

2020 COMPUTER VISION





Capture Frequency - Rolling `Shutter'





FOURIER SERIES & FOURIER TRANSFORMS

Thinking in frequency

I understand frequency as in waves...



...but how does this relate to the complex signals we see in natural images? ...to image frequency?

Fourier series

A bold idea (1807):

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

Our building block:

$$\sum_{n=1}^{\infty} (a_n \cos(nt) + b_n \sin(nt))$$

Add enough of them to get any signal *g(t)* you want!

Jean Baptiste Joseph Fourier (1768-1830)



$$t = [0,2], f = 1$$

$$g(t) = (1)\sin(2\pi f t) + (1/3)\sin(2\pi(3f) t)$$

$$g\left(\underbrace{\int_{t} \int_{0}^{t} \int_$$

Slides: Efros















Jean Baptiste Joseph Fourier (1768-1830)

A bold idea (1807):

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

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
 - Applies to periodic signals

...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.



Example: Music

We think of music in terms of pitches (frequencies) at different loudnesses (amplitudes).

Sine/cosine and circle



Square wave (approx.)



Mehmet E. Yavuz

Sawtooth wave (approx.)



Mehmet E. Yavuz

One series in each of x and y



<u>Generative Art, じゃがりきん, Video, 2018</u>



Wikipedia – Fourier transform



Wikipedia – Fourier transform

Amplitude-phase form

Add phase term to shift cos values into sin values



Morse [<u>Peppergrower;</u> Wikipedia]

Fourier Transform

- Stores the amplitude and phase at each frequency:
 - For mathematical convenience, this is often notated in terms of real and complex numbers
 - Related by Euler's formula



Fourier Transform

- Stores the amplitude and phase at each frequency:
 - For mathematical convenience, this is often notated in terms of real and complex numbers
 - Related by Euler's formula



Amplitude encodes how much signal there is at a particular frequency:

$$A = \pm \sqrt{\operatorname{Re}(\varphi)^2 + \operatorname{Im}(\varphi)^2}$$

Phase encodes spatial information (indirectly):

$$\phi = \tan^{-1} \frac{\operatorname{Im}(\varphi)}{\operatorname{Re}(\varphi)}$$

Amplitude-phase form

Add component of infinite frequency = mean of signal over period

= value around which signal fluctuates



Average of signal over period

Electronics signal processing calls this the 'DC offset'

How to read Fourier transform images

We display the space such that inf-frequency coefficient is in the center.

Fourier transform component image

$$\frac{a_0}{2} + \sum_{n=1}^{N} (a_n \sin(nx + \emptyset_n))$$



How to read Fourier transform images Image is rotationally symmetric about center because of negative frequencies

 $\frac{a_0}{2} + \sum_{n=1}^{N} (a_n \sin(nx + \emptyset_n))$

e.g., wheel rotating one way or the other

For real-valued signals, positive and negative frequencies are complex conjugates (see additional slides).



How to read Fourier transform images Image is as large as maximum frequency by resolution of input under Nyquist frequency

 $\frac{a_0}{2} + \sum_{n=1}^{N} (a_n \sin(nx + \phi_n))$

Nyquist frequency is half the sampling rate of the signal. Sampling rate is size of image (X and Y), so Fourier transform images are $(\pm X/2, \pm Y/2)$.



A few questions

How is the Fourier decomposition computed?

Intuitively, by correlating the signal with a set of waves of increasing frequency!

2D Discrete Fourier Transform $F[u,v] = \frac{1}{MN} \sum_{m=0}^{M-1} \sum_{n=0}^{N-1} I[m,n] \cdot e^{-i2p \left(\frac{um}{M} + \frac{vn}{N}\right)}$



Source: Seul et al, Practical Algorithms for Image Analysis, 2000, p. 249, 262.

2D FFT can be computed as two discrete Fourier transforms in 1 dimension

Earl F. Glynn

2D Discrete Fourier Transform

Edge represents highest frequency, smallest resolvable length (2 pixels)



Earl F. Glynn

Fourier analysis in images

Spatial domain images



Fourier decomposition amplitude images

http://sharp.bu.edu/~slehar/fourier/fourier.html#filtering More: http://www.cs.unm.edu/~brayer/vision/fourier.html

Signals can be composed

Spatial domain images



Fourier decomposition amplitude images

http://sharp.bu.edu/~slehar/fourier/fourier.html#filtering More: http://www.cs.unm.edu/~brayer/vision/fourier.html

Natural image

Natural image



Fourier decomposition Amplitude image



What does it mean to be at pixel x,y? What does it mean to be more or less bright in the Fourier decomposition image?
Brian Pauw demo

- Live Fourier decomposition images
 Using FFT2 function
- I hacked it a bit
- http://www.lookingatnothing.com/index.php/ archives/991

Hoiem

Think-Pair-Share

Match the spatial domain image to the Fourier amplitude image





Fourier Bases

Teases away 'fast vs. slow' changes in the image.



This change of basis is the Fourier Transform

Basis reconstruction



Full image



First 9 basis fns

First 1 basis fn



First 16 basis fns



First 4 basis fns



First 400 basis fns

Danny Alexander

Now we can edit frequencies!



Low and High Pass filtering



Removing frequency bands



Brayer

High pass filtering + orientation









Brian Pauw demo

- Live Fourier decomposition images
 Using FFT2 function
- This time, for editing images.

Fourier Transform

- Stores the amplitude and phase at each frequency:
 - For mathematical convenience, this is often notated in terms of real and complex numbers
 - Related by Euler's formula



Fourier Transform

- Stores the amplitude and phase at each frequency:
 - For mathematical convenience, this is often notated in terms of real and complex numbers
 - Related by Euler's formula



Amplitude encodes how much signal there is at a particular frequency:

$$A = \pm \sqrt{\operatorname{Re}(\varphi)^2 + \operatorname{Im}(\varphi)^2}$$

Phase encodes spatial information (indirectly):

$$\phi = \tan^{-1} \frac{\operatorname{Im}(\varphi)}{\operatorname{Re}(\varphi)}$$

Amplitude / Phase



- Amplitude tells you "how much"
- Phase tells you "where"
- Translate the image?
 - Amplitude unchanged
 - Adds a constant to the phase.

Morse [Peppergrower; Wikipedia]



Amplitude

Phase





Amplitude

Phase

Efros

John Brayer, Uni. New Mexico

 "We generally do not display PHASE images because most people who see them shortly thereafter succumb to hallucinogenics or end up in a Tibetan monastery."

https://www.cs.unm.edu/~brayer/vision/fourier.html

Think-Pair-Share

- In Fourier space, where is more of the information that we see in the visual world?
 - Amplitude
 - Phase

Cheebra

Zebra phase, cheetah amplitude

Cheetah phase, zebra amplitude



The frequency amplitude of natural images are quite similar.

- Heavy in low frequencies, falling off in high frequencies
- Will any image be like that, or is it a property of the world we live in?

Most information in the image is carried in the phase, not the amplitude.

What is the relationship to phase in audio?

- In audio perception, frequency is important but phase is not.
- In visual perception, both are important.

• ??? :(

Properties of Fourier Transforms

• Linearity F[ax(t)+by(t)] = a F[x(t)]+b F[y(t)]

• Fourier transform of a real signal is symmetric about the origin

• The energy of the signal is the same as the energy of its Fourier transform

The Convolution Theorem

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

$$\mathbf{F}[g * h] = \mathbf{F}[g]\mathbf{F}[h]$$

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

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

Filtering in spatial domain

10-120-210-1



Convolution



Slide: Hoiem

Why does the Gaussian filter give a nice smooth image, but the square filter give edgy artifacts?



Why do we have those lines in the image?

Sharp edges in an image or kernel need _all_ frequencies to represent them.



Box filter / sinc filter duality

What is the box filter in the frequency domain? http://madebyevan.com/dft/



Hays

Evan Wallace demo

- Made for CSCI1230
- 1D example
- Forbes 30 under 30

- Figma (collaborative design tools)

http://madebyevan.com/dft/





with Dylan Field

Box filter / sinc filter duality

What is the box filter in the frequency domain? What is the sinc filter in the frequency domain?



Spatial Domain \iff Frequency Domain Frequency Domain \iff Spatial Domain



Gaussian filter duality

- Fourier transform of one Gaussian... ...is another Gaussian (with inverse variance).
- Why is this useful?
 - Smooth degradation in frequency components
 - No sharp cut-off
 - No negative values
 - Never zero (infinite extent)



Frequency domain magnitude



Ringing artifacts -> 'Gibbs effect' Where *infinite* series can never be reached



Is convolution invertible?

• If convolution is just multiplication in the Fourier domain, isn't deconvolution just division?

- Sometimes, it clearly is invertible (e.g. a convolution with an identity filter)
- In one case, it clearly isn't invertible (e.g. convolution with an all zero filter)
- What about for common filters like a Gaussian?

Convolution



10

8

6

2

0

-2

-4





FFT











Deconvolution?



-40

10
But under more realistic conditions



200

400

600

1000

800

Hays

-10

-20

-30

-40

But under more realistic conditions



-10

-20

-30

-40

Random noise, .0001 magnitude

But under more realistic conditions



Hays

Deconvolution is hard.

• Active research area.

• Even if you know the filter (non-blind deconvolution), it is still hard and requires strong *regularization* to counteract noise.

• If you don't know the filter (blind deconvolution), then it is harder still.

Index cards before you leave

Fourier decomposition is tricky I want to know what is confusing to you!

- Take 1 minute to talk to your partner about the lecture

- Write down on the card something that you'd like clarifying.

Thinking in Frequency - Compression

How is it that a 4MP image can be compressed to a few hundred KB without a noticeable change?

Lossy Image Compression (JPEG)

8x8 blocks



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



Block-based Discrete Cosine Transform (DCT)

Slides: Efros

Image compression using DCT

• Compute DCT filter responses in each 8x8 block

Filter responses

- -30.19 61.2027.2456.13 - 20.10 - 2.39-415.380.464.47 - 21.86 - 60.76 10.2513.15-7.09 -8.544.88 $G = \begin{bmatrix} 4.47 & -21.80 & -00.70 & 10.25 & 10.15 \\ -46.83 & 7.37 & 77.13 & -24.56 & -28.91 \\ -48.53 & 12.07 & 34.10 & -14.76 & -10.24 \\ 12.12 & -6.55 & -13.20 & -3.95 & -1.88 \\ -7.73 & 2.91 & 2.38 & -5.94 & -2.38 \\ -1.03 & 0.18 & 0.42 & -2.42 & -0.88 \\ -0.17 & 0.14 & -1.07 & -4.19 & -1.17 \end{bmatrix}$ 9.935.42-5.656.301.831.951.75-2.793.140.944.301.85-3.024.12-0.66-0.100.501.68
- Quantize to integer (div. by magic number; round)
 - More coarsely for high frequencies (which also tend to have smaller values)
 - Many quantized high frequency values will be zero

Quantization divisers (element-wise)

	16	11	10	16	24	40	51	61	
Q =	12	12	14	19	26	58	60	55	
	14	13	16	24	40	57	69	56	
	14	17	22	29	51	87	80	62	
	18	22	37	56	68	109	103	77	
	24	35	55	64	81	104	113	92	
	49	64	78	87	103	121	120	101	
	72	92	95	98	112	100	103	99	

Quantized values

JPEG Encoding

• Entropy coding (Huffman-variant)

Quantized values

Linearize *B* like this.



Helps compression:

- We throw away the high frequencies ('0').
 - The zig zag pattern increases in frequency space, so long runs of zeros.







T (Cb=0.5,Cr=0.5)





Cr

(Y=0.5,Cb=05)



Most JPEG images & videos subsample chroma



PSP Comp 3 2x2 Chroma subsampling 285K Original 1,261K lossless 968K PNG

JPEG Compression Summary

- 1. Convert image to YCrCb
- 2. Subsample color by factor of 2
 - People have bad resolution for color
- 3. Split into blocks (8x8, typically), subtract 128
- 4. For each block
 - a. Compute DCT coefficients
 - b. Coarsely quantize
 - Many high frequency components will become zero
 - c. Encode (with run length encoding and then Huffman coding for leftovers)