

Discrete Fourier Transform

DD2423 Image Analysis and Computer Vision

Mårten Björkman

Computational Vision and Active Perception
School of Computer Science and Communication

November 13, 2013

$$\mathcal{F}(f(x)) = \int_{x \in \mathbb{R}^n} f(x) e^{-i\omega^T x} dx = \hat{f}(\omega)$$

$$\mathcal{F}^{-1}(\hat{f}(\omega)) = \frac{1}{(2\pi)^n} \int_{\omega \in \mathbb{R}^n} \hat{f}(\omega) e^{i\omega^T x} d\omega$$

Terminology:

Frequency spectrum: $\hat{f}(\omega) = |\hat{f}(\omega)| e^{i\phi(\omega)}$

Fourier spectrum: $|\hat{f}(\omega)|$

Power spectrum: $|\hat{f}(\omega)|^2$

Why Fourier transforms?

- Why are we interested in a decomposition of an image into complex exponential functions?
- Sinusoids and cosinusoids are eigenfunctions of convolutions!

$$e^{j\omega t} \xRightarrow{\mathcal{L}} A(\omega)e^{j\omega t}$$

- Note: $A(\omega)$ is complex (change in magnitude and phase).
- Thus we can understand what the filter does to the different frequencies of the image.

Theorem 2

- Convolution in the spatial domain is same as multiplication in the Fourier (frequency) domain

$$\mathcal{F}(h * f) = \mathcal{F}(h)\mathcal{F}(f)$$

$$f \rightarrow \boxed{*h} \rightarrow g = h * f \qquad \hat{f} \rightarrow \boxed{\hat{h}} \rightarrow \hat{g} = \hat{h} \hat{f}$$

Usage:

- For analysis and understanding of convolution operators.
- Some filters may be easily represented in the Fourier domain.
- Implementation: when size of the filter is too large it is more effective to use multiplication in the Fourier domain.

$$\mathcal{F}(h \star f) = \mathcal{F}(h) \mathcal{F}(f)$$

Proof:

$$\begin{aligned}\mathcal{F}(h \star f)(\omega) &= \int_{x \in \mathbb{R}^n} \left(\int_{\eta \in \mathbb{R}^n} h(x - \eta) f(\eta) d\eta \right) e^{-i\omega^T x} dx && \{\text{rewrite}\} \\&= \int_{\eta \in \mathbb{R}^n} \left(\int_{x \in \mathbb{R}^n} h(x - \eta) e^{-i\omega^T (x - \eta)} dx \right) f(\eta) e^{-i\omega^T \eta} d\eta && \{\text{with } (x - \eta) = \zeta\} \\&= \int_{\eta \in \mathbb{R}^n} \left(\int_{\zeta \in \mathbb{R}^n} h(\zeta) e^{-i\omega^T \zeta} d\zeta \right) f(\eta) e^{-i\omega^T \eta} d\eta && \{\text{separate}\} \\&= \left(\int_{\zeta \in \mathbb{R}^n} h(\zeta) e^{-i\omega^T \zeta} d\zeta \right) \left(\int_{\eta \in \mathbb{R}^n} f(\eta) e^{-i\omega^T \eta} d\eta \right) = \mathcal{F}(h)(\omega) \mathcal{F}(f)(\omega)\end{aligned}$$

- Given

$$h(x, y) = h_1(x)h_2(y)$$

$$\begin{aligned}(h * f)(x, y) &= \int_{\eta \in \mathbb{R}^n} \int_{\zeta \in \mathbb{R}^n} h(\eta, \zeta) f(x - \eta, y - \zeta) d\eta d\zeta \\ &= \int_{\eta \in \mathbb{R}^n} h_1(\eta) \underbrace{\left(\int_{\zeta \in \mathbb{R}^n} h_2(\zeta) f(x - \eta, y - \zeta) d\zeta \right)}_{***} d\eta\end{aligned}$$

*** convolution of a column (fixed value of x) in y -direction

- If convolution mask $h(x, y)$ can be separated as above \Rightarrow 2D convolution can be performed as a series of 1D convolutions.
- Discrete case: If the mask is m^2 in size $\Rightarrow 2m$ operations / pixel instead of m^2 operations per pixel.

$$\begin{aligned}\hat{f} &= \int_{\omega_1 \in \mathbb{R}^n} \int_{\omega_2 \in \mathbb{R}^n} f(x, y) e^{-i(\omega_1 x_1 + \omega_2 x_2)} dx_1 dx_2 \\ &= \int_{\omega_1 \in \mathbb{R}^n} e^{-i\omega_1 x_1} \left(\int_{\omega_2 \in \mathbb{R}^n} f(x, y) e^{-i\omega_2 x_2} dx_2 \right) dx_1\end{aligned}$$

- A Fourier transform in 2D can always be performed as a series of two 1D Fourier transforms.

Filtering techniques typically modify frequency characteristics:

- Enhance edges (increase high frequency)
- Remove noise (decrease high frequency)
- Smooth (decrease high frequency, increase low frequency)

Observation: Convolution is a spatial operation on an image.

Images can be converted to their frequency component prior to filtering to facilitate direct manipulation of image frequency characteristics.

Spatial versus frequency domain

The Fourier Transform converts spatial image data into a frequency representation. Both representations contain equivalent information.

Spatial Domain

- + Intuitive Representation
- Designing filters can be hard
- Filtering with large kernels may result in long processing times.
- + Kernels applied directly to spatial data.

Frequency Domain

- Non-intuitive representation
- + Designing filters often easier.
- + Filtering with large kernels can be performed very quickly
- Image and Kernel must first be converted to frequency domain, modified, then reconverted.

$$\hat{f}(u, v) = \frac{1}{\sqrt{MN}} \sum_{m=0}^{M-1} \sum_{n=0}^{N-1} f(m, n) e^{-2\pi i (\frac{mu}{M} + \frac{nv}{N})} \quad (1)$$

$$f(m, n) = \frac{1}{\sqrt{MN}} \sum_{u=0}^{M-1} \sum_{v=0}^{N-1} \hat{f}(u, v) e^{+2\pi i (\frac{mu}{M} + \frac{nv}{N})} \quad (2)$$

Terminology:

- Fourier spectrum: $|F(u, v)| = \sqrt{Re^2(u, v) + Im^2(u, v)}$
- Phase angle: $\phi(u, v) = \tan^{-1} \frac{Im(u, v)}{Re(u, v)}$
- Power spectrum: $P(u, v) = |F(u, v)|^2 = Re^2(u, v) + Im^2(u, v)$
- The magnitude is simply the peak value, and the phase determines where the origin is, or where the sinusoid starts.

- Continuous

$$\hat{f}(\omega) = \int_{x \in \mathbb{R}^n} f(x) e^{-i\omega^T x} dx$$

- Discrete

$$\hat{f}(u) = \frac{1}{\sqrt{M}^n} \sum_{x \in I^n} f(x) e^{-\frac{2\pi i u^T x}{M}}$$


- Frequency variables are related (in 1D) by

$$\omega = \frac{2\pi u}{M}$$


- Note: u assumes values $0 \dots M-1 \Rightarrow \omega \in [0, 2\pi)$.
- By periodic extension, we can map this integral to $[-\pi, \pi)$.

More insight (2 pixel wide stripes)


- The maximum frequency which can be represented in the spatial domain are one pixel wide stripes (period=2): $\omega_{max} = 2\pi\frac{1}{2} = \pi$
- So, 2 pixel wide stripes (period=4) give $\omega = 2\pi\frac{1}{4} = \frac{1}{2}\omega_{max}$
- Plotting magnitude of the Fourier transform:
 - Two points are halfway between the center and the edge of the image, i.e. the represented frequency is half of the maximum.
 - One point in the middle shows the DC-value (image mean).



- If you take the logarithm of the Fourier transform, you see many minor frequencies. The reason is that since an image can only be represented by square pixels, the diagonals cannot be represented without discretization noise.



Example images and Fourier transforms



Given a simple 4-pixel “image” $I(c) = [3, 2, 2, 1]$, what is its Fourier Transform $F(v)$?

Given a simple 4-pixel “image” $I(c) = [3, 2, 2, 1]$, what is its Fourier Transform $F(v)$?

$$F(v) = \frac{1}{N} \sum_{c=0}^{N-1} I(c) e^{-j2\pi \frac{vc}{N}}$$

$$F(0) = \frac{1}{4} \sum_{c=0}^3 I(c) e^{-j2\pi \frac{0c}{4}} = \frac{1}{4} \sum_{c=0}^3 I(c) e^0 = \frac{1}{4} [I(0) + I(1) + I(2) + I(3)] = \frac{1}{4} [3 + 2 + 2 + 1] = 2$$

$$F(1) = \frac{1}{4} \sum_{c=0}^3 I(c) e^{-j2\pi \frac{1c}{4}} = \frac{1}{4} [3e^0 + 2e^{-j\pi/2} + 2e^{-j\pi} + 1e^{-j3\pi/2}] = \frac{1}{4} [3 + 2(-j) + 2(-1) + 1(j)] = \frac{1}{4} [1 - j]$$

$$F(2) = \frac{1}{4} \sum_{c=0}^3 I(c) e^{-j2\pi \frac{2c}{4}} = \frac{1}{4} [3e^0 + 2e^{-j\pi} + 2e^{-j2\pi} + 1e^{-j3\pi}] = \frac{1}{4} [3 + (-2) + 2 + (-1)] = \frac{1}{2}$$

$$F(3) = \frac{1}{4} \sum_{c=0}^3 I(c) e^{-j2\pi \frac{3c}{4}} = \frac{1}{4} [3e^0 + 2e^{-j3\pi/2} + 2e^{-j3\pi} + 1e^{-j9\pi/2}] = \frac{1}{4} [3 + 2j + 2(-1) + 1(-j)] = \frac{1}{4} [1 + j]$$

$$\hat{f}(u, v) = \frac{1}{\sqrt{MN}} \sum_{m=0}^{M-1} \sum_{n=0}^{N-1} f(m, n) e^{-2\pi i (\frac{mu}{M} + \frac{nv}{N})} \quad (3)$$

$$\hat{f}(u, v) = \frac{1}{\sqrt{M}} \sum_{m=0}^{M-1} \left(\frac{1}{\sqrt{N}} \sum_{n=0}^{N-1} f(m, n) e^{-2\pi i \frac{nv}{N}} \right) e^{-2\pi i \frac{mu}{M}}$$

$$f(m, n) = \frac{1}{\sqrt{M}} \sum_{m=0}^{M-1} \left(\frac{1}{\sqrt{N}} \sum_{n=0}^{N-1} \hat{f}(u, v) e^{2\pi i \frac{nv}{N}} \right) e^{2\pi i \frac{mu}{M}}$$

2D DFT can be implemented as a series of 1D DFTs along each column, followed by 1D DFTs along each row.

$$\mathcal{F} [a f_1(m, n) + b f_2(m, n)] = a \hat{f}_1(u, v) + b \hat{f}_2(u, v)$$
$$a f_1(m, n) + b f_2(m, n) = \mathcal{F}^{-1} [a \hat{f}_1(u, v) + b \hat{f}_2(u, v)]$$

You can add two functions (images) or rescale a function, either before or after computing the Fourier transform. It leads to the same result.

- If the original function is multiplied with an exponential like the one below and transformed, it will result in a shift of the origin of the frequency plane to point (u_0, v_0) .

$$\mathcal{F} \left[f(m, n) e^{2\pi i (\frac{mu_0}{M} + \frac{nv_0}{N})} \right] = \hat{f}(u - u_0, v - v_0)$$

- For $(u_0, v_0) = (M/2, N/2)$

$$e^{2\pi i (\frac{mu_0}{M} + \frac{nv_0}{N})} = e^{\pi i (m+n)} = (-1)^{m+n}$$


and

$$f(m, n)(-1)^{m+n} \iff \hat{f}(u - M/2, v - N/2)$$

- Conclusion: the origin of the Fourier transform can be moved to the center by multiplying the original function by $(-1)^{m+n}$.

Property III - Modulation/Frequency translation

From left: Original image, magnitude of the Fourier spectrum, original multiplied by $1 + 2\cos \omega y$ at a relative frequency of 16, magnitude of the Fourier spectrum.




$$\text{Note: } 1 + 2\cos \omega y = 1 + e^{i\omega y} + e^{-i\omega y}.$$

Property IV - Translation

- If the image is moved, the Fourier spectrum undergoes a phase shift, but magnitude of the spectrum remains the same.


$$\mathcal{F} [f(m - m_0, n - n_0)] = \hat{f}(u, v) e^{-2\pi i (\frac{m_0 u}{M} + \frac{n_0 v}{N})}$$

$$|\hat{f}(u, v) e^{-2\pi i (\frac{m_0 u}{M} + \frac{n_0 v}{N})}| = |\hat{f}(u, v)|$$



Property V - Rotation

Rotation of the original image rotates \hat{f} by the same angle.



Exercise: Introduce polar coordinates and perform direct substitution.

$$A = \begin{pmatrix} S_1 & & \\ & \ddots & \\ & & S_n \end{pmatrix} \text{ (diagonal)}$$

$$g(x) = f(S_1 x_1, \dots, S_n x_n)$$

$$\hat{g}(\omega) = \frac{1}{|S_1 \dots S_n|} \hat{f}\left(\frac{\omega_1}{S_1}, \dots, \frac{\omega_n}{S_n}\right)$$

Conclusion: compression in spatial domain is same as expansion in Fourier domain (and vice versa).


- The DFT and its inverse are periodic with period N , for an $N \times N$ image. This means:

$$\hat{f}(u, v) = \hat{f}(u + N, v) = \hat{f}(u, v + N) = \hat{f}(u + N, v + N)$$

- This property defines the implied symmetry in the Fourier spectrum as well as that \hat{f} repeats itself infinitely.
- However, only one period is enough to reconstruct the original function f .

Property VIII - Conjugate Symmetry

- The Fourier transform satisfies $\hat{f}(u, v) = \hat{f}^*(-u, -v)$ and $|\hat{f}(u, v)| = |\hat{f}(-u, -v)|$.
- With periodicity and above, we have that \hat{f} has period N and is (conjugate) symmetric around the origin.

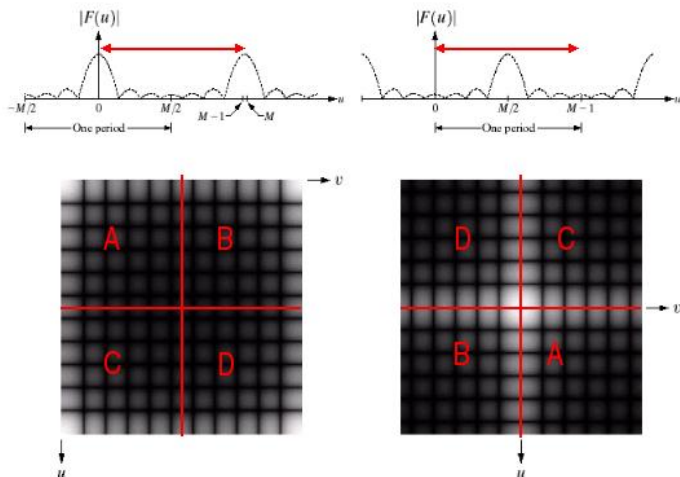


- Thus we don't need $2N^2$ (N^2 real and N^2 imaginary) values to represent a $N \times N$ image in Fourier domain, but N^2 if we exploit the symmetry.

a b
c d


FIGURE 4.34

(a) Fourier spectrum showing back-to-back half periods in the interval $[0, M - 1]$.
 (b) Shifted spectrum showing a full period in the same interval.
 (c) Fourier spectrum of an image, showing the same back-to-back properties as (a), but in two dimensions.
 (d) Centered Fourier spectrum.



Wavelength in two dimensions

- The wavelength of the sinusoid is: $\lambda = \frac{1}{\sqrt{u^2 + v^2}}$, where (u, v) are the frequencies along (r, c) and the periods are $1/u$ and $1/v$.



Multiplication in the spatial domain is same as convolution in the Fourier domain.

$$\mathcal{F}(hf) = \mathcal{F}(h) * \mathcal{F}(f)$$

Exercise: Prove it!

- A linear, shift invariant system (such as a filter) is completely specified by its response to an impulse, which is called the **impulse response**.
- The **transfer function** H is the Fourier transform of the impulse response.
- Using the convolution theorem to describe the effects of the system:

$$g(x, y) = h(x, y) * f(x, y)$$



$$G(u, v) = H(u, v) \cdot F(u, v)$$

- Convolution with an impulse function


$$h(x, y) * \delta(x, y) = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} h(x_0, y_0) \delta(x - x_0, y - y_0) dx_0 dy_0 = h(x, y)$$

results in a “copy” of $h(x, y)$ to the location of the impulse


Transfer function examples




$$e^{-t^2/2\sigma^2}$$




$$\sigma\sqrt{2\pi}e^{-\sigma^2\omega^2/2}$$




$$\cos \omega_0 t$$



$$\pi[\delta(\omega + \omega_0) - \delta(\omega - \omega_0)]$$




$$\sin \omega_0 t$$









$$j\pi[\delta(\omega + \omega_0) - \delta(\omega - \omega_0)]$$







$$\theta(\omega) = -\pi/2$$




Transfer function examples



Transfer function examples




- How many samples are required to describe the given signal without loss of information?
- What signal can be reconstructed given the current sampling rate?




- A signal is band limited if its highest frequency is bounded. This frequency is called the **bandwidth**.
- The sine/cosine component of the highest frequency determines the highest “frequency content” of the signal.
- If the signal is sampled at a rate equal or greater to than twice its highest frequency, the original signal can be completely recovered from its samples (Shannon).
- The minimum sampling rate for a band limited function is called **Nyquist rate**.

Reconstruction



- For reconstruction, we need to convolve with a sinc function.
 - It is the Fourier transform of the box function.
 - It has infinite support.
- May be approximated by a Gaussian, cubic or even “tent” function.



- If the signal is undersampled, aliasing occurs.
- Prevent aliasing by:
 - increasing sampling rate, or [infeasible]
 - decreasing highest frequency before sampling. [blurring]

Parseval's equality

The total energy contained in an image summed across all (x, y) is equal to the total energy of its Fourier Transform summed across all frequencies.

$$\int_{x \in \mathbb{R}^n} f(x)g(x)dx = \frac{1}{(2\pi)^n} \int_{\omega \in \mathbb{R}^n} \hat{f}(\omega)\hat{g}^*(\omega)d\omega$$

Proof:

$$\begin{aligned} \int_{\omega \in \mathbb{R}^n} \hat{f}(\omega)\hat{g}^*(\omega)d\omega &= \int_{\omega \in \mathbb{R}^n} \hat{f}(\omega) \left(\int_{x \in \mathbb{R}^n} g(x)e^{i\omega^T x} dx \right) d\omega \\ &= \int_{x \in \mathbb{R}^n} g(x) \underbrace{\left(\int_{\omega \in \mathbb{R}^n} \hat{f}(\omega)e^{i\omega^T x} d\omega \right)}_{(2\pi)^n f(x)} dx = (2\pi)^n \int_{x \in \mathbb{R}^n} f(x)g(x)dx \end{aligned}$$

If $f = g$ then

$$\int_{x \in \mathbb{R}^n} f(x)^2 dx = \frac{1}{(2\pi)^n} \int_{\omega \in \mathbb{R}^n} |\hat{f}|^2 d\omega.$$

Summary of good questions

- What is the Fourier transform of a convolution? Why is it important?
- What does separability of filters mean?
- How do you interpret a point in the Fourier domain in the spatial domain?
- How do you apply a discrete Fourier transform?
- What happens to the Fourier transform, if you translate an image?
- What happens to the Fourier transform, if you rotate an image?
- In what sense is the Fourier transform symmetric?
- What does the Sampling Theorem mean in practice?
- What can you do to get rid of aliasing in the sampling process?

- Gonzalez and Woods: Chapter 4
- Szeliski: Chapter 3.4
- Introduction to Lab 2