each high-pass filtered image, for both the real and fake images, similar to Zhang *et al.* [50]. We observe periodic patterns (dots or lines) in most of the synthetic images, while BigGAN and ProGAN contains relatively few such artifacts.

S.-Y. Wang et al. [CNN-generated images are surprisingly easy to spot... for now.](#) CVPR 2020