


The Frequency Domain



Randall Munroe, <http://xkcd.com/26/>

CSC320: Introduction to Visual Computing

Michael Guerzhoy


Many slides borrowed


from Steve Seitz, Alexei Efros, Derek Hoiem, Allan Jepson, David Fleet, John M. Brayer



Salvador Dalí


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





A set of basis vectors

Teases away fast vs. slow changes in the image.



This change of basis has a special name...

Jean Baptiste Joseph Fourier (1768-1830)

had crazy idea (1807)

Any univariate function can be rewritten as a weighted sum of sines and cosines of different frequencies.

...the manner in which the author arrives at these equations is not exempt of difficulties and...his analysis to integrate them still leaves something to be desired on the score of generality and even rigour.

Don't believe it?

- Neither did Lagrange, Laplace, Poisson and other big wigs
- Not translated into English until 1878!

But it's (mostly) true!

- called Fourier Series
- there are some subtle restrictions



A sum of sines

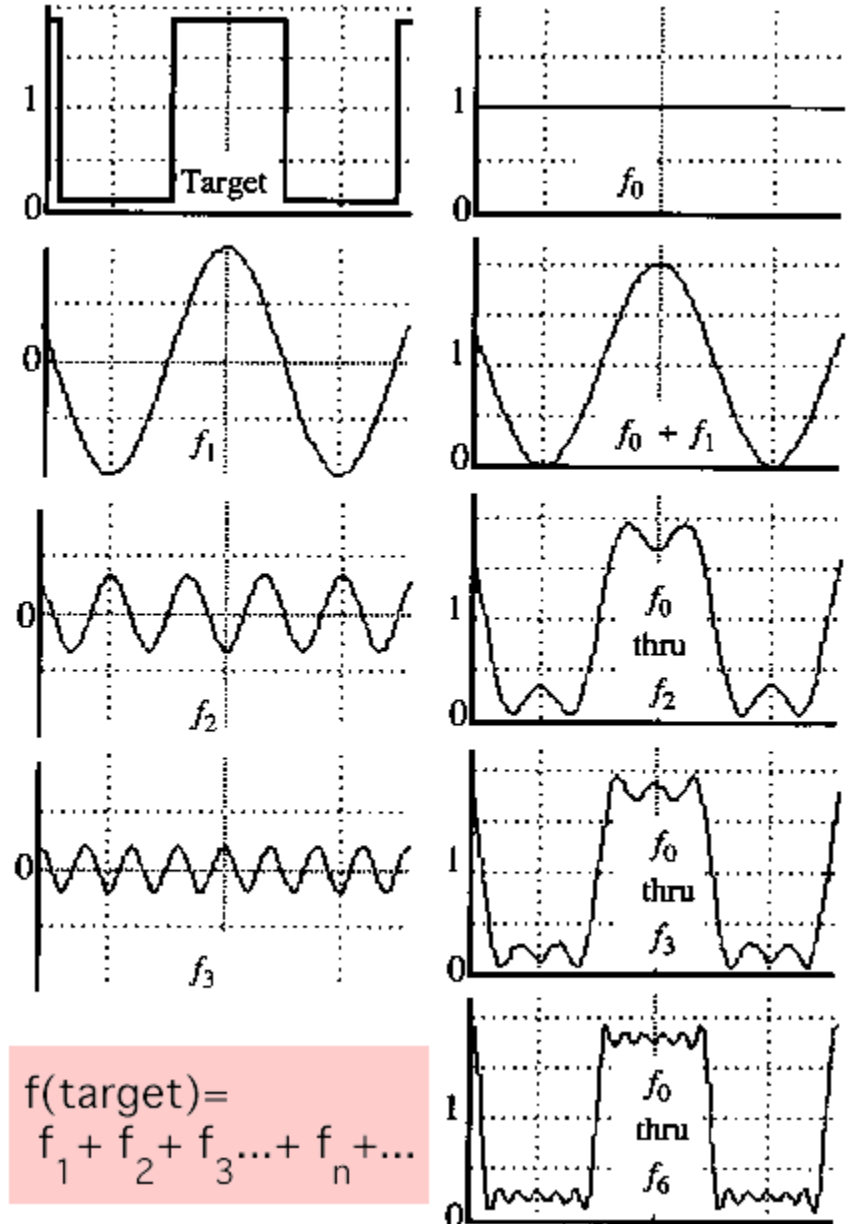
Our building block:

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

Add enough of them to get any signal $f(x)$ you want!

What does each control?

Which one encodes the coarse vs. fine structure of the signal?



$$f(\text{target}) = f_1 + f_2 + f_3 + \dots + f_n + \dots$$

Fourier Transform

We want to understand the frequency ω of our signal. So, let's reparametrize the signal by ω instead of x :



For every ω from 0 to ∞ , $F(\omega)$ holds the amplitude A and phase ϕ of the corresponding sine $A \sin(\omega x + \phi)$

- How can F hold both? Complex number trick!

$$F(\omega) = R(\omega) + iI(\omega)$$


$$A = \pm \sqrt{R(\omega)^2 + I(\omega)^2} \quad \phi = \tan^{-1} \frac{I(\omega)}{R(\omega)}$$

We can always go back:




Time and Frequency

example : $g(t) = \sin(2pf t) + (1/3)\sin(2p(3f) t)$





Time and Frequency

example : $g(t) = \sin(2pf t) + (1/3)\sin(2p(3f) t)$




Frequency Spectra

example : $g(t) = \sin(2pf t) + (1/3)\sin(2p(3f) t)$




Frequency Spectra


Usually, frequency is more interesting than the phase




Frequency Spectra




=





+




=




Frequency Spectra




Frequency Spectra




=




+




=




Frequency Spectra




=




+




=




Frequency Spectra




Frequency Spectra





$$= A \sum_{k=1}^{\infty} \frac{1}{k} \sin(2\pi kt)$$





Frequency Spectra




(a)



(b)




(c)




FT: Just a change of basis


$$M * f(x) = F(\omega)$$



*




||




IFT: Just a change of basis

$$M^{-1} * F(\omega) = f(x)$$



•
•
•



Finally: Scary Math

$$\text{Fourier Transform : } F(\omega) = \int_{-\infty}^{+\infty} f(x)e^{-i\omega x} dx$$

$$\text{Inverse Fourier Transform : } f(x) = \frac{1}{2\pi} \int_{-\infty}^{+\infty} F(\omega)e^{i\omega x} d\omega$$

Finally: Scary Math

$$\text{Fourier Transform : } F(\omega) = \int_{-\infty}^{+\infty} f(x)e^{-i\omega x} dx$$

$$\text{Inverse Fourier Transform : } f(x) = \frac{1}{2\pi} \int_{-\infty}^{+\infty} F(\omega)e^{i\omega x} d\omega$$


...not really scary: $e^{i\omega x} = \cos(\omega x) + i \sin(\omega x)$

is hiding our old friend: $A \sin(\omega x + \phi)$

$$\begin{array}{l} \text{phase can be encoded} \\ \text{by sin/cos pair} \end{array} \rightarrow \begin{array}{l} P \cos(x) + Q \sin(x) = A \sin(x + \phi) \\ A = \pm \sqrt{P^2 + Q^2} \quad \phi = \tan^{-1}\left(\frac{P}{Q}\right) \end{array}$$

So it's just our signal $f(x)$ times sine at frequency ω

Extension to 2D



2D Discrete Fourier Transform

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

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

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:


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

2D Discrete Fourier Transform (DFT)


- If $\hat{h}(k, l)$ is the 2D DFT of h , then
 - $|\hat{h}(k, l)|^2 = \text{Re}(\hat{h}(k, l))^2 + \text{Im}(\hat{h}(k, l))^2$ is the **power spectrum** of h

2D Fourier Basis Functions

Grating for $(k,l) = (1,-3)$




Real




Imag

Grating for $(k,l) = (7,1)$



Real


Note:
(0, 0) is in the top-left corner in the images here.
In other images in this lecture, the origin is on the bottom-left



Zero crossings



- Zero crossing of $\sin(\omega_k n + \omega_l m)$:
 - $\omega_k n + \omega_l m = q\pi$ for some q
 - Corresponds to an “edge” in the grating (switch from positive numbers to negatives, or vice versa)
- What is the orientation of the edge?
 - $\omega_k n + \omega_l m = q\pi$
 - $m = \frac{q\pi}{\omega_l} + (-\omega_k/\omega_l)n$
 - The slope $(-\omega_k/\omega_l)$ determines the direction
 - $(-\omega_k/\omega_l)$ positive \Rightarrow the edge goes top-left to bottom-right
 - $(-\omega_k/\omega_l)$ negative \Rightarrow the edge goes bottom-left to top-right
 - $(-\omega_k/\omega_l)$ zero \Rightarrow the edge is horizontal
 - $(-\omega_k/\omega_l)$ is infinity (i.e., $\omega_l = 0$) \Rightarrow edge is vertical

Fourier analysis in images




- The dot at the centre is the “zero frequency” term (the sum of the image): $\sum_{n,m} e^{-i(0 \times n + 0 \times m)} h(n, m)$
- Why two dots off side?

Why Two Dots?



- One for $(k, l) = (1, 1)$
- One for $(k, l) = (-1, -1)$
- Note: the gratings for those look the same for the real part
- Say F is a 2D Fourier transform of a *real* image f . Then:
 - $|F(u, v)| = |F(-u, -v)|$
 - (Since $\cos(x) = \cos(-x)$, $\sin(-x) = -\sin(x)$. Details: exercise)


Signals can be composed



<https://web.archive.org/web/20130513181427/http://sharp.bu.edu/~slehar/fourier/fourier.html>

More: <http://www.cs.unm.edu/~brayer/vision/fourier.html>

Man-made Scene




The FT of man-made scenes

- Edges have to be represented as sums of gratings, with the gratings in the same general direction as the edges
 - Partly* explains the high magnitude of the FT along the n and m axis in most outdoor images
- Other edges also need to be represented
 - In the Colosseum photo on the previous slide, we have two edge directions that are not vertical or horizontal
 - They correspond to the lines you see on the FT magnitude image (note: while in the photo $(0,0)$ is on the top right and x and y increase going to the bottom-right, in the FT magnitude image,


*Small print on the “+” pattern(aside, not covered in detail): Edge Effects

Two FTs of similar patterns



Aside, continued (not covered in detail)

- We treat the image as if it were tiling the entire plane, so if the left/right, top/bottom ends are different, we get “edge effects”



Can change spectrum, then reconstruct




Low Pass filtering

- Only keep the **low frequencies**, set the coefficients of the high frequencies to 0
- Similar to blurring with a Gaussian kernel



High Pass filtering

- Only keep the **high frequencies**, set the coefficients of the low frequencies to 0
- We'll talk later how to implement this using filters






Strong Vertical Frequency
(Sharp Horizontal Edge)

Diagonal Frequencies

Strong Horz. Frequency
(Sharp Vert. Edge)

Log Magnitude

Low Frequencies



The Convolution Theorem

The greatest thing since sliced (banana) bread!

- The Fourier transform of the convolution of two functions is the product of their Fourier transforms

$$F[g * h] = F[g]F[h]$$

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

$$F^{-1}[gh] = F^{-1}[g] * F^{-1}[h]$$

- **Convolution** in spatial domain is equivalent to **multiplication** in frequency domain!


2D convolution theorem example

$f(x,y)$




*

$h(x,y)$




⇓

$g(x,y)$




×

$|F(s_x, s_y)|$




⇓




$|G(s_x, s_y)|$

The Ringing Artifact, Again

Gaussian



Box filter




Gaussian Filtering

- The FT of a Gaussian(σ) is Gaussian($1/\sigma$)
 - You can prove it with integrals, or experimentally, or just trust me
- The Convolution Theorem gives us:


intensity image




filter: gaussian




filtered image




log ft magnitude of image



filter: gaussian




log ft magnitude of filtered image



Box Filtering isn't as nice







- The FT of a box filter is similar to what we saw for the step function: high frequencies remain



Fourier Transform Pairs




Spatial domain

Frequency domain





Low-pass, Band-pass, High-pass filters


low-pass:



Band-pass:



Edges in images



What does blurring take away?




original

What does blurring take away?



smoothed (5x5 Gaussian)

High-Pass filter




smoothed – original


Band-pass filtering

Gaussian Pyramid (low-pass images)




Da Vinci and Peripheral Vision





coarse components
(peripheral vision)



medium components
(near peripheral vision)



fine details
(central vision)


Leonardo playing with peripheral vision

Unsharp Masking



Unsharp Masking


- The blurred image has all the **low frequency** components intact
- If we subtract the blurred image from the original, we only keep the **high frequencies**



Unsharp masking

- If in the original, a pixel is brighter than its neighbours, the difference image will be *positive*
- If in the original, a pixel is darker than its neighbours, the difference image will be *negative*
- Otherwise, the difference image is around 0
- So adding the difference image had little effect away from the edges, and will tend to exaggerate the edges
- That is called “sharpening:” the edges will become less blurry


Lossy Image Compression (JPEG)



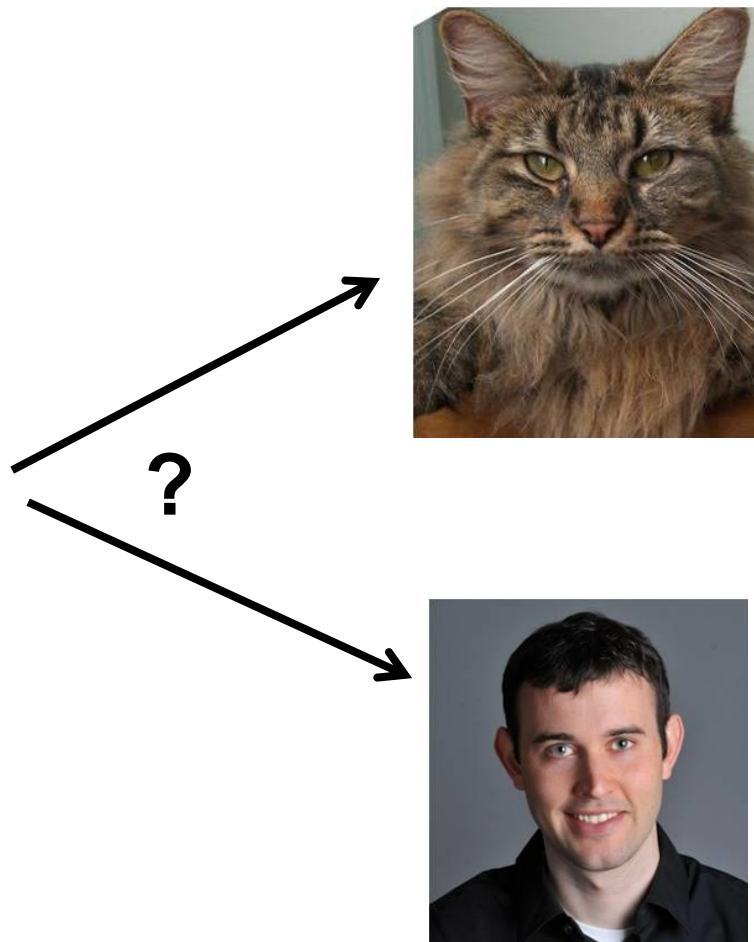
Block-based Discrete Cosine Transform (DCT) (similar to DFT)

Using DCT in JPEG

- The first coefficient $B(0,0)$ is the DC component, the average intensity
- The top-left coeffs represent low frequencies, the bottom right represent high frequencies
- Store as much of the top left as you can



Hybrid Images



http://cvcl.mit.edu/hybrid_gallery/gallery.html