

6.3000: Signal Processing

2D Fourier Transforms 2

- Structure of 2D Transforms
- Directionality and Rotation
- Magnitudes of Fourier Transforms
- Phases of Fourier Transforms
- 2D Convolution



Last Time: Introduction to 2D Signal Processing

2D signal processing builds on simple extensions of 1D.

Domain: time (t) $\rightarrow$ space (x, y)

Basis Functions: $e^{j\omega t}$ $\rightarrow$ $e^{j(\omega_x x + \omega_y y)}$

Transform: $\int$ $\rightarrow$ $\int \int$



Many similarities between 1D and 2D.

Some new issues/considerations.

2D Discrete Fourier Transform

Analysis and synthesis equations are similar.

One dimensional DFT:

$$F[k] = \frac{1}{N} \sum_{n=0}^{N-1} f[n] e^{-j\frac{2\pi k}{N}n}$$

$$f[n] = \sum_{k=0}^{N-1} F[k] e^{j\frac{2\pi k}{N}n}$$

Two dimensional DFT:

$$F[k_x, k_y] = \frac{1}{N_x N_y} \sum_{n_x=0}^{N_x-1} \sum_{n_y=0}^{N_y-1} f[n_x, n_y] e^{-j\left(\frac{2\pi k_x}{N_x}n_x + \frac{2\pi k_y}{N_y}n_y\right)}$$

$$f[n_x, n_y] = \sum_{k_x=0}^{N_x-1} \sum_{k_y=0}^{N_y-1} F[k_x, k_y] e^{j\left(\frac{2\pi k_x}{N_x}n_x + \frac{2\pi k_y}{N_y}n_y\right)}$$

But there are some surprising relations.

2D Discrete Fourier Transform

We can break a 2D DFT into a sequence of 1D DFTs.

$$\begin{aligned} F[k_x, k_y] &= \frac{1}{N_x N_y} \sum_{n_x=0}^{N_x-1} \sum_{n_y=0}^{N_y-1} f[n_x, n_y] e^{-j\left(\frac{2\pi k_x}{N_x} n_x + \frac{2\pi k_y}{N_y} n_y\right)} \\ &= \underbrace{\frac{1}{N_y} \sum_{n_y=0}^{N_y-1} \left(\underbrace{\frac{1}{N_x} \sum_{n_x=0}^{N_x-1} f[n_x, n_y] e^{-j\frac{2\pi k_x}{N_x} n_x}}_{\text{first take DFTs of rows}} \right) e^{-j\frac{2\pi k_y}{N_y} n_y}}_{\text{then take DFTs of resulting columns}} \end{aligned}$$

Start with a 2D function of space $f[n_x, n_y]$.



- Replace each row by the DFT of that row.
- Replace each column by the DFT of that column.

The result is $F[k_x, k_y]$, the 2D DFT of $f[n_x, n_y]$.

→ new directionality relations

Ocean

We see these same trends in real images.





horizontal features in the ocean view $\rightarrow$ strong vertical line in the DFT

significant vertical spectral content; little horizontal spectral content

Trees

We see these same trends in real images.



vertical features in forest image $\rightarrow$ strong horizontal content in DFT
significant horizontal spectral content; little vertical spectral content


Directionality and Rotation

2D signals have **directionality** (vertical versus horizontal structure).




Rotation is a property of 2D without counterpart in 1D.

Rotation of Images

Example: Find the DFT of a triplet.




DFT
 $\Rightarrow$






Rotation of Images

Example: Find the DFT of a triplet.



$\xRightarrow{\text{DFT}}$



Rotation of Images

Rotating an image rotates its Fourier transform by the same angle.

Start with the definition of the 2D CTFT.

$$F(\omega_x, \omega_y) = \int \int f(x, y) e^{-j(\omega_x x + \omega_y y)} dx dy$$

Express points (x, y) in space as (r, θ) and points (ω_x, ω_y) in the frequency plane as (ω, ϕ) .

$$\begin{aligned}\omega_x x + \omega_y y &= \underbrace{\omega \cos \phi}_{\omega_x} \underbrace{r \cos \theta}_x + \underbrace{\omega \sin \phi}_{\omega_y} \underbrace{r \sin \theta}_y \\ &= \omega r (\cos \phi \cos \theta + \sin \phi \sin \theta) = \omega r \cos(\phi - \theta)\end{aligned}$$

Then

$$F_p(\omega, \phi) = \int \int f_p(r, \theta) e^{-j\omega r \cos(\phi - \theta)} r dr d\theta$$

where $f_p(r, \theta)$ and $F_p(\omega, \phi)$ are polar equivalents of $f(x, y)$ and $F(\omega_x, \omega_y)$.

Rotation of Images

Rotating an image rotates its Fourier transform by the same angle.

$$F_p(\omega, \phi) = \int \int f_p(r, \theta) e^{-j\omega r \cos(\phi - \theta)} r dr d\theta$$

If

$$f_1(r, \theta) \xrightarrow{\text{CTFT}} F_1(\omega, \phi)$$

and

$$f_2(r, \theta) = f_1(r, \theta - \psi)$$



Then

$$\begin{aligned} F_2(\omega, \phi) &= \int \int f_2(r, \theta) e^{-j\omega r \cos(\phi - \theta)} r dr d\theta \\ &= \int \int f_1(r, \theta - \psi) e^{-j\omega r \cos(\phi - \theta)} r dr d\theta \\ &= \int \int f_1(r, \lambda) e^{-j\omega r \cos(\phi - (\lambda + \psi))} r dr d\lambda \\ &= F_1(\omega, \phi - \psi) \end{aligned}$$

Rotating a picture by ψ rotates its Fourier transform by ψ .

Moon


What are the dominant features of the DFT magnitude of the moon?



Large distribution of frequencies. Concentration along $r = c$ axis due to illumination from upper left.

Coordinate Transformations (e.g., in numpy)

If $x \rightarrow r$ and $y \rightarrow c$ then $k_x \rightarrow k_r$ and $k_y \rightarrow k_c$.




Using r, c coordinates (in `numpy`) is equivalent to using x, y coordinates.

Magnitude and Phase


So far, we have only considered magnitude. Does phase matter?



Magnitude




Phase




There is clearly structure in the magnitude; phase looks random.

Magnitude and Phase

Zeroing out the phase has an enormous impact on the image.



Magnitude




Uniform Phase



Phase is clearly important.

Magnitude and Phase


Flattening the magnitude has a big effect.



Uniform Mag




Phase




But the image is still recognizable!

Magnitude and Phase


Substituting the magnitude from a different image has a big effect.



Different Mag




Phase



But the boat is recognizable. What magnitude was used?

Discrete Fourier Series of Sounds

We previously looked at Fourier representations for sounds.
Phase played a (relatively) minor role in auditory perception.
These signals have the same magnitudes but different phases.



But they all sound very similar to each other.

Visual Perception of Phase


Why are images so sensitive to phase?

Convolution and Filtering

As with 1D, one of the most important applications of the 2D DFT is in computing the responses of signal processing systems.

Convolution and Filtering (1D)

If a system is **linear and time invariant**, then its response to any input $x[n]$ is $(x * h)[n]$ where $h[n]$ is the unit-sample response of the system.



Convolution in time is equivalent to multiplication in frequency $\rightarrow$ **filtering**.


$$\left. \begin{array}{l} x[n] \xrightarrow{\text{DTFT}} X(\Omega) \\ h[n] \xrightarrow{\text{DTFT}} H(\Omega) \end{array} \right\} (h * x)[n] \xrightarrow{\text{DTFT}} H(\Omega)X(\Omega)$$

Using the DFT **speeds** computation, but makes convolution **“circular.”**


$$\left. \begin{array}{l} x[n] \xrightarrow{\text{DFT}} X[k] \\ h[n] \xrightarrow{\text{DFT}} H[k] \end{array} \right\} (h \circledast x)[n] \xrightarrow{\text{DFT}} H[k]X[k]$$

Comparing Regular and Circular Convolution (1D)

Convolve $h[n]$ with $x[n]$ given below.




$$(h * x)[n] = \sum_{m=-\infty}^{\infty} h[m]x[n-m] = x[n] + x[n-1] + x[n-2]$$




Comparing Regular and Circular Convolution (1D)

Find the circular convolution of $h[n]$ with $x[n]$ using DFT (length $N = 5$).






$$(h \circledast x)[n] = (h * x_p)[n] \quad \text{where} \quad x_p[n] = \sum_{m=-\infty}^{\infty} x[n + mN]$$



Circular convolution of $h[n]$ with $x[n]$ is equivalent to conventional convolution of $h[n]$ with a periodically extended version of $x[n]$.

Comparing Regular and Circular Convolution (1D)

Circular convolution of $h[n]$ with $x[n]$ is equivalent to conventional convolution of $h[n]$ with a periodically extended version of $x[n]$.




Comparing Regular and Circular Convolution (1D)

Circular convolution is also equivalent to an aliased version of $(h * x)[n]$.

$$(h \circledast x)[n] = \sum_{m=-\infty}^{\infty} (h * x)[n+mN]$$

where N is the length of the DFT analysis window (here $N = 5$).



Comparing Regular and Circular Convolution (1D)

Two ways to relate circular convolution to conventional convolution:

- convolution of $h[n]$ with a periodically extended version of $x[n]$

$$(h \circledast x)[n] = (h * x_p)[n] \quad \text{where} \quad x_p[n] = \sum_{m=-\infty}^{\infty} x[n+mN]$$

- aliased version of $(h * x)[n]$

$$(h \circledast x)[n] = \sum_{m=-\infty}^{\infty} (h * x)[n+mN]$$

Why should this be true?

Comparing Regular and Circular Convolution (1D)

Two ways to relate circular convolution to conventional convolution:

- convolution of $h[n]$ with a periodically extended version of $x[n]$

$$(h \circledast x)[n] = (h * x_p)[n] \quad \text{where} \quad x_p[n] = \sum_{m=-\infty}^{\infty} x[n+mN]$$

- aliased version of $(h * x)[n]$

$$(h \circledast x)[n] = \sum_{m=-\infty}^{\infty} (h * x)[n+mN]$$


Aliasing and periodic extension are equivalent operations!

Aliasing and Periodic Extension

If a system is LTI, then periodic extension of its input is equivalent to aliasing its output.

Let $\mathcal{G}$ represent an LTI system with unit-sample response

$$g[n] = \sum_{m=-\infty}^{\infty} \delta[n+mN]$$



2D Convolution


2D convolution is similar, but both input and unit-sample response are 2D.

If a system is linear and shift-invariant, its response to input $f[r, c]$ is a superposition of shifted and scaled versions of unit-sample response $h[r, c]$.

$$\begin{aligned}(h * f)[r, c] &= \sum_{m=-\infty}^{\infty} \sum_{n=-\infty}^{\infty} h[m, n] f[m - r, n - c] \\ &= \sum_{m=-\infty}^{\infty} \sum_{n=-\infty}^{\infty} f[m, n] h[m - r, n - c]\end{aligned}$$


2D Convolution

Graphical representation of 2D convolution.




Check Yourself

Let white represent 1 and black represent 0.




Which of the following represents the convolution $f_a * f_b$?




where black and white represents the smallest and largest pixel value.

Comparison of Conventional and Circular 2D Convolution

Conventional convolution: convolve in space or implement with DTFT




Circular convolution: implement with DFT



Aliasing view: $\circledast$ wraps $*$ result vertically, horizontally, and diagonally.



Comparison of Conventional and Circular 2D Convolution

Periodic extension of input view.



Comparison of Conventional and Circular 2D Convolution


Periodic extension result.



Comparison of Conventional and Circular 2D Convolution

The output of conventional convolution can be bigger than the input, while that of circular convolution aliases to the same size as the input.

conventional



circular



Summary

Rotating an image rotates its transform.

Many features of an image (such as the orientations of structures) are apparent in the **magnitude** of the Fourier transform, but the **phase** of the Fourier transform is crucial to representing sharp edges.

Convolution in 2D requires flip-and-shift in both x and y directions. Circular convolution in 2D wraps the result of conventional convolution so as to match the size of the result with that of the 2D DFT.

Question of the Day

Two 3×3 images are shown below with white representing 1 and black representing 0.



Determine the 2D signal that results from circular convolution of these images, where the circular convolution is implemented with $N_x=N_y=3$.