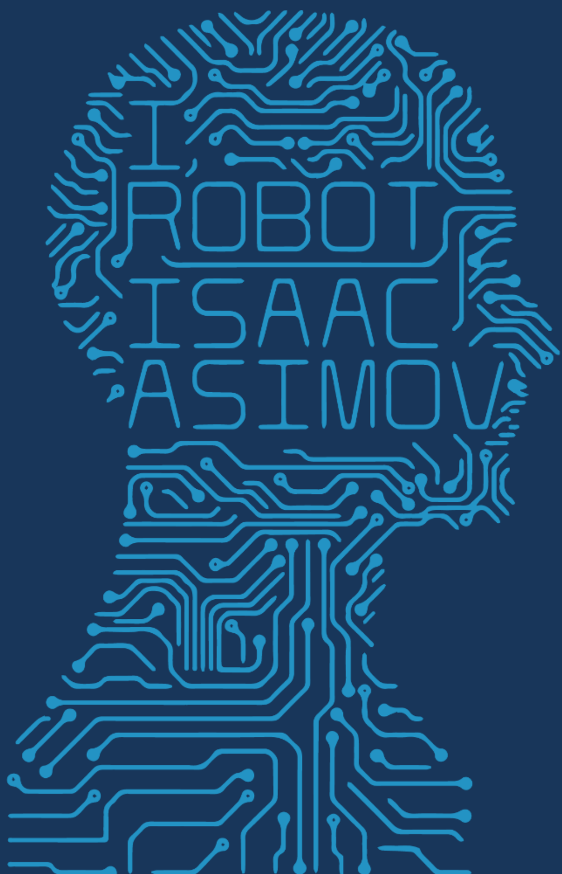

1950

Future Vision

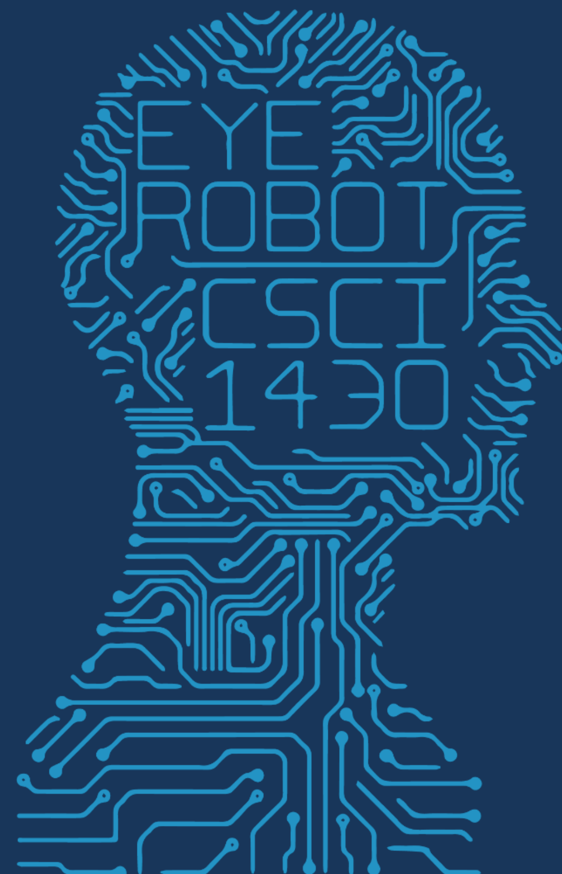

2017 MWF 1pm

Computer Vision


Ken-ichi Ueda
godutchbaby
flapack
tueda
Mark Wilkinson
Ken-ichi Ueda
antonyw
flapack
Mark Wilkinson
annetanne
drew_avery
flapack
Mark Wilkinson

Sharir
Snaps
Putting our vi
As you walk dow
curiosity for life.
a global commun
and noticing patt
citizen scientists
organisms occur-
Every observatior
in your own back
"Our citizen scientists lear
organism. Coloration ma
species, but there are sub
and texture—critical for
— Alison Young
Citizen Science Engagement Co
California Academy of Sciences
From tide pools to urban parks,
diversity is everywhere. On a sp
morning, a crowd fanned out in
Gate Park in search of spiders,
songbirds, and more. A few hou
the group marveled at their feat
than 1,700 observations of 260 s


Sharing Snapshots
Putting our vision to work
As you walk down the street or visit a park, bring your curiosity for life. The Academy and iNaturalist are fostering a global community of ordinary people making observations and noticing patterns in the world around them. These citizen scientists are building a record of where and when organisms occur—data that helps protect Earth's biodiversity. Every observation matters. Who knows what you may find in your own backyard?
"Our citizen scientists learn to look very closely at each organism. Coloration may be extremely similar in some species, but there are subtle differences—in pattern, spotting, and texture—critical for a correct identification."
— Alison Young
California Academy of Sciences

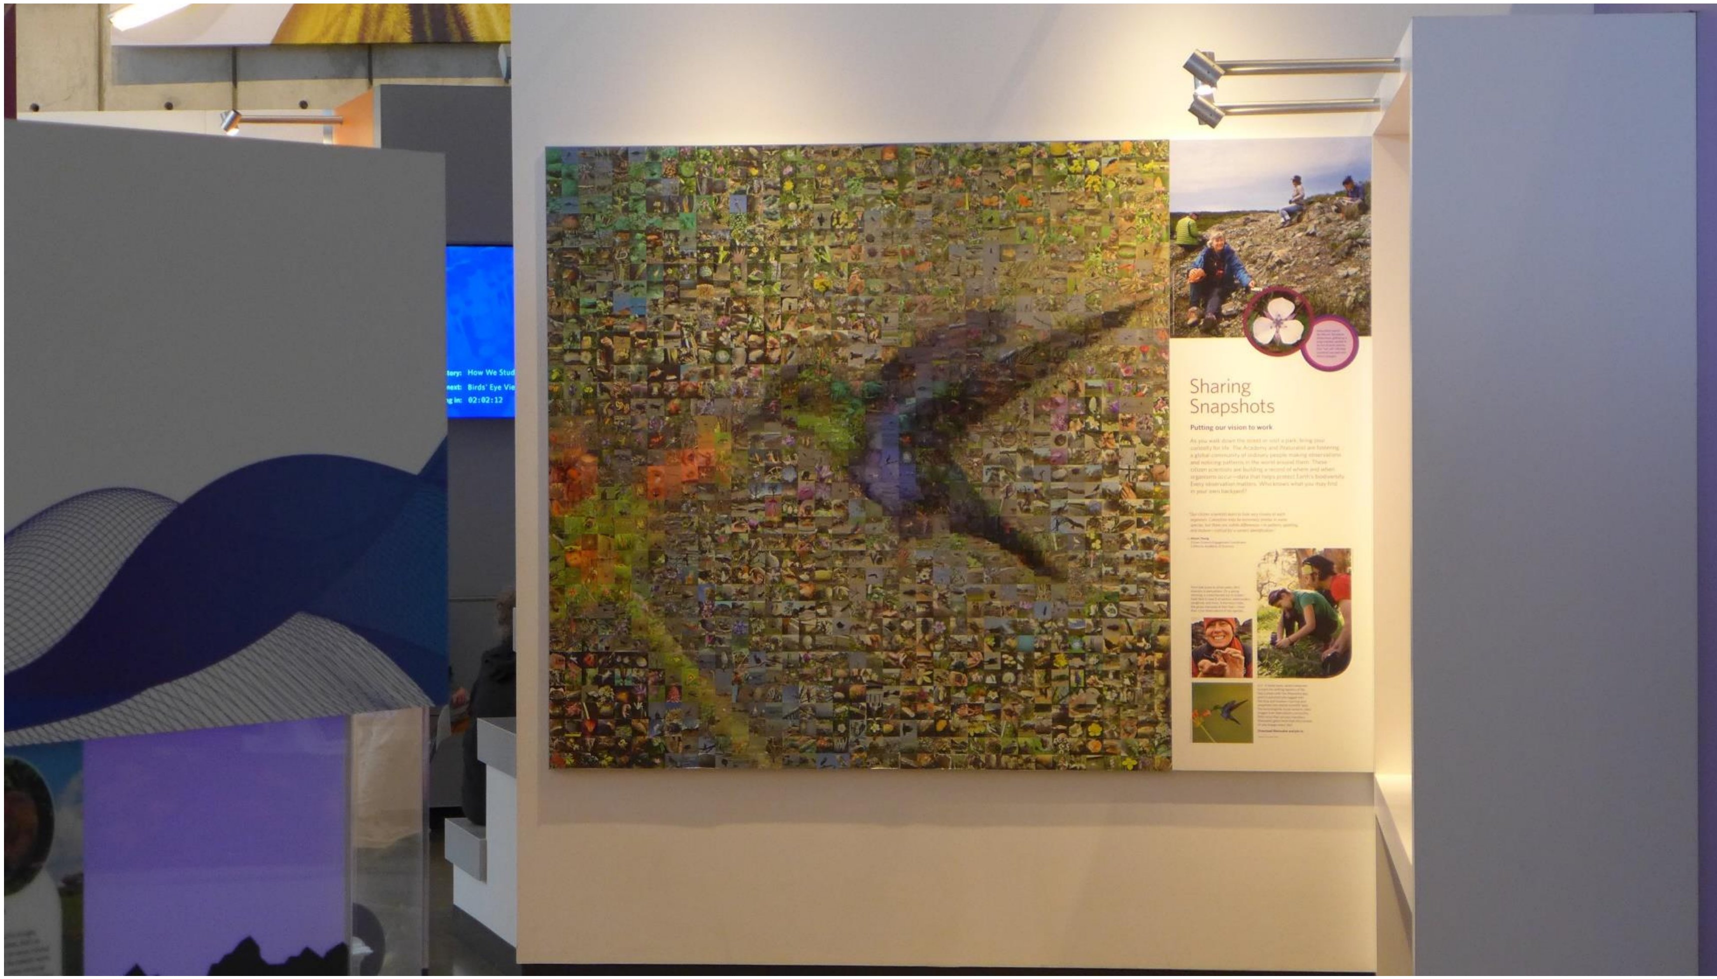

Sharing Snapshots
Putting our vision to work


Seeing in Color
Want to see the world differently? Try on another animal's eyes. Birds and bees see colors of light invisible to the human eye. Other animals detect only a few colors—dogs see shades of yellows, blues, and grays. Explore how every animal eye, from simple to complex, is built to pick up a certain slice of light.
Sending Secret Signals

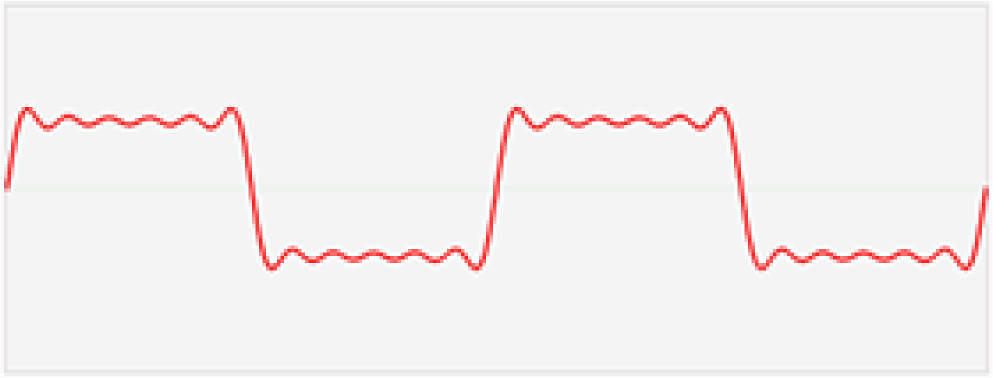

Wikipedia – Fourier transform

Natural image

Natural image

Fourier decomposition
Frequency coefficients (amplitude)

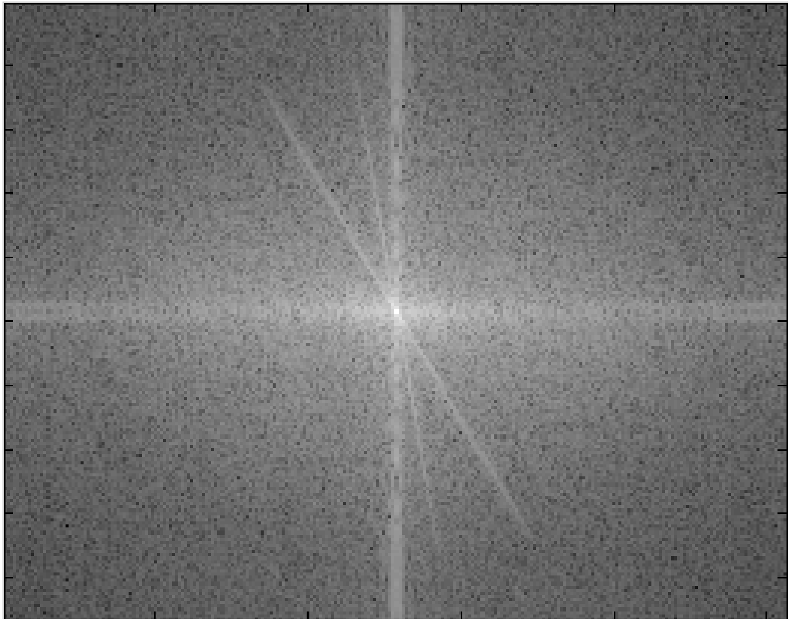

What does it mean to be at pixel x,y?
What does it mean to be more or less bright in the Fourier decomposition image?

Basis reconstruction

Full image

First 1 basis fn

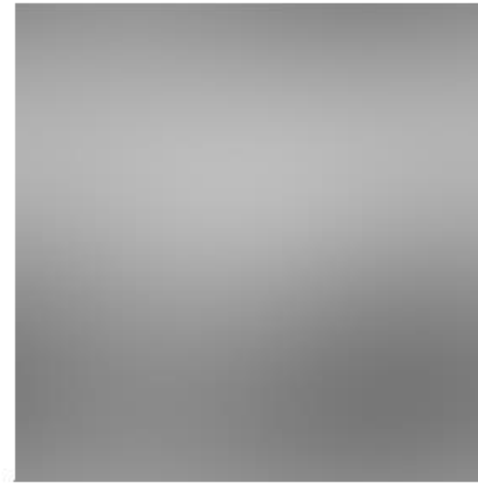

First 4 basis fns

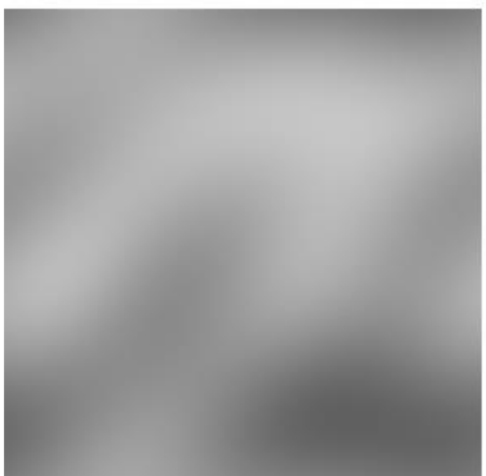

First 9 basis fns

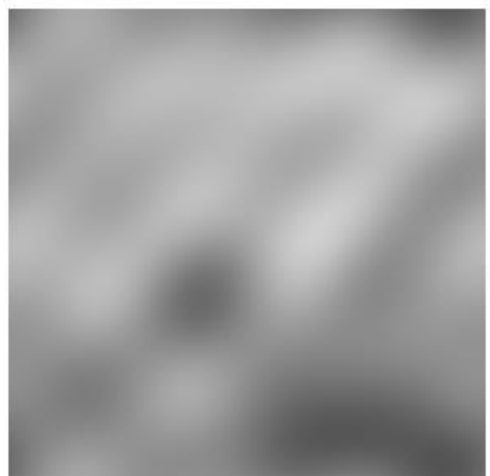

First 16 basis fns

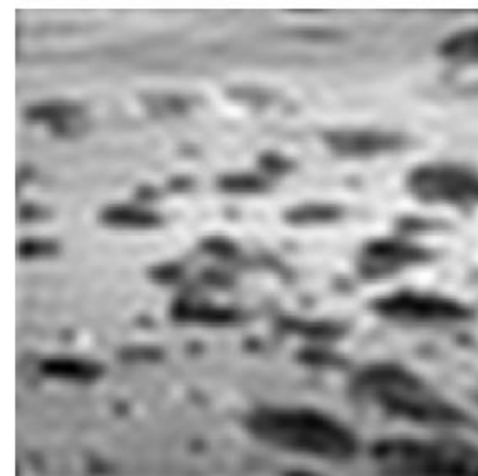

First 400 basis fns

Danny Alexander

Properties of Fourier Transforms

- Linearity $\mathrm{F}[ax(t)+by(t)] = a\,\mathrm{F}[x(t)] + b\,\mathrm{F}[y(t)]$
- 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

See Szeliski Book (3.4)

The Convolution Theorem

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

$$\mathrm{F}[g * h] = \mathrm{F}[g]\,\mathrm{F}[h]$$

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

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

Hays

Filtering in spatial domain

1	0	-1
2	0	-2
1	0	-1

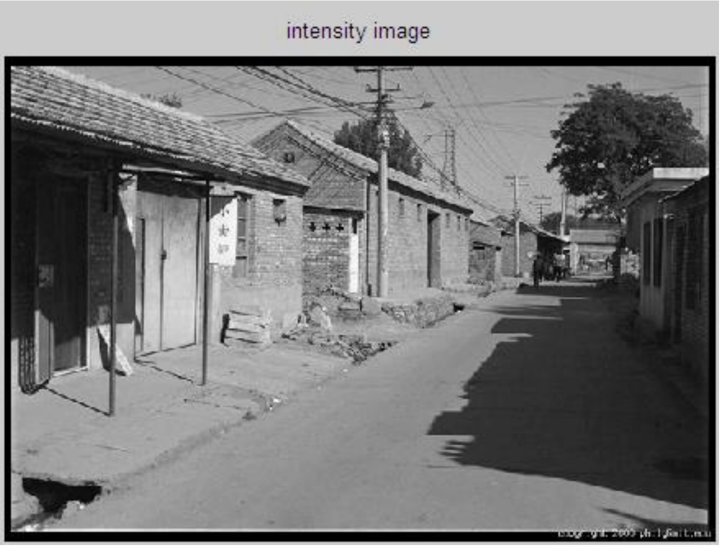

*

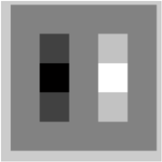

=

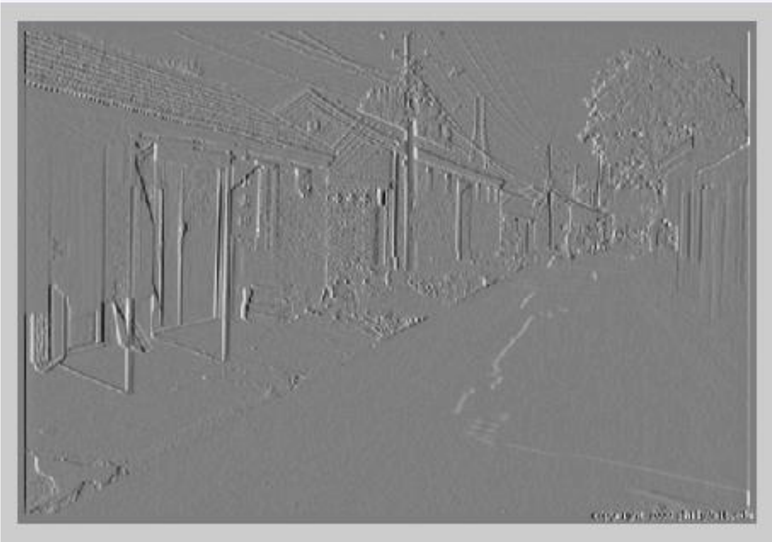

Hays

Filtering in frequency domain

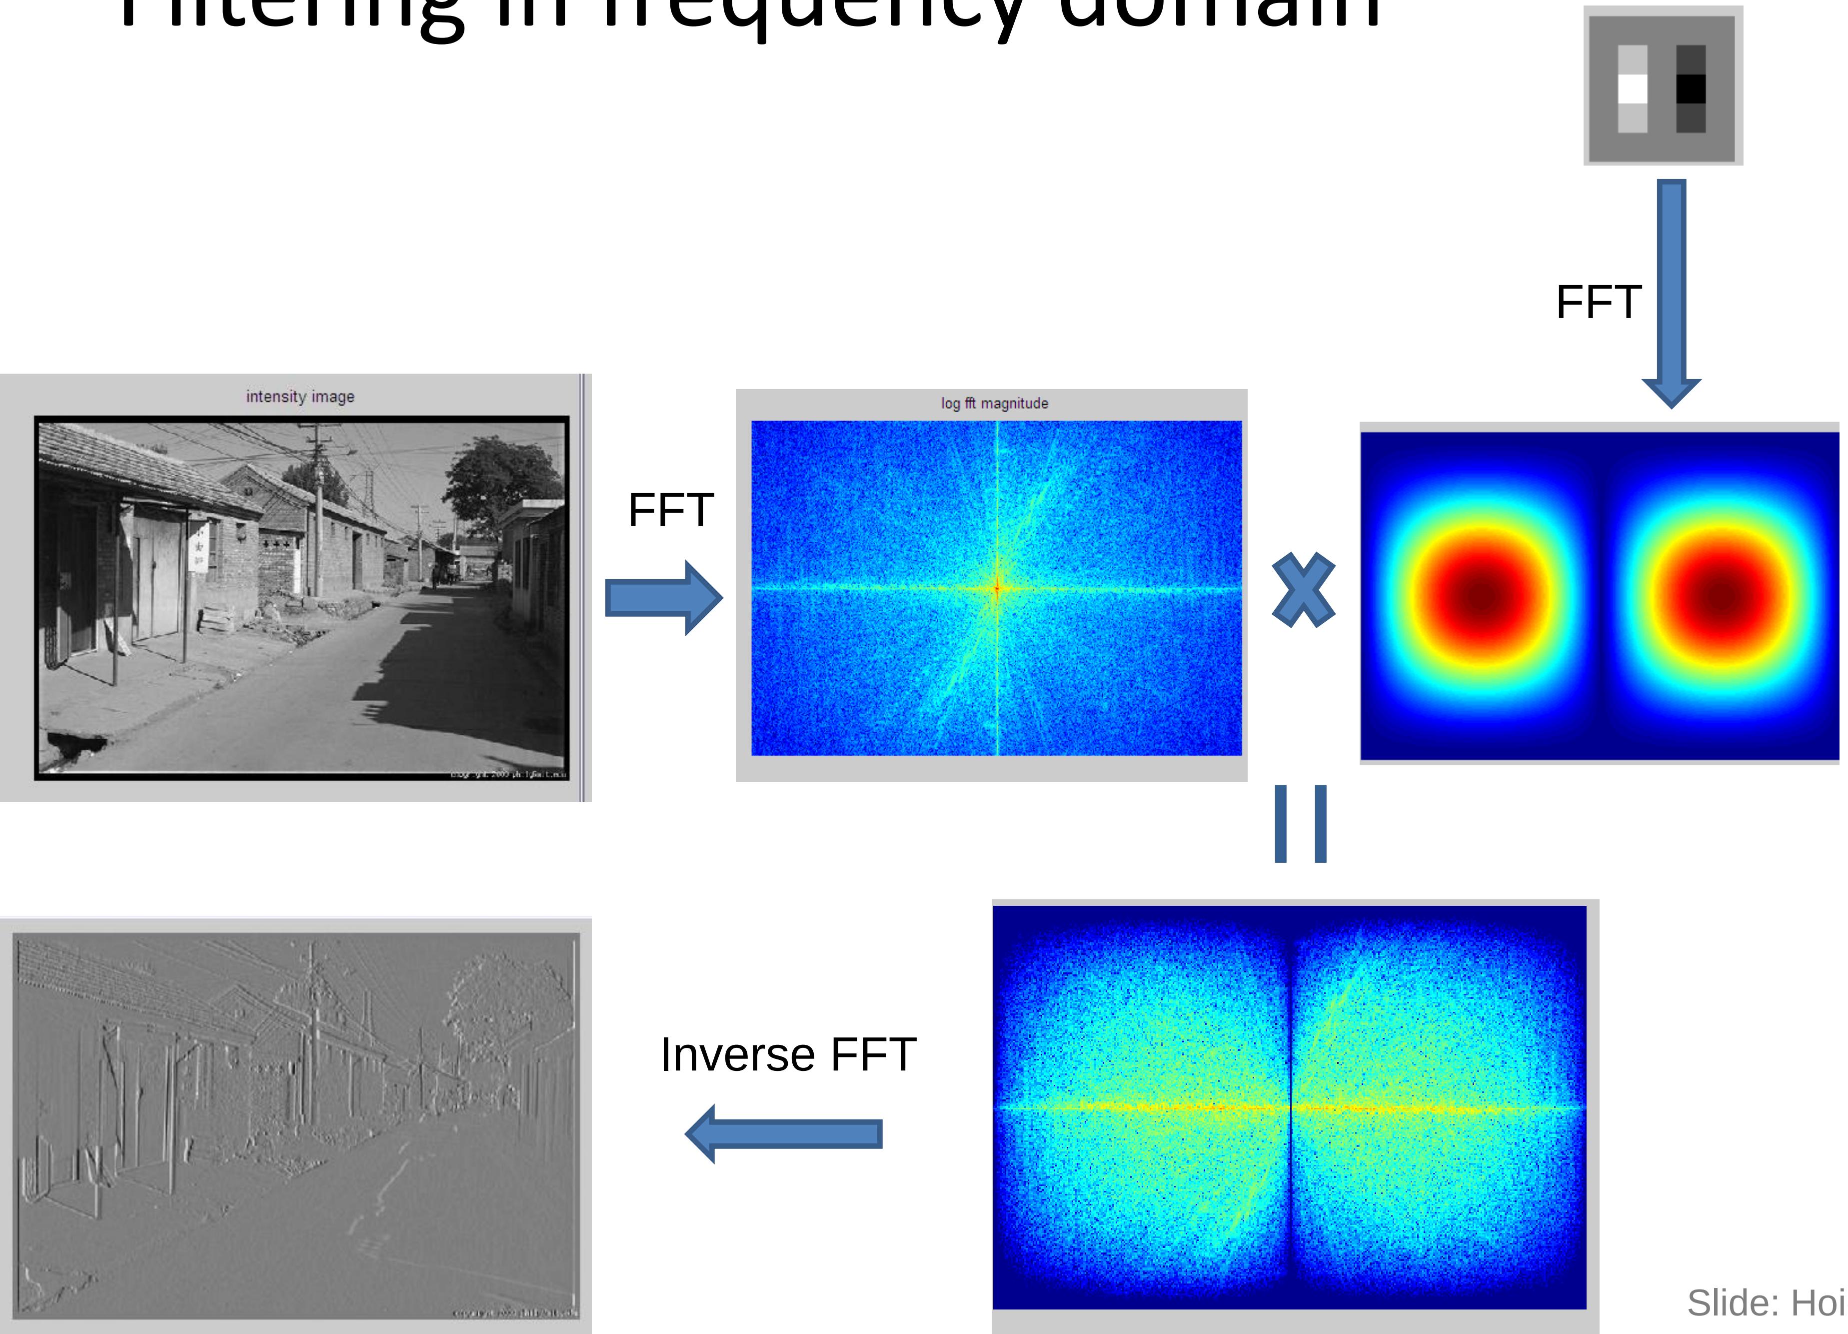

Slide: Hoiem

Now we can edit frequencies!

Low and High Pass filtering

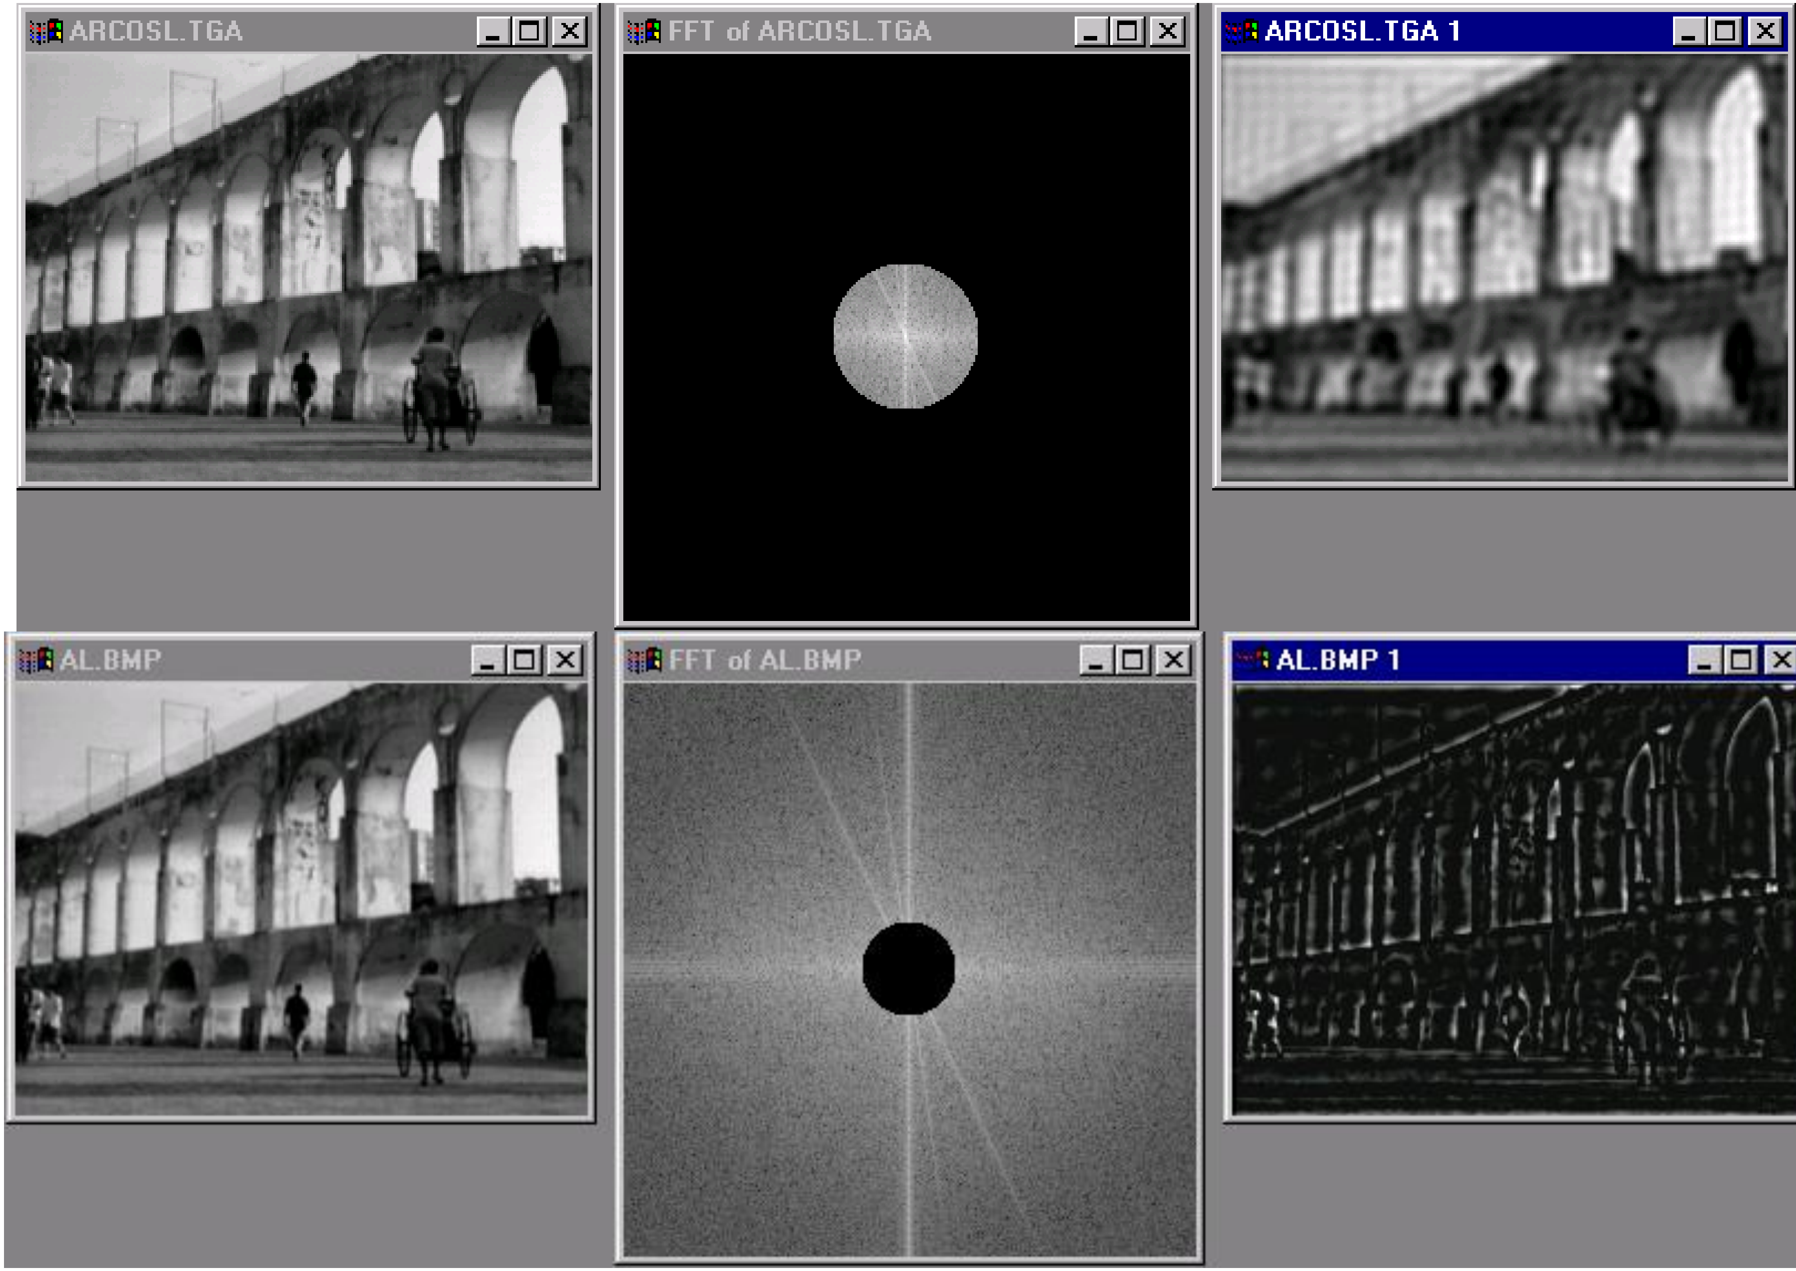

Removing frequency bands

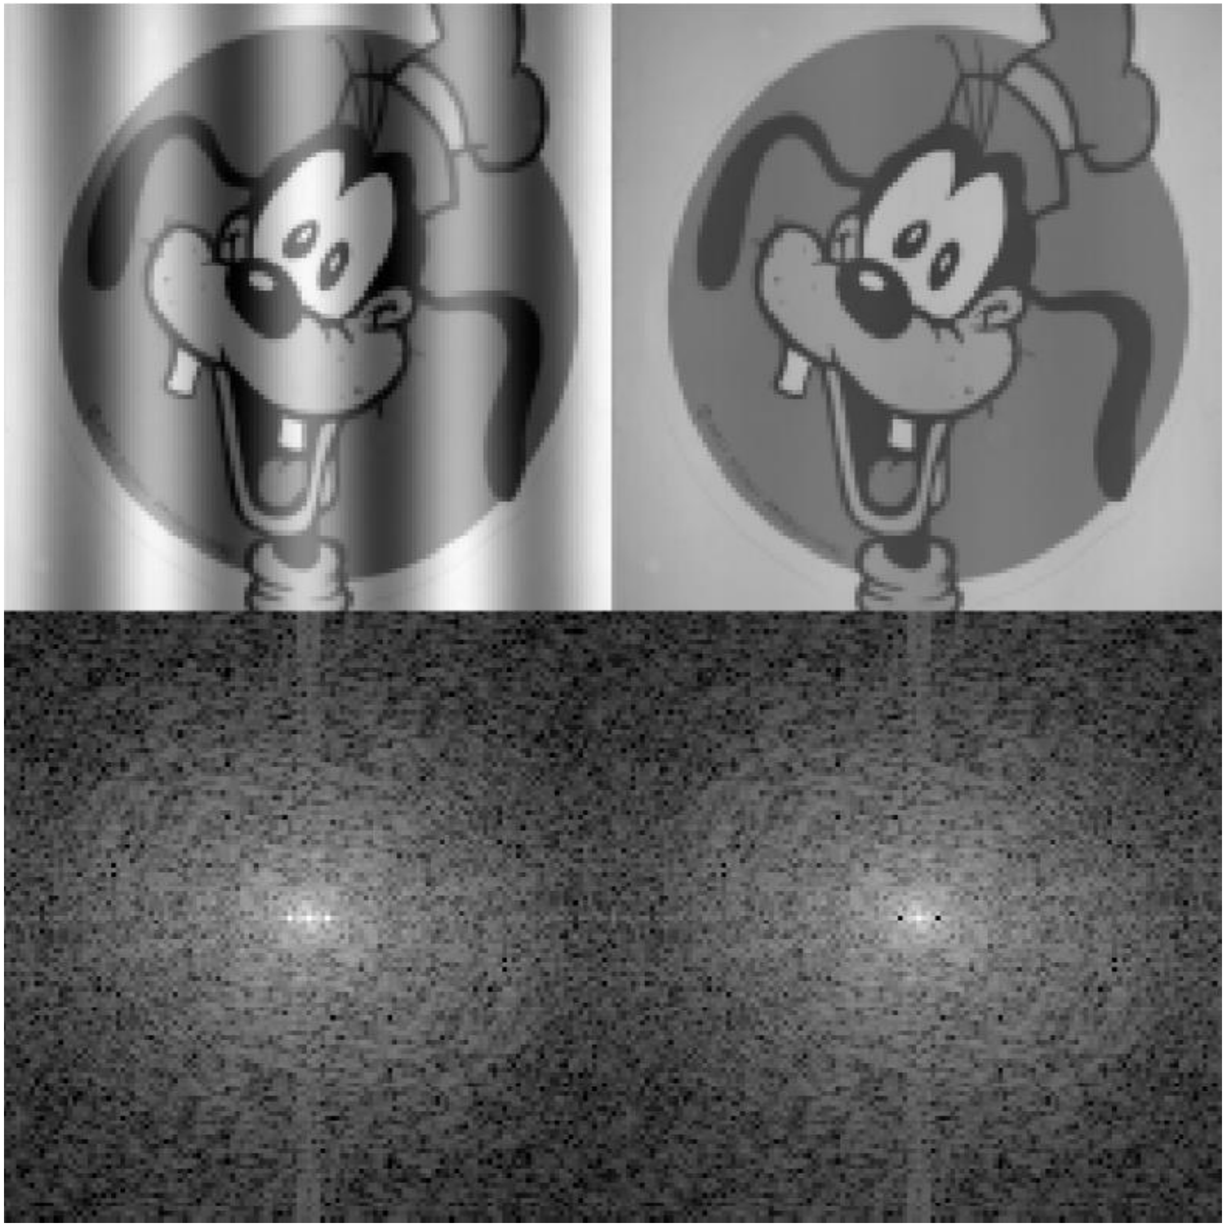

Brayer

High pass filtering + orientation

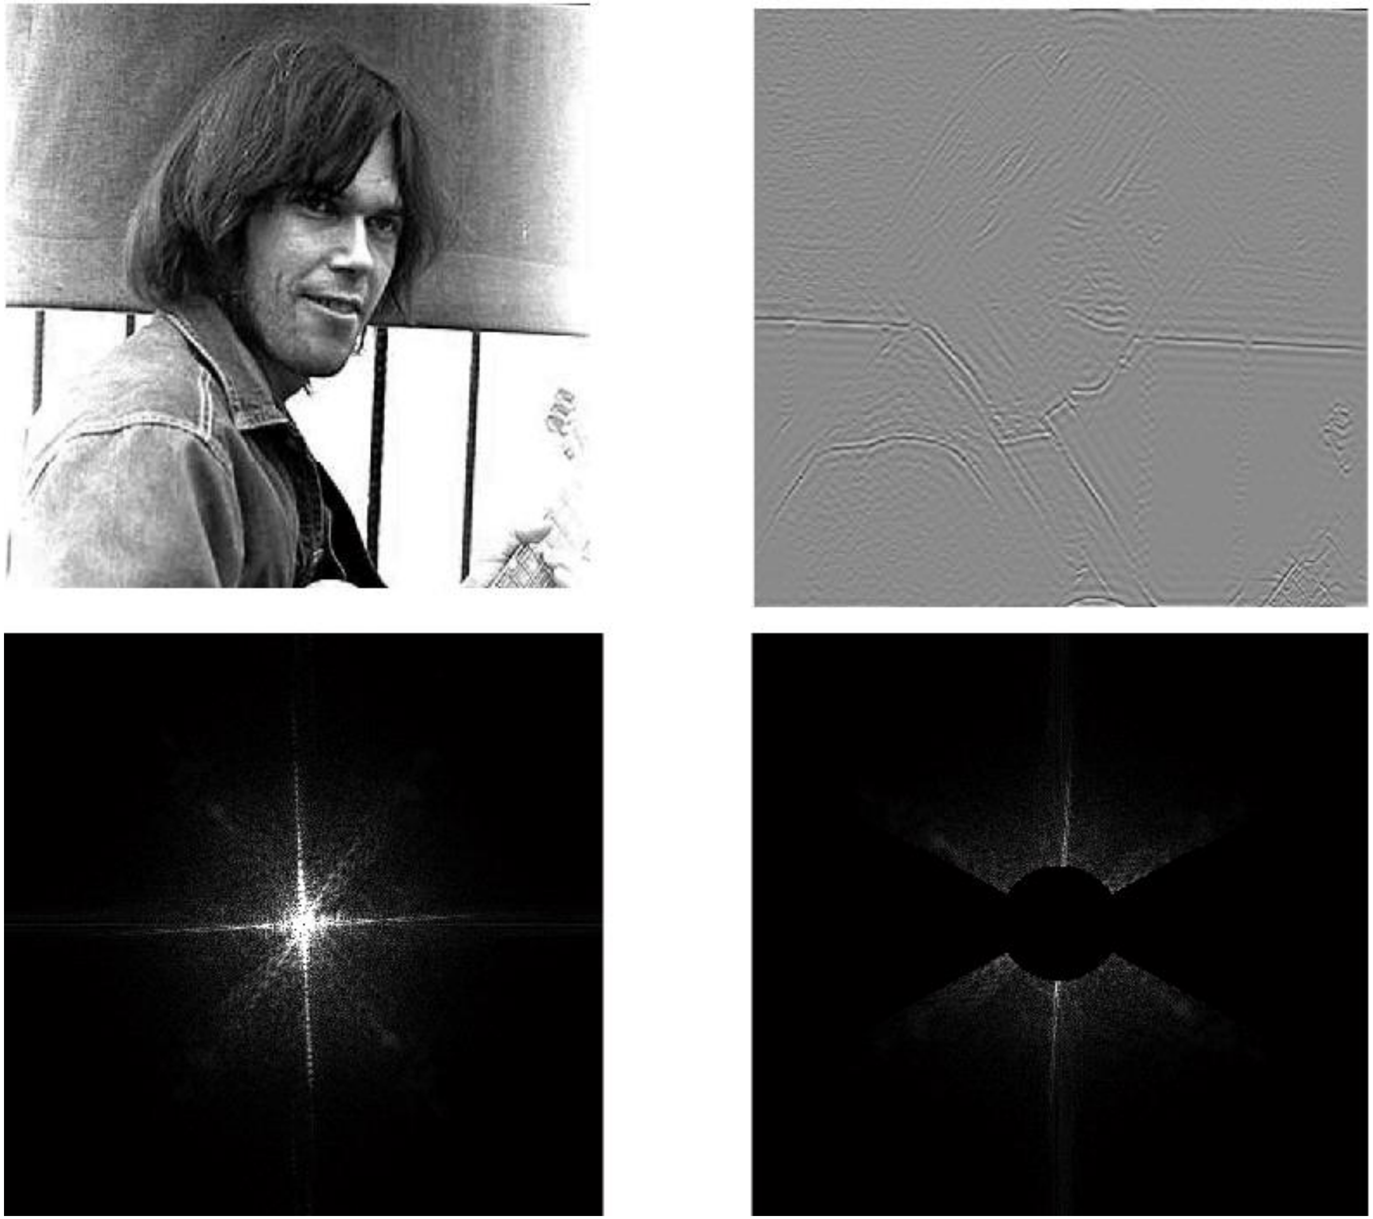

Application: Hybrid Images

When we see an image from far away, we are effectively subsampling it!

A. Oliva, A. Torralba, P.G. Schyns, SIGGRAPH 2006

Hybrid Image in FFT

Hybrid Image

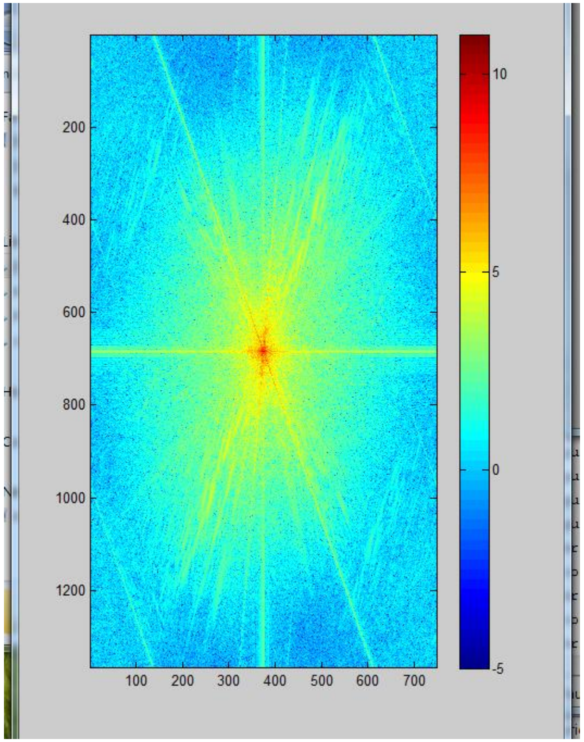

Low-passed Image

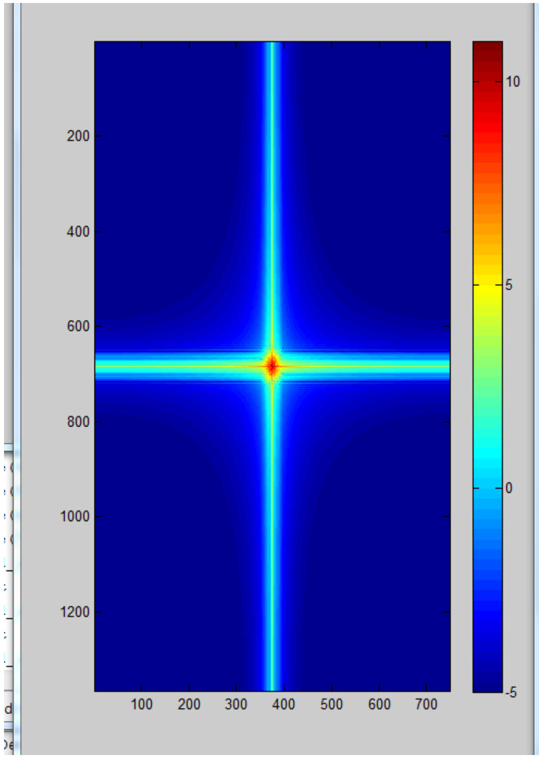

High-passed Image

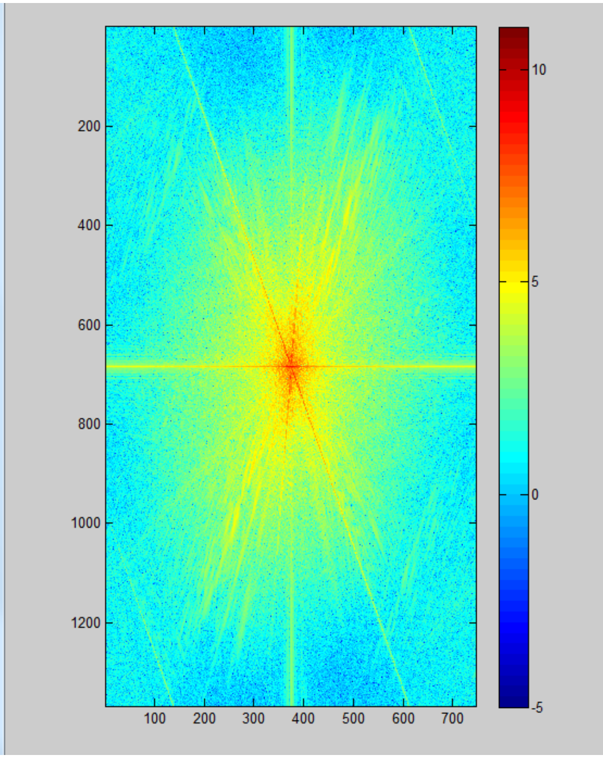

Hays

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

Gaussian 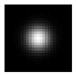

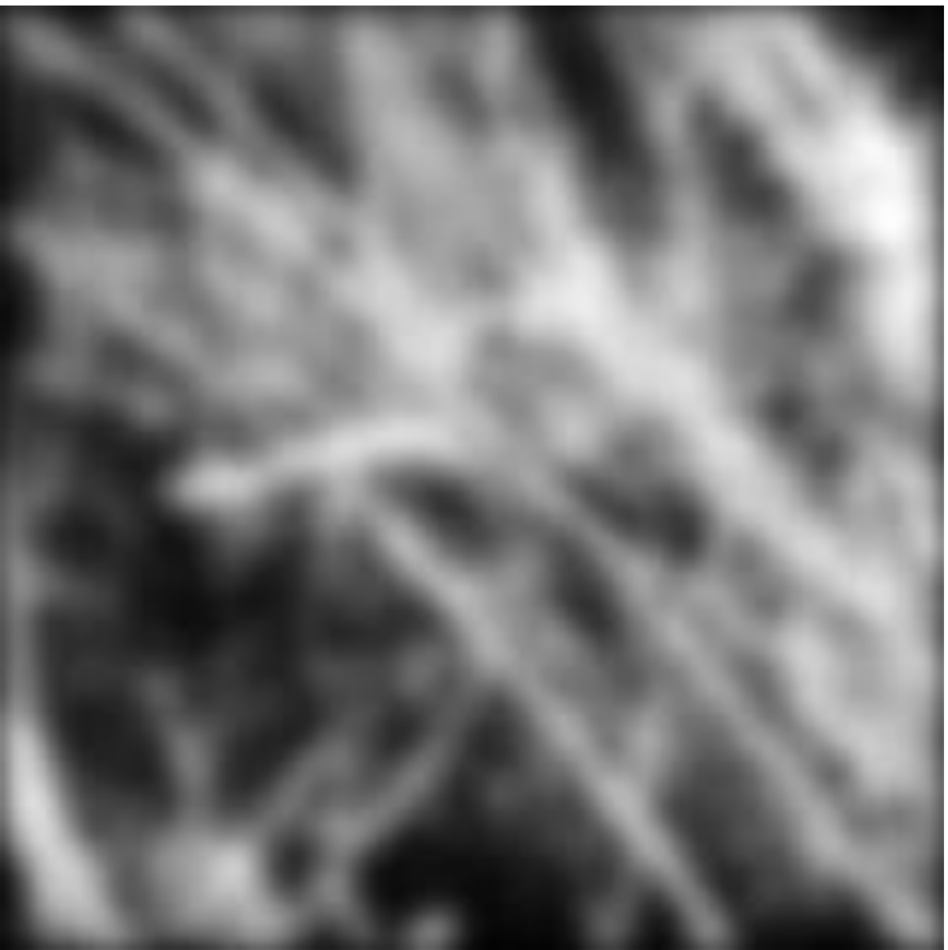

Box filter 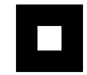

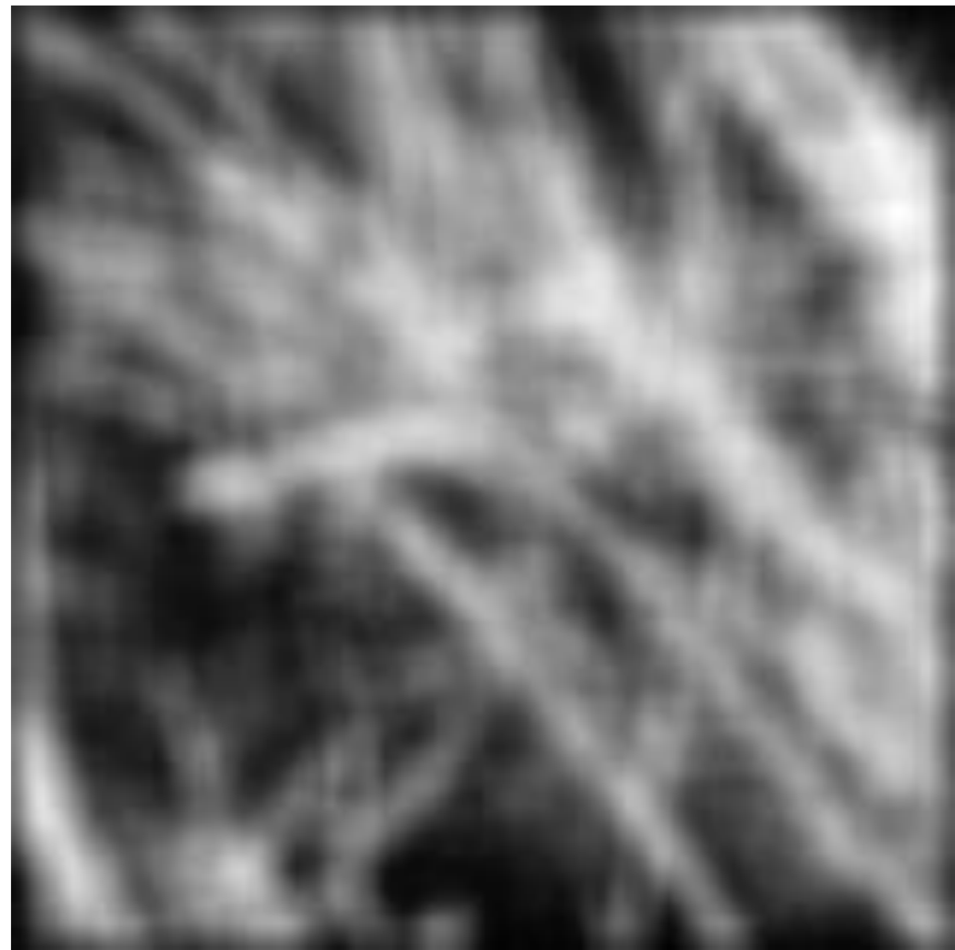

Why do we have those lines in the image?

- Sharp edges in the image need _all_ frequencies to represent them.

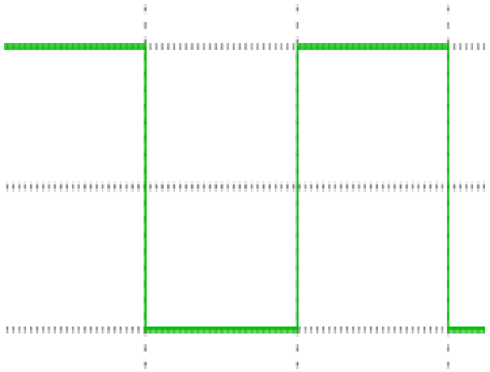

$$= A\sum_{k=1}^{\infty}\frac{1}{k}\sin(2\pi kt)$$

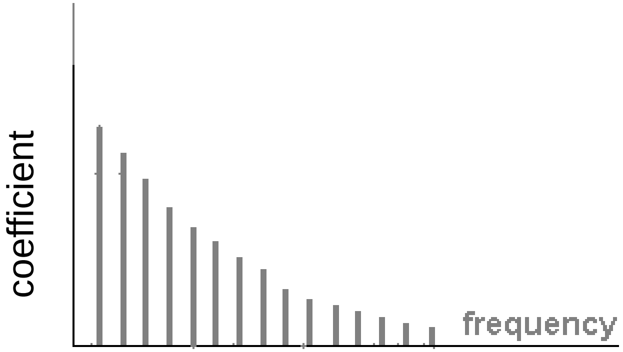

Efros

Box filter / sinc filter duality

- What is the spatial representation of the hard cutoff (box) in the frequency domain?
- http://madebyevan.com/dft/

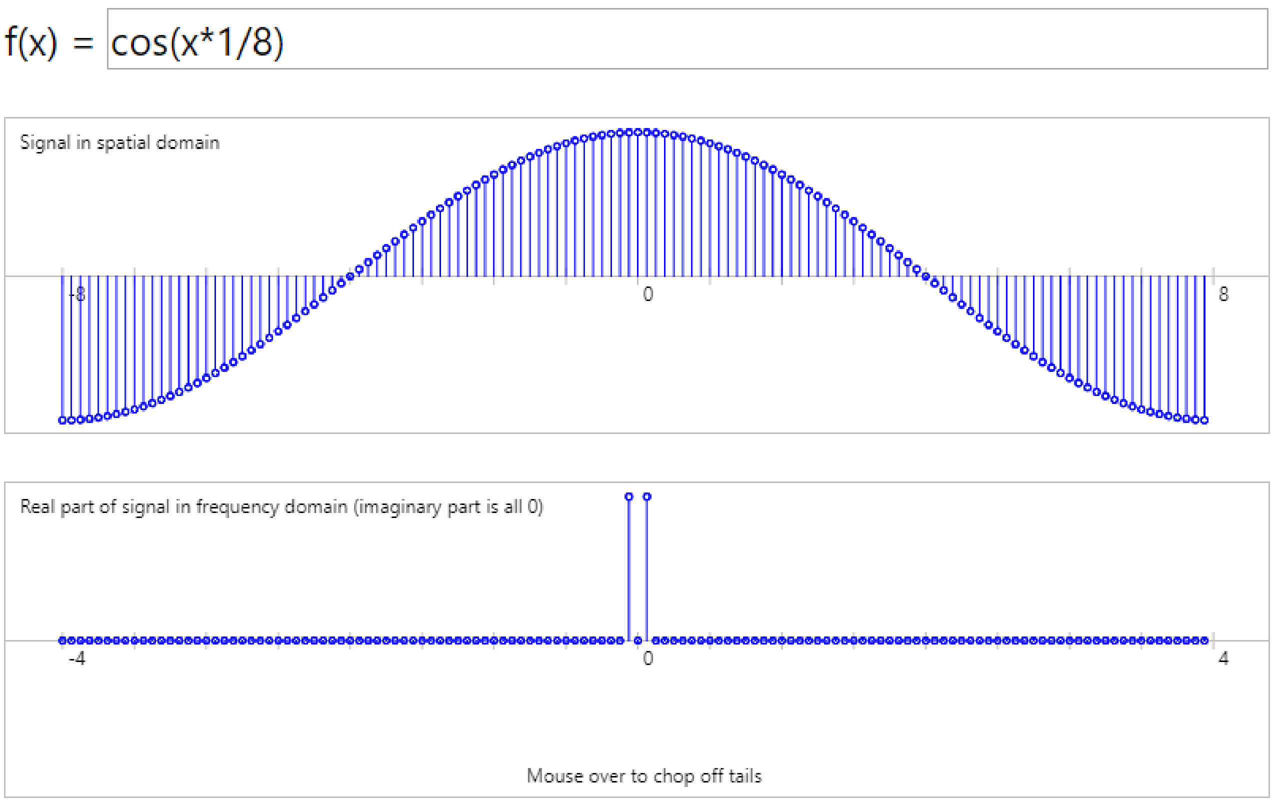

Evan Wallace demo

- Made for CS123
- 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 spatial representation of the hard cutoff (box) in the frequency domain?
- http://madebyevan.com/dft/

Box filter

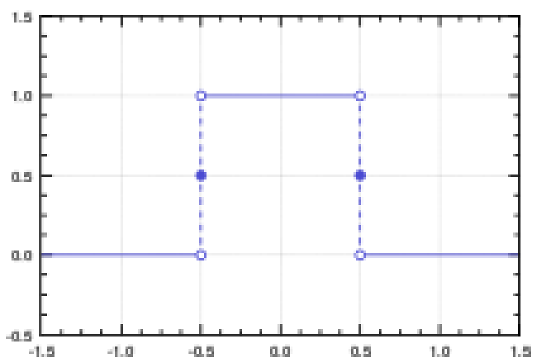

Sinc filter

sinc(x) = sin(x) / x

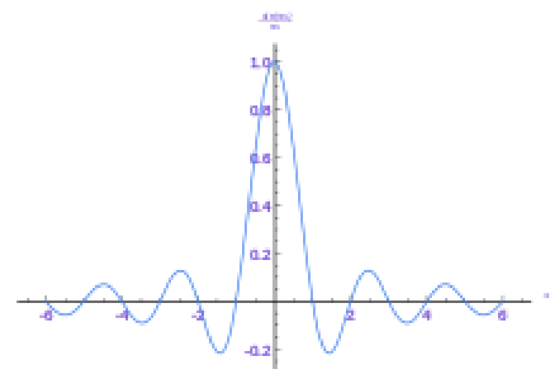

Spatial Domain ⟺ Frequency Domain

Frequency Domain ⟺ Spatial Domain

Box filter (spatial)

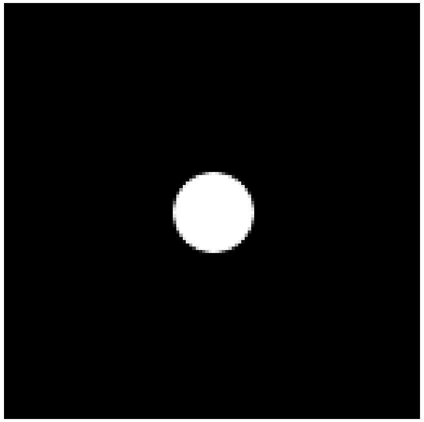

Frequency domain magnitude

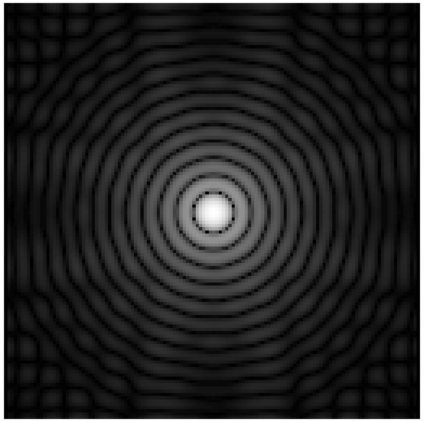

Gaussian filter duality

- 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)

Box filter (spatial)

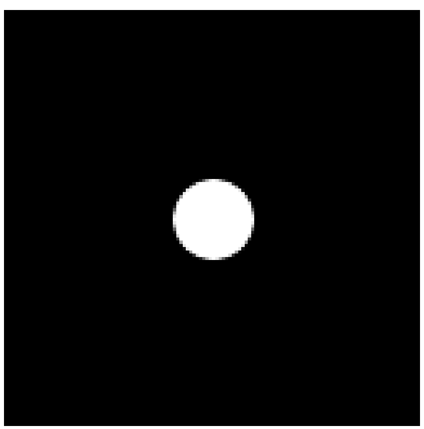

Frequency domain magnitude

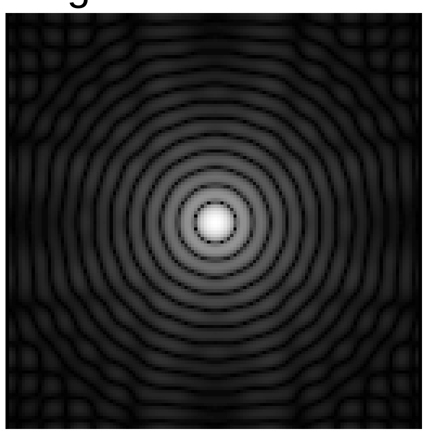

Gaussian filter (spatial)

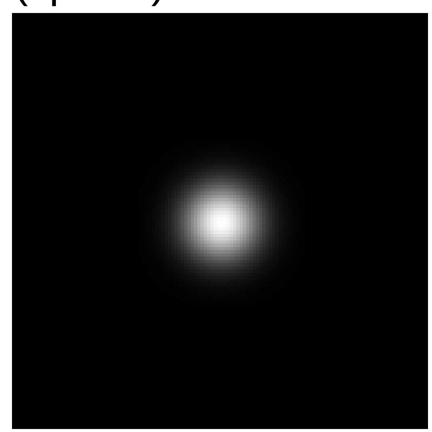

Frequency domain magnitude

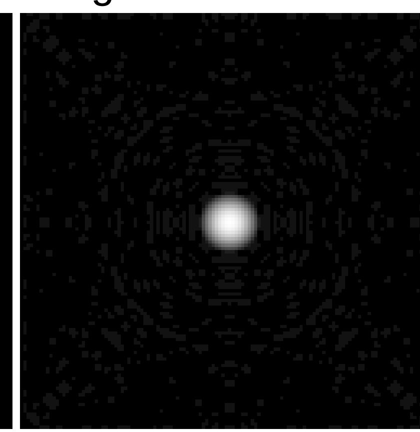

Is convolution invertible?

- 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

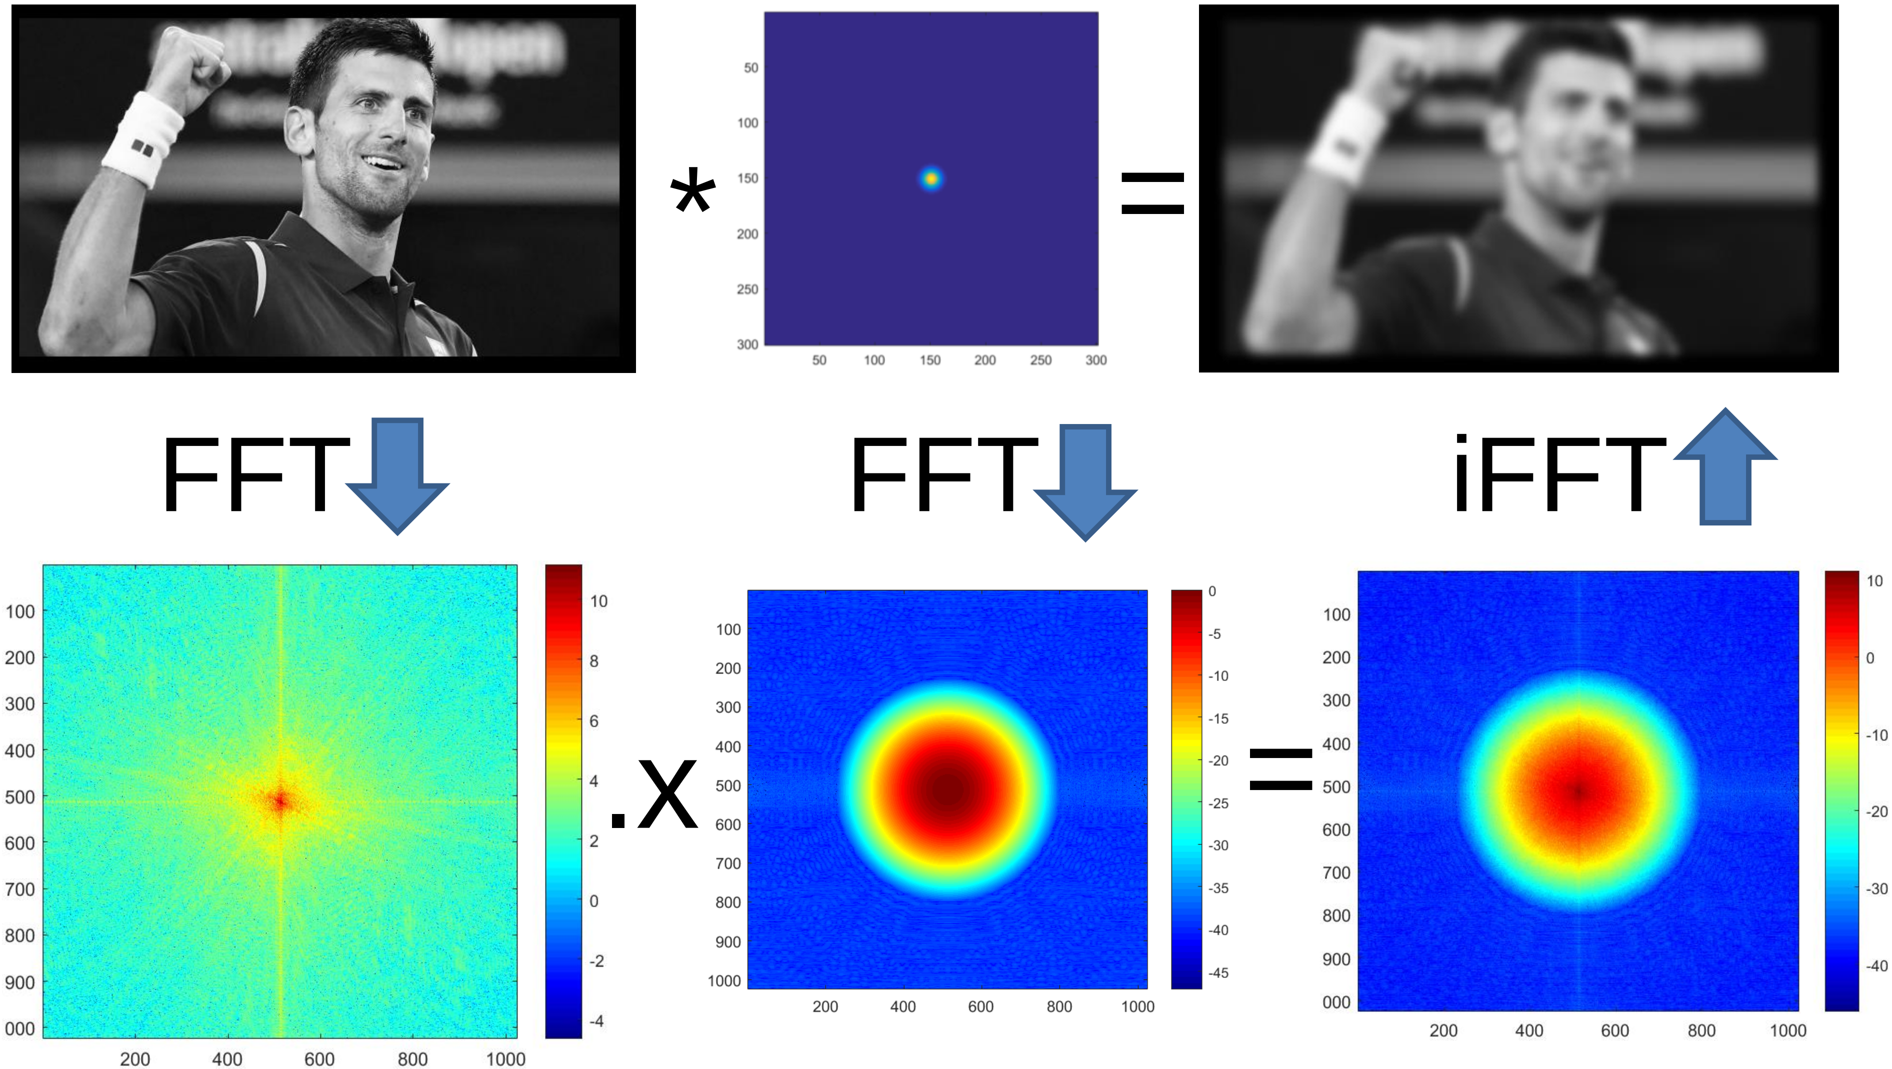

Hays

Deconvolution?

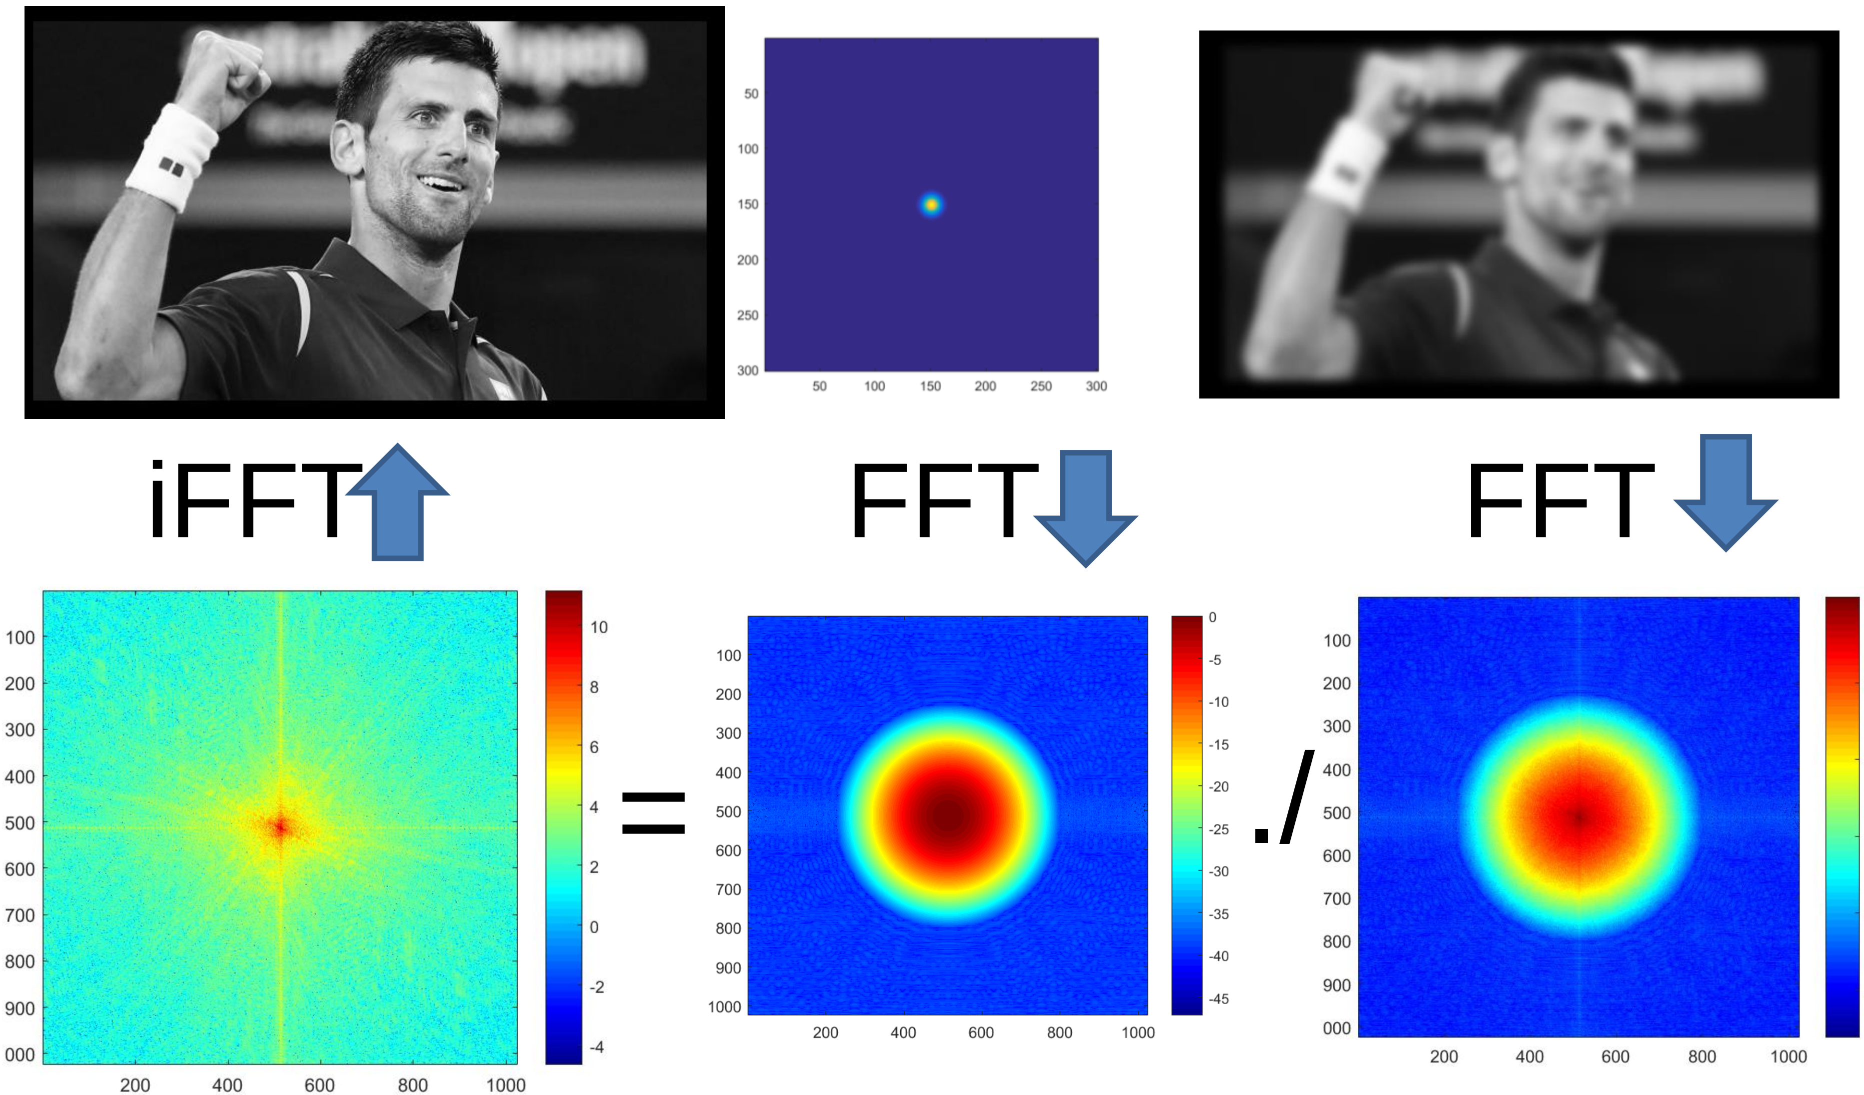

Hays

But under more realistic conditions

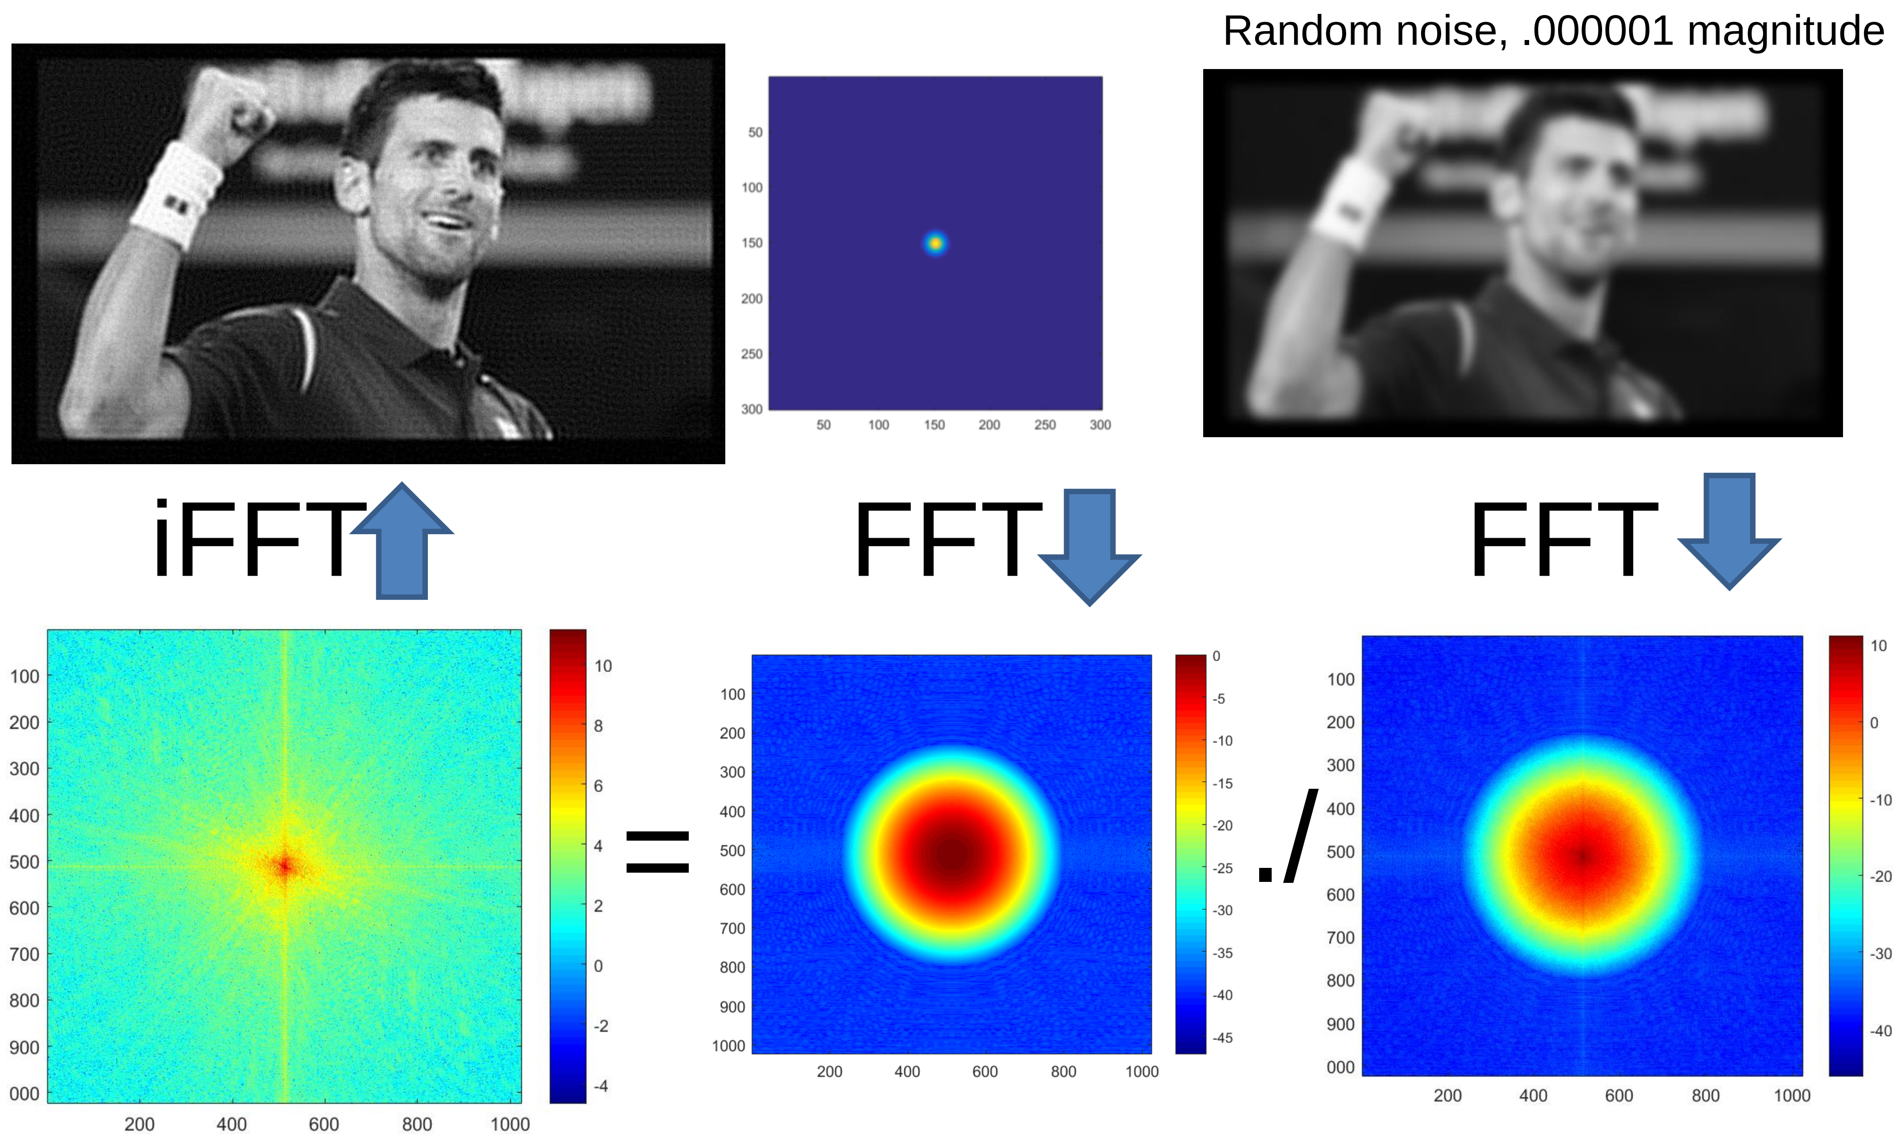

But under more realistic conditions

Random noise, .0001 magnitude

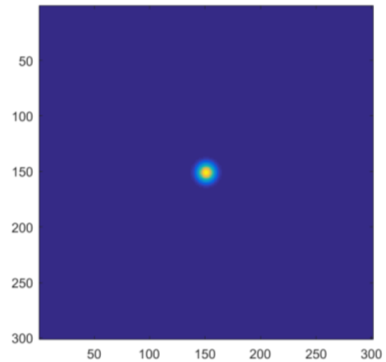

iFFT

FFT

FFT

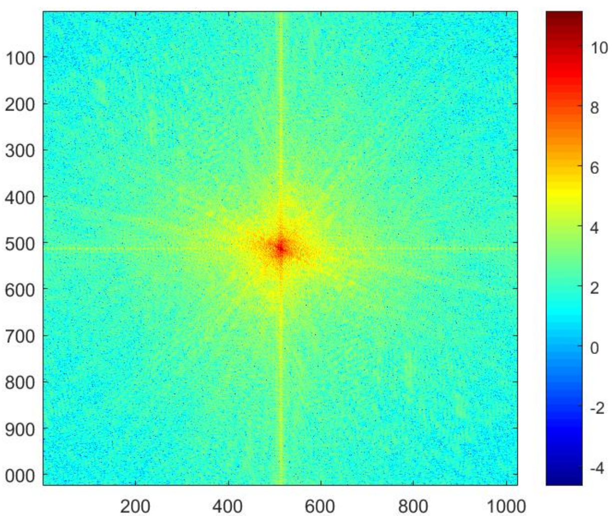

=

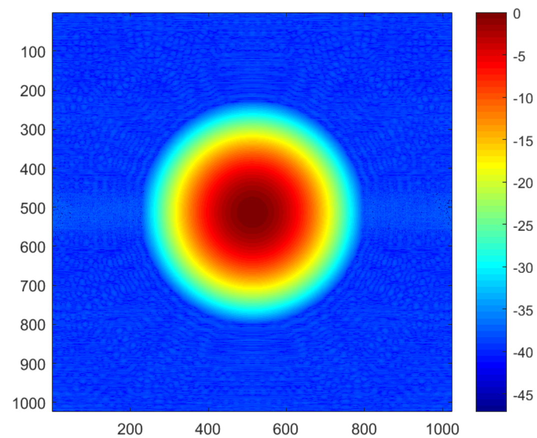

./

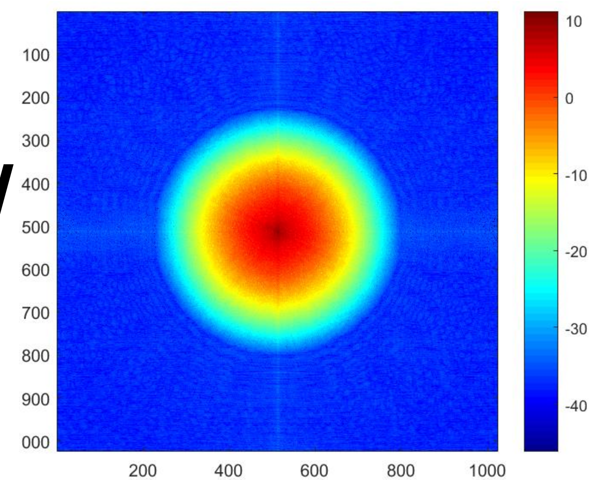

Hays

But under more realistic conditions

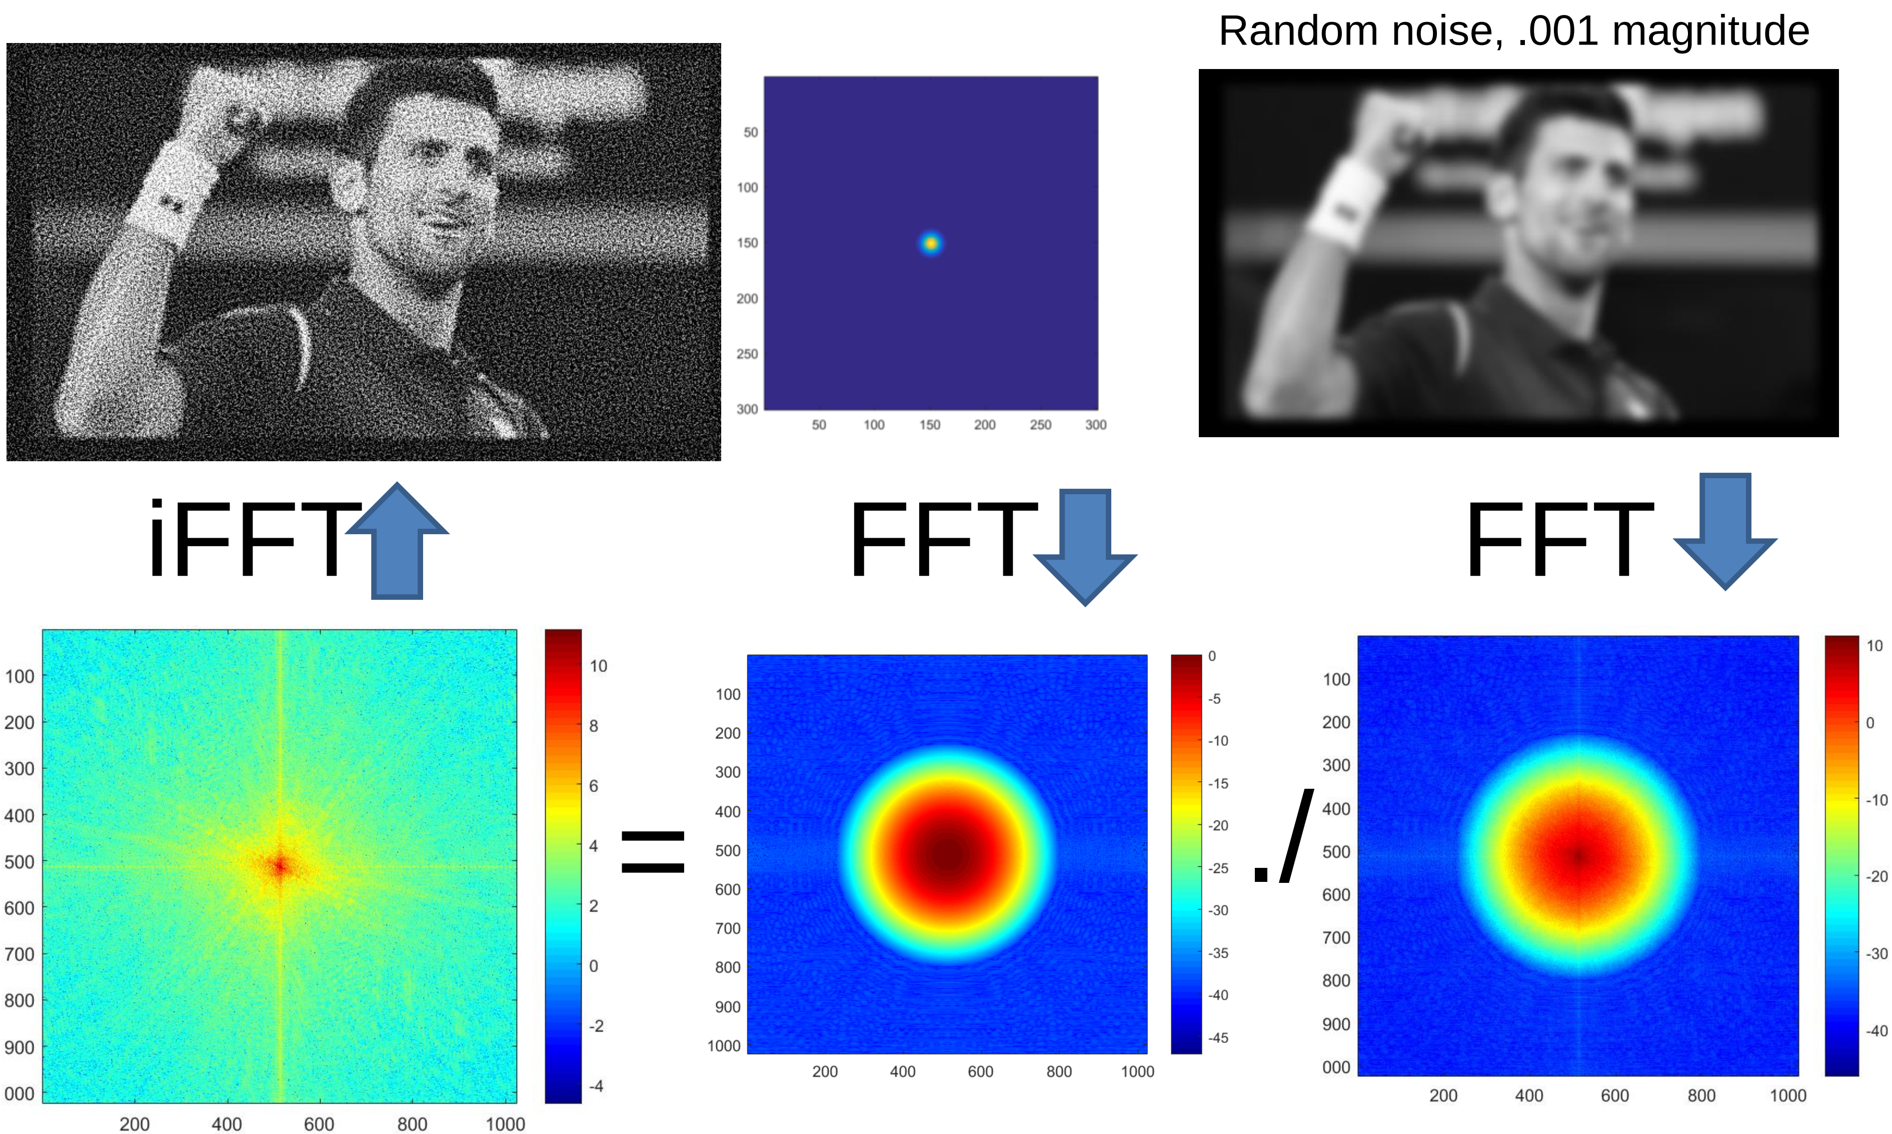

Hays

But you can’t invert multiplication by 0!

- A Gaussian is only zero at infinity…

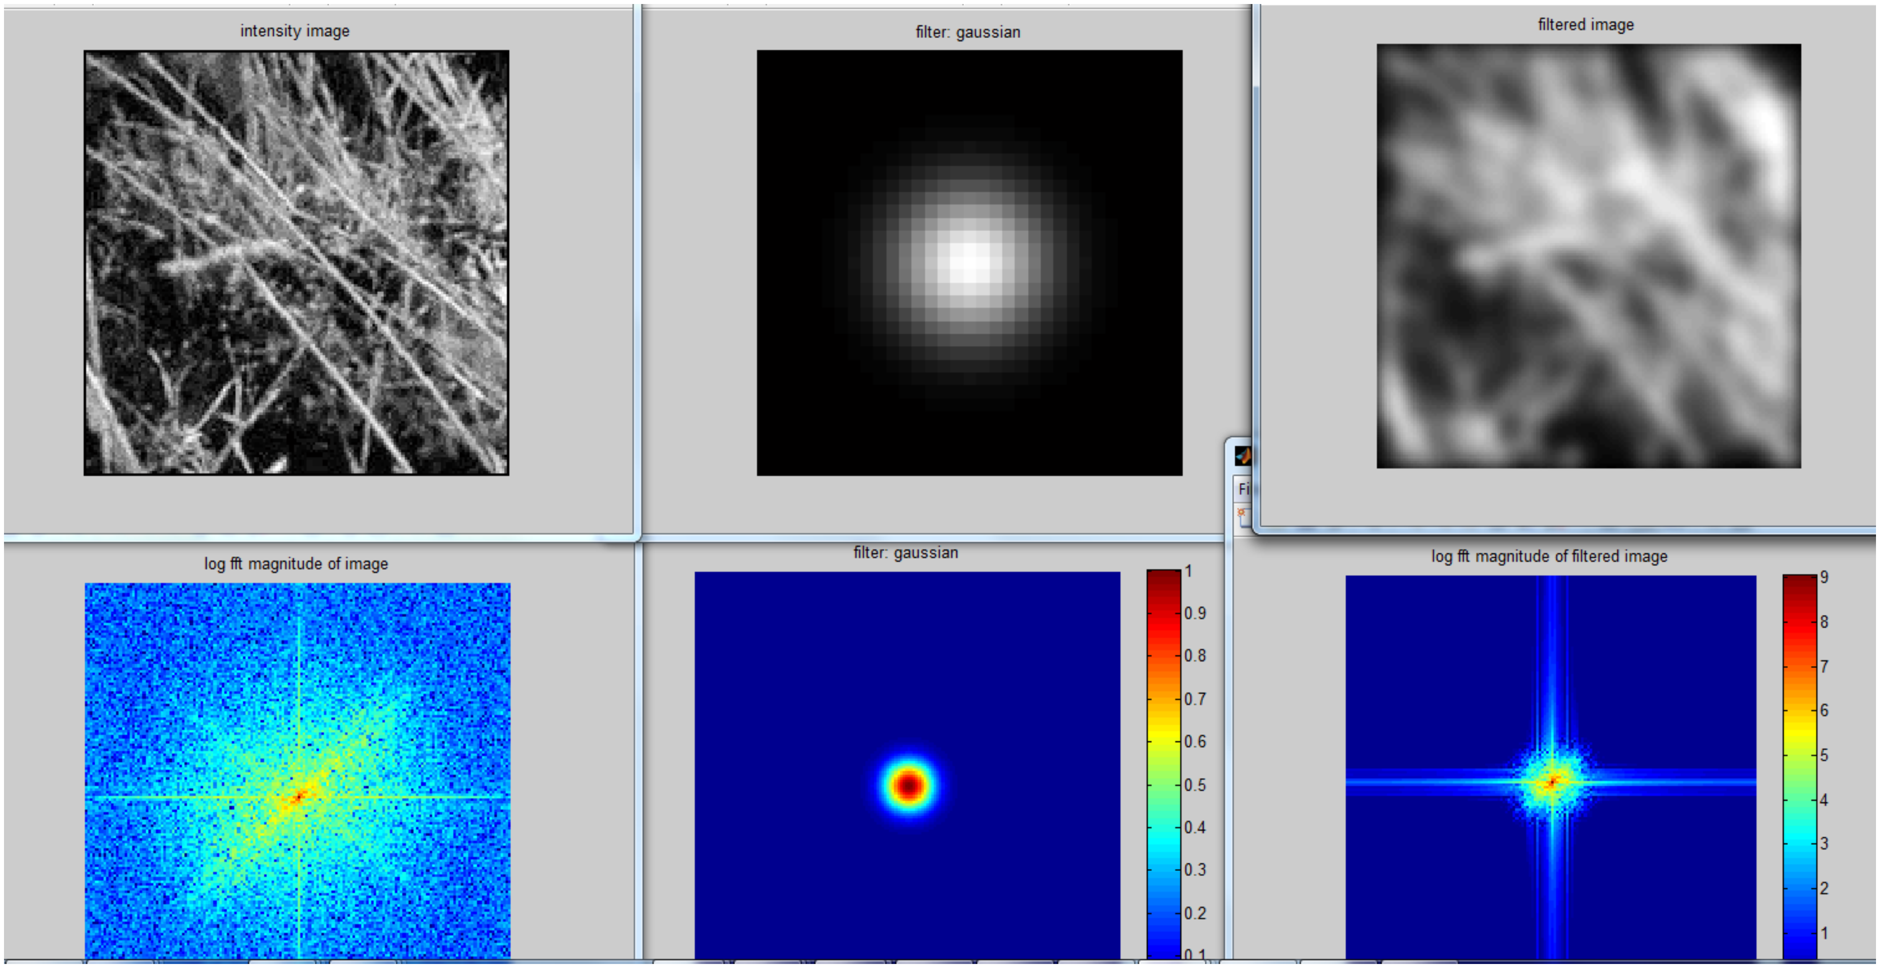

Hays

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.

Fourier, Joseph (1768-1830)

French mathematician who discovered that any periodic motion can be written as a superposition of sinusoidal and cosinusoidal vibrations. He developed a mathematical theory of heat in *Théorie Analytique de la Chaleur (Analytic Theory of Heat)*, (1822), discussing it in terms of differential equations.

Fourier was a friend and advisor of Napoleon. Fourier believed that his health would be improved by wrapping himself up in blankets, and in this state he tripped down the stairs in his house and killed himself. The paper of Galois which he had taken home to read shortly before his death was never recovered.

SEE ALSO: Galois

Additional biographies: MacTutor (St. Andrews), Bonn

How would math have changed if the onesie had been invented?!?!

: (

A few questions…

If we have infinite frequencies, why does the image end?

- Sampling theory. Frequencies higher than Nyquist frequency end up falling on an existing sample.
  - i.e., they are ‘aliases’ for existing samples!
- Nyquist frequency is half the sampling frequency.

- (Nyquist frequency is not Nyquist-Shannon rate, which is sampling required to reconstruct alias-free signal. Both are derived from same theory.)

A few questions

Why is frequency decomposition centered in middle, and duplicated and rotated?

- From Euler:
  - $\cos(x) + i\sin(x) = e^{ix}$
  - $\cos(\omega t) = \frac{1}{2}\left(e^{-i\omega t} + e^{+i\omega t}\right)$
- Coefficients for negative frequencies (i.e., ‘backwards traveling’ waves)

- FFT of a real signal is conjugate symmetric
  - i.e., f(-x) = f*(x)

A few questions

How is the Fourier decomposition computed?

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

Notes in hidden slides.

Plus: http://research.stowers-institute.org/efg/Report/FourierAnalysis.pdf

Applications of Fourier analysis

- Fast filtering with large kernels

- Fourier Optics
  - Fraunhofer diffraction is Fourier transform of slit in the 'far field'.

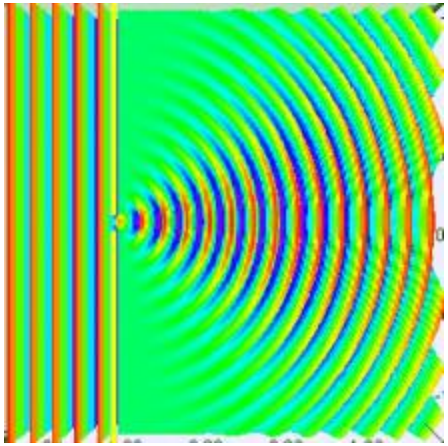

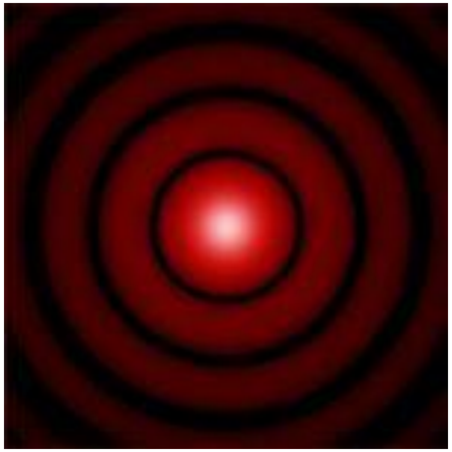

Circular aperture = Airy disc diffraction pattern

  - Light spectrometry for astronomy

Wikipedia for graphics

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)

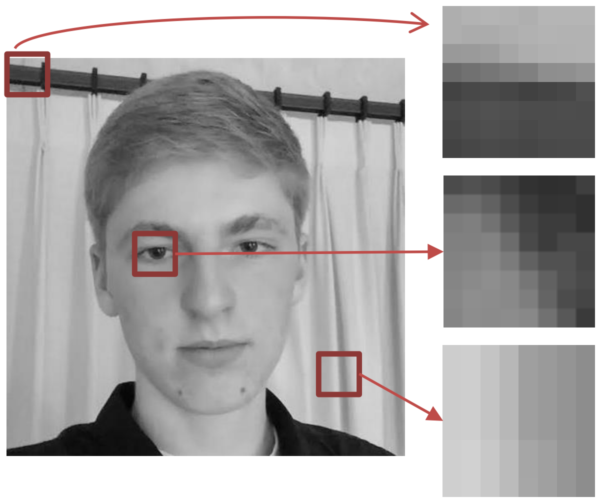

The first coefficient B(0,0) is the DC component, the average intensity

The top-left coeffs represent low frequencies, the bottom right represent high frequencies

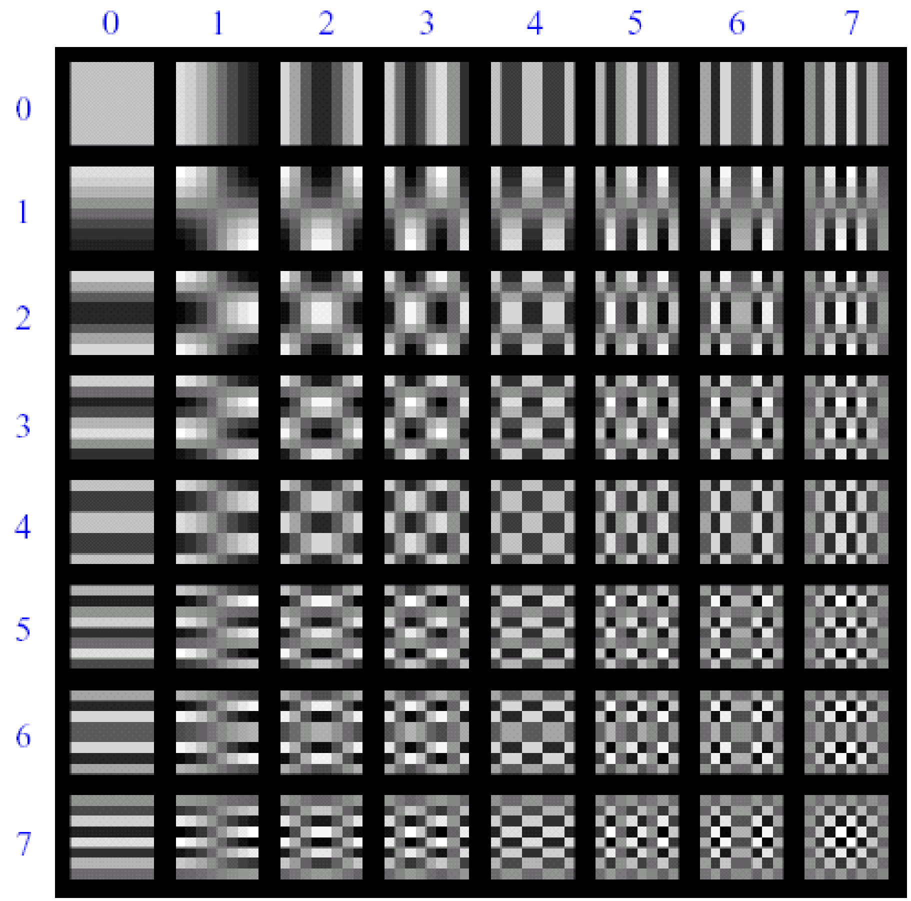

Block-based Discrete Cosine Transform (DCT)

Slides: Efros

Image compression using DCT

- Compute DCT filter responses in each 8x8 block

Filter responses

$$G = \begin{array}{c} \xrightarrow{u} \\ \begin{bmatrix} -415.38 & -30.19 & -61.20 & 27.24 & 56.13 & -20.10 & -2.39 & 0.46 \\ 4.47 & -21.86 & -60.76 & 10.25 & 13.15 & -7.09 & -8.54 & 4.88 \\ -46.83 & 7.37 & 77.13 & -24.56 & -28.91 & 9.93 & 5.42 & -5.65 \\ -48.53 & 12.07 & 34.10 & -14.76 & -10.24 & 6.30 & 1.83 & 1.95 \\ 12.12 & -6.55 & -13.20 & -3.95 & -1.88 & 1.75 & -2.79 & 3.14 \\ -7.73 & 2.91 & 2.38 & -5.94 & -2.38 & 0.94 & 4.30 & 1.85 \\ -1.03 & 0.18 & 0.42 & -2.42 & -0.88 & -3.02 & 4.12 & -0.66 \\ -0.17 & 0.14 & -1.07 & -4.19 & -1.17 & -0.10 & 0.50 & 1.68 \end{bmatrix} \end{array} \Big\downarrow v$$

- 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)

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

Quantized values

$$B = \begin{bmatrix} -26 & -3 & -6 & 2 & 2 & -1 & 0 & 0 \\ 0 & -2 & -4 & 1 & 1 & 0 & 0 & 0 \\ -3 & 1 & 5 & -1 & -1 & 0 & 0 & 0 \\ -3 & 1 & 2 & -1 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix}$$

JPEG Encoding

- Entropy coding (Huffman-variant)

Quantized values

$$B = \begin{bmatrix} -26 & -3 & -6 & 2 & 2 & -1 & 0 & 0 \\ 0 & -2 & -4 & 1 & 1 & 0 & 0 & 0 \\ -3 & 1 & 5 & -1 & -1 & 0 & 0 & 0 \\ -3 & 1 & 2 & -1 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix}$$

Linearize *B* like this.

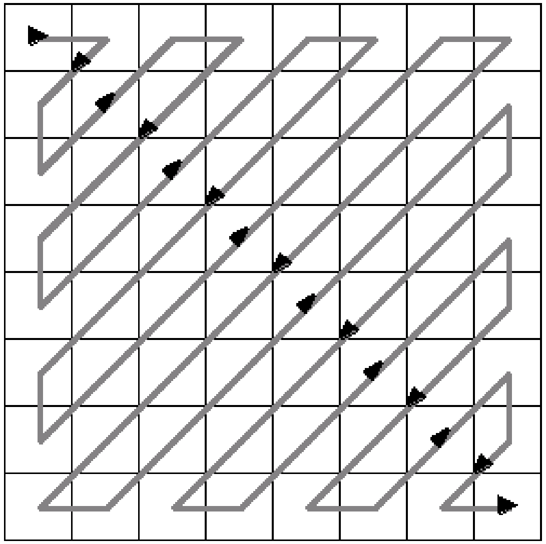

Helps compression:

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

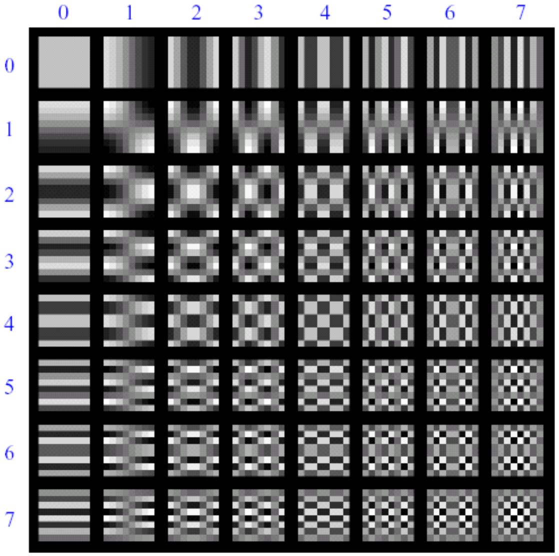

Color spaces: YCbCr

Fast to compute, good for compression, used by TV

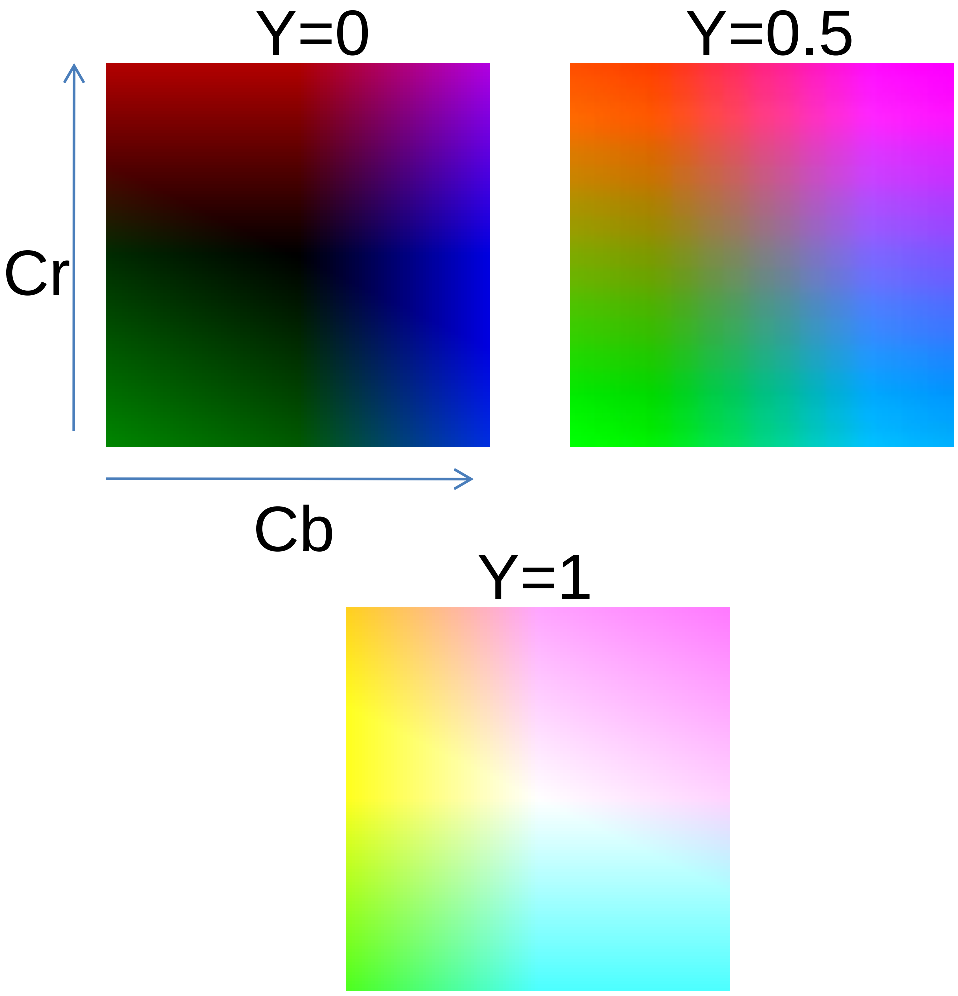

Y
(Cb=0.5,Cr=0.5)

Cb
(Y=0.5,Cr=0.5)

Cr
(Y=0.5,Cb=05)

James Hays

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)

http://en.wikipedia.org/wiki/YCbCr
http://en.wikipedia.org/wiki/JPEG