1950

Future Vision

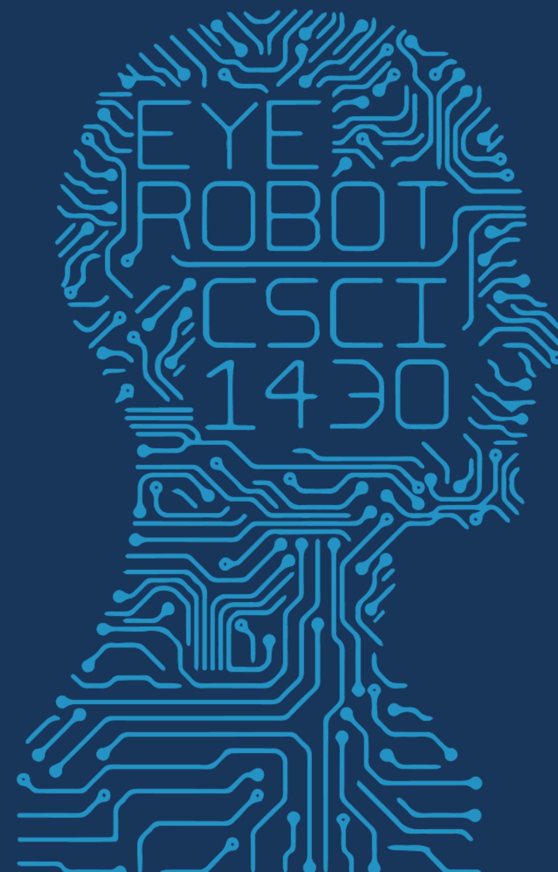

2017 MWF 1pm 368

Computer Vision

Recap: Fourier transform

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

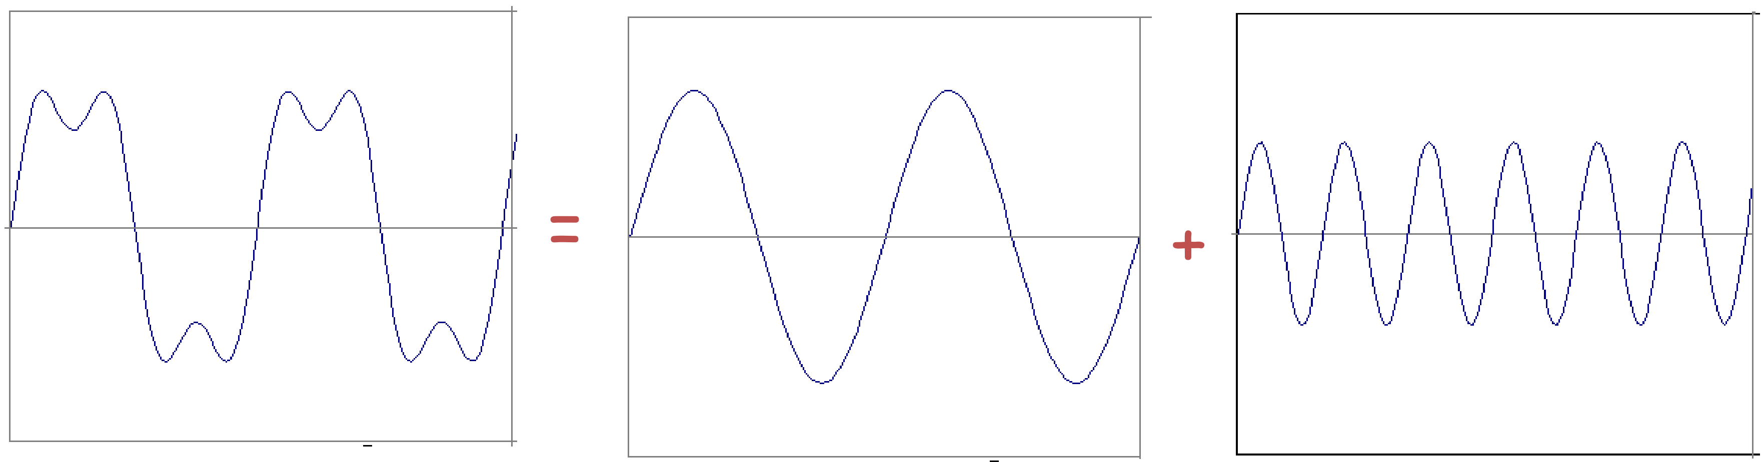

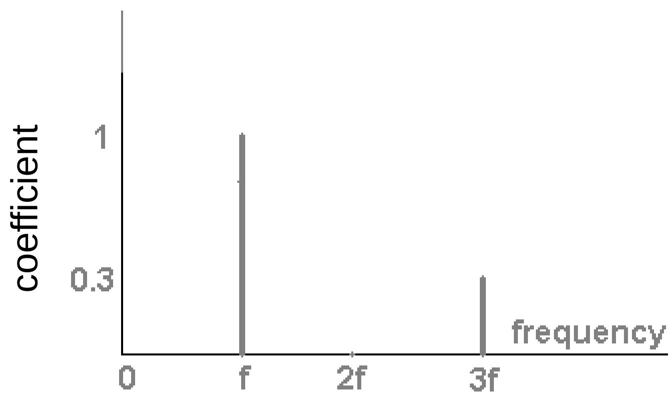

Efros

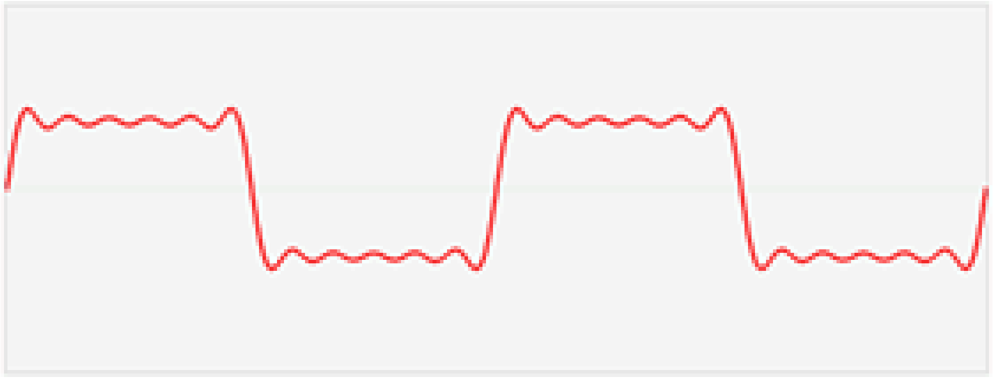

Wikipedia – Fourier transform

Sine/cosine and circle

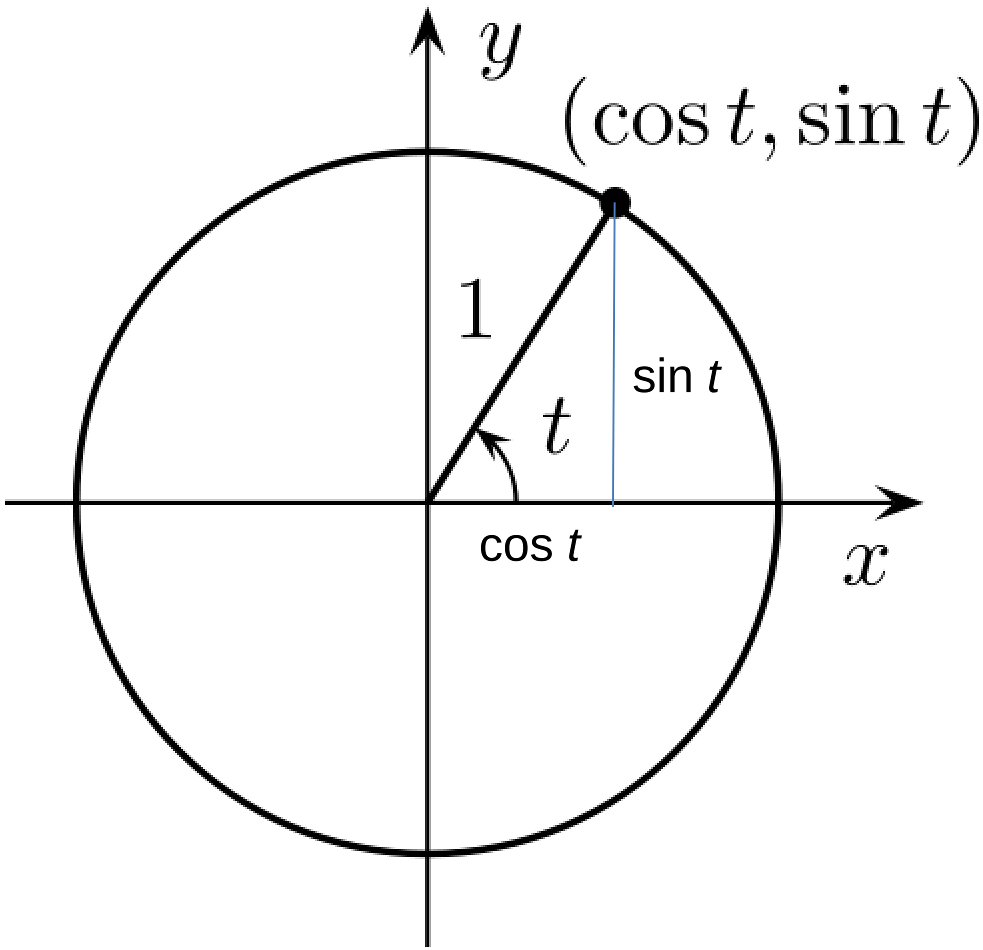

Wikipedia – Unit Circle

Square wave (approx.)

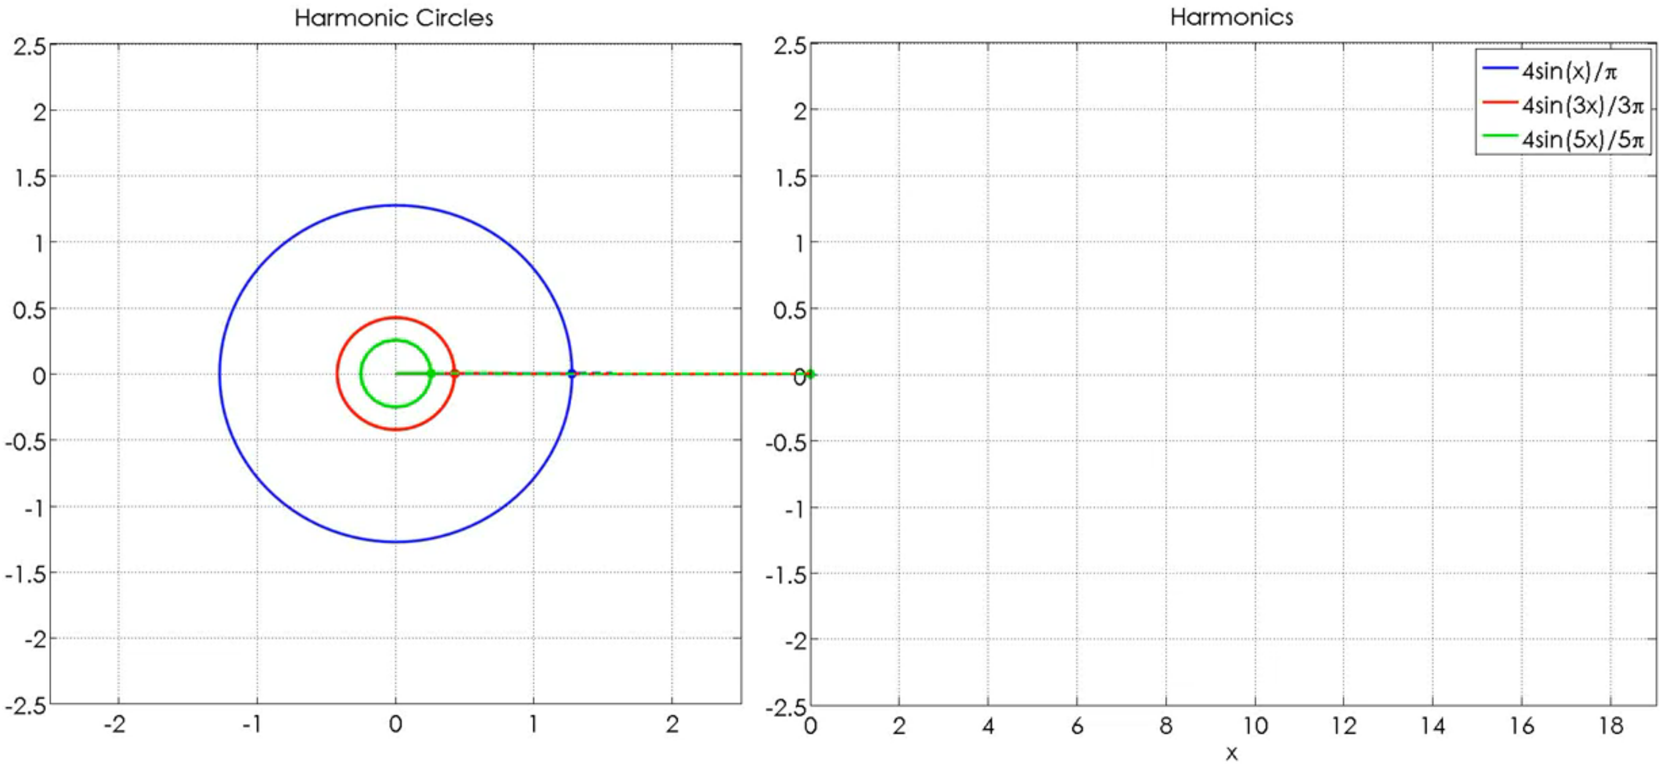

Sawtooth wave (approx.)

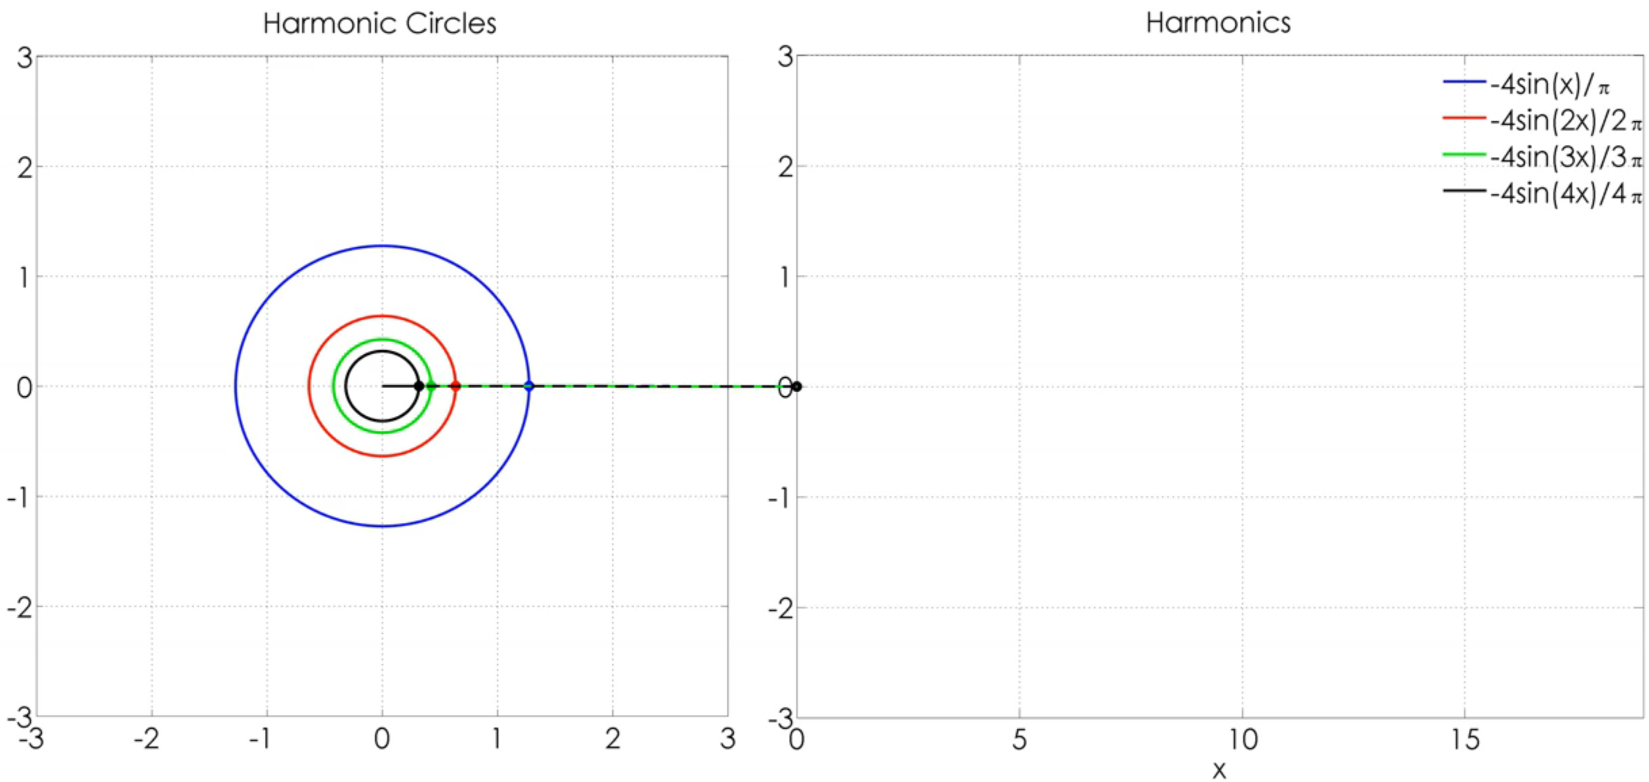

Mehmet E. Yavuz

Euler's formula

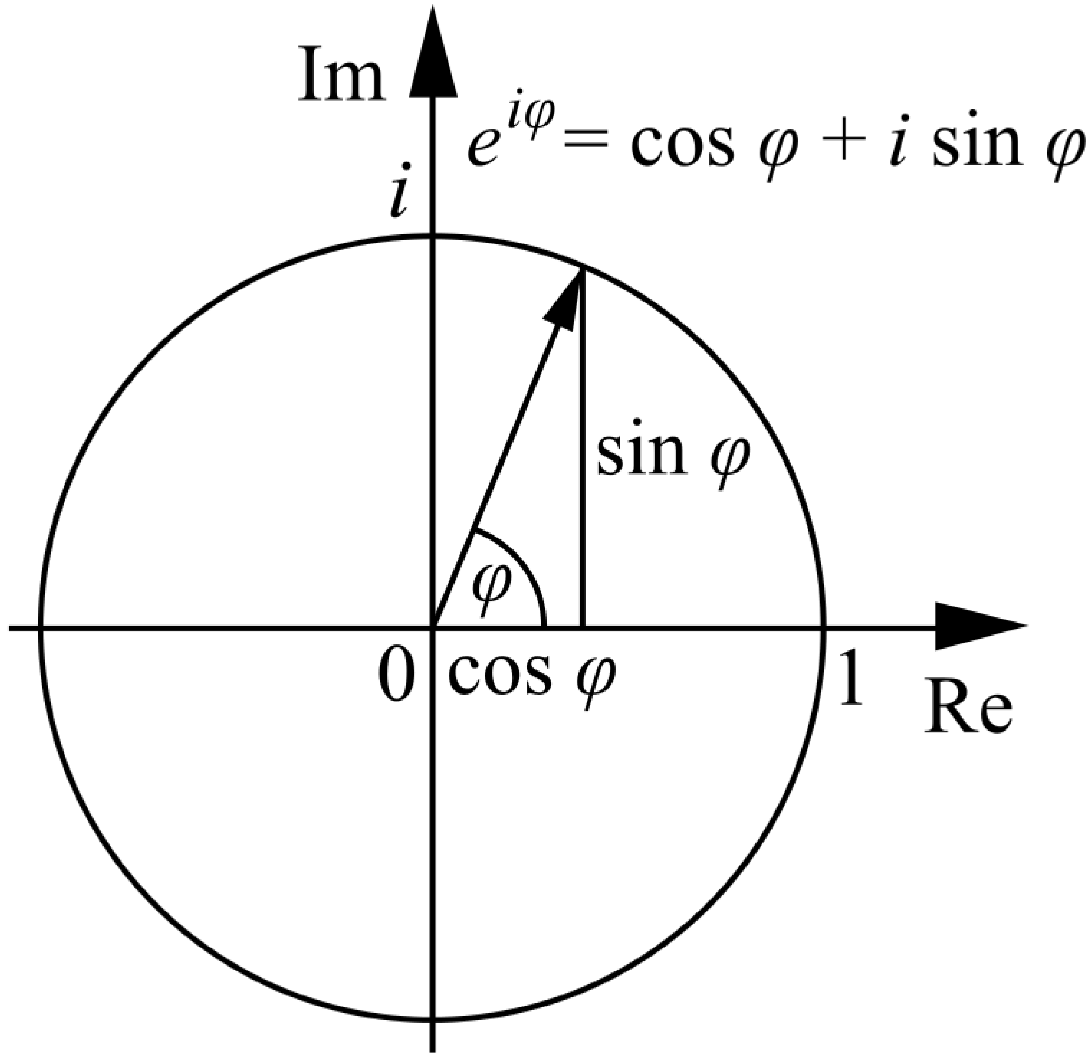

Wikipedia

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

Amplitude: $$A = \pm\sqrt{\mathrm{Re}(\omega)^2 + \mathrm{Im}(\omega)^2}$$

  - Phase encodes spatial information (indirectly)

Phase: $$\phi = \tan^{-1}\frac{\mathrm{Im}(\omega)}{\mathrm{Re}(\omega)}$$

Hays

Brian Pauw demo

- Live FFT2 demo
- I hacked it a bit
- http://www.lookingatnothing.com/index.php/archives/991

Amplitude / Phase

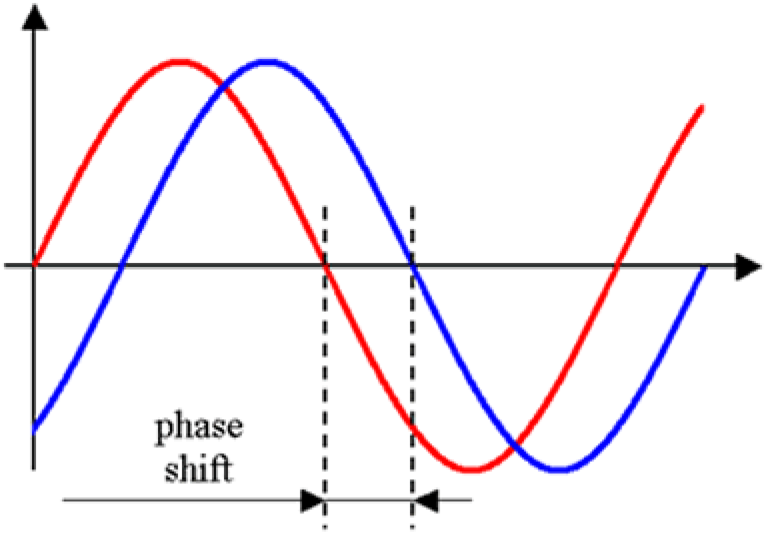

- Amplitude tells you “how much”
- Phase tells you “where”

- Translate the image?
  - Amplitude unchanged
  - Adds a constant to the phase.

Hays

Filtering

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

Gaussian

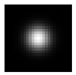

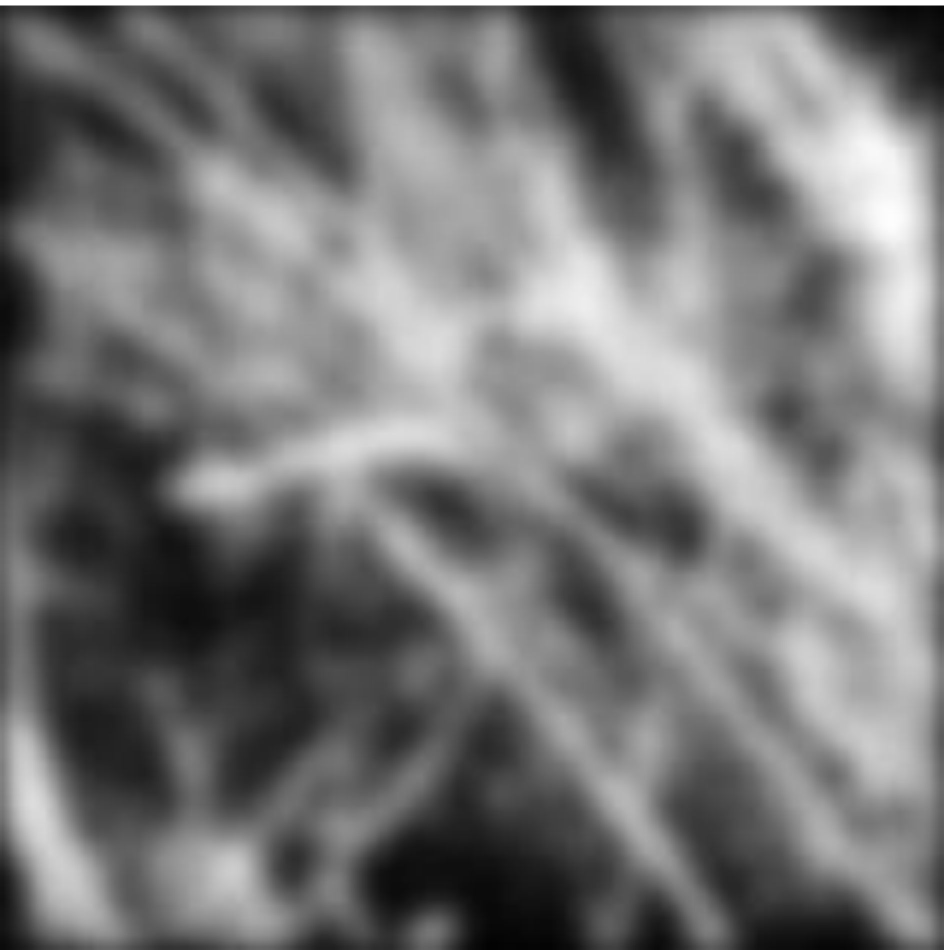

Box filter

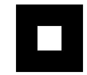

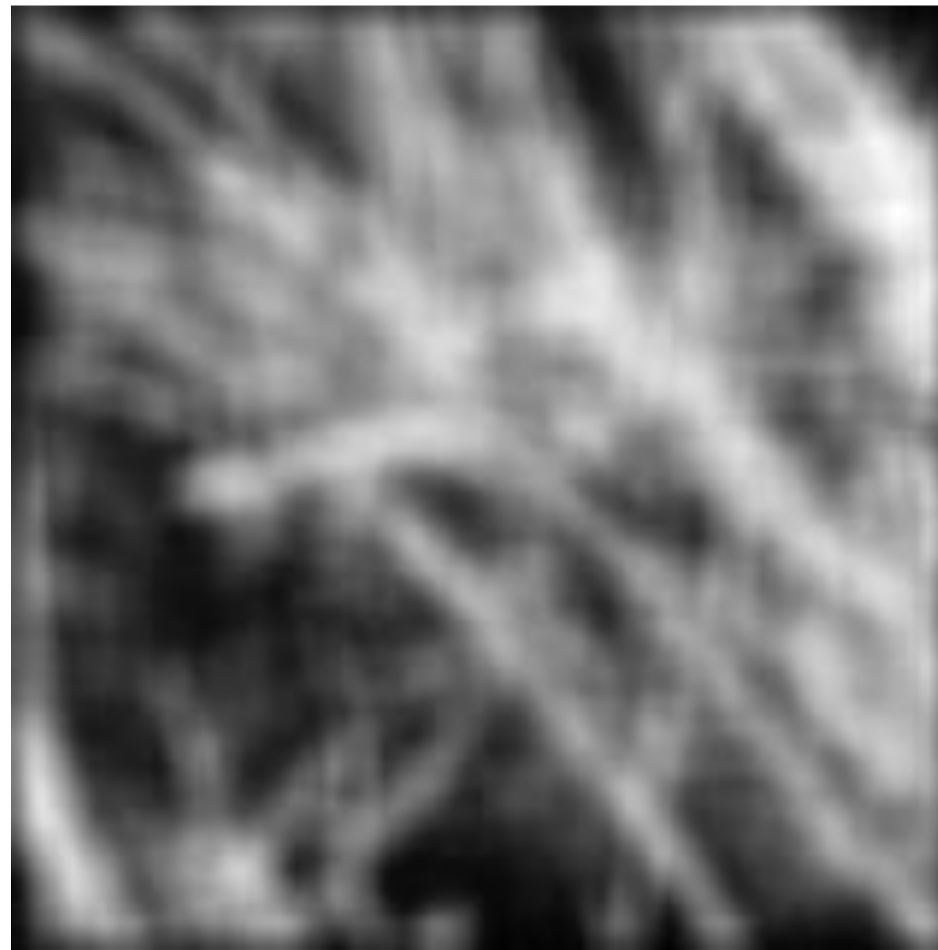

Why do we have those lines in the image?

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

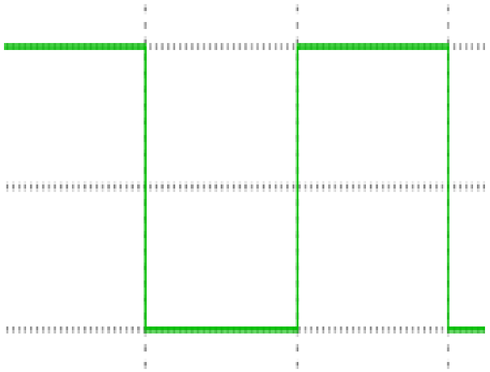

$$= A\sum_{k=1}^{\infty}\frac{1}{k}\sin(2\pi kt)$$

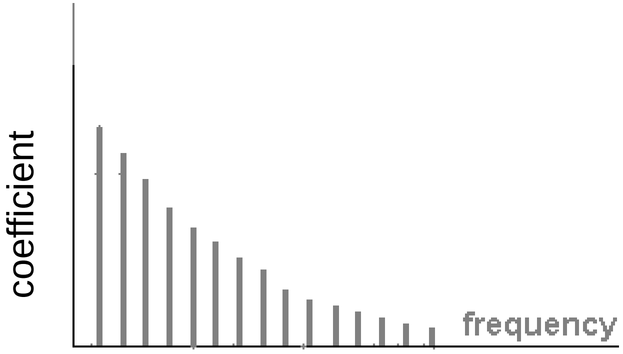

Efros

Hays

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

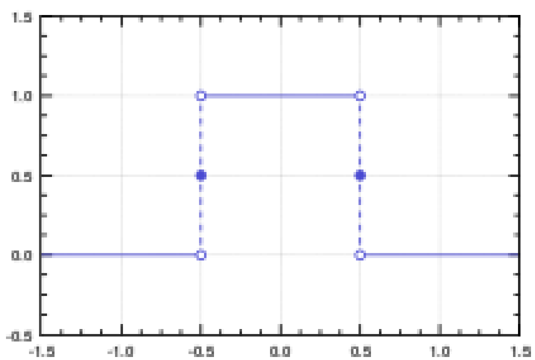

Sinc filter

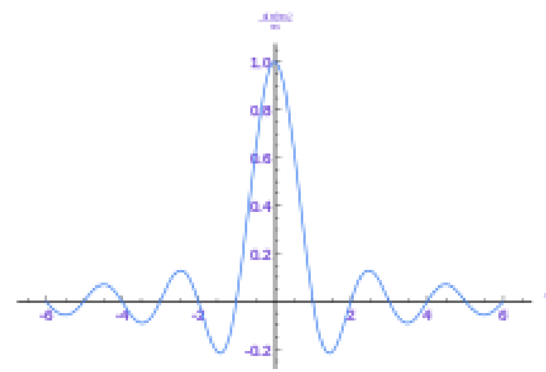

sinc(x) = sin(x) / x

Spatial Domain ⟺ Frequency Domain

Frequency Domain ⟺ Spatial Domain

Box filter (spatial)

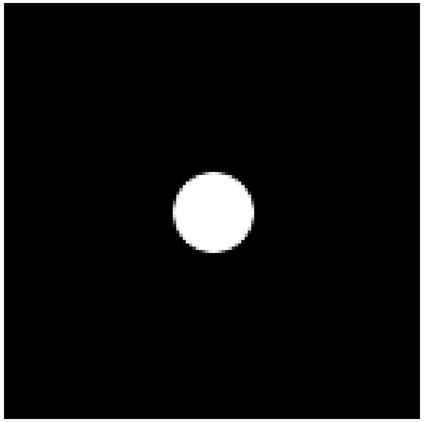

Frequency domain magnitude

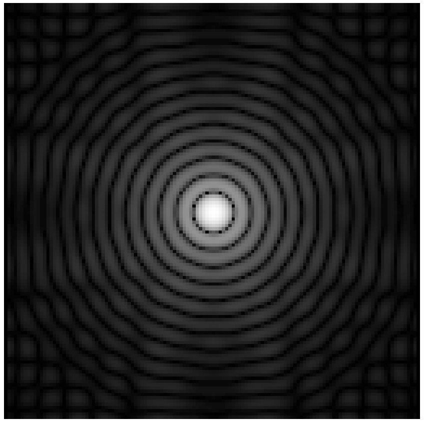

Gaussian filter duality

- Fourier transform of one Gaussian…

  …is *another Gaussian (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)

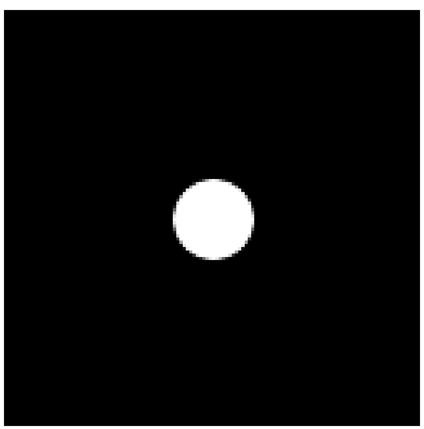

Frequency domain magnitude

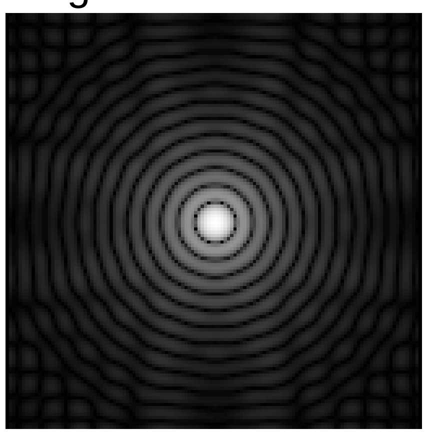

Gaussian filter (spatial)

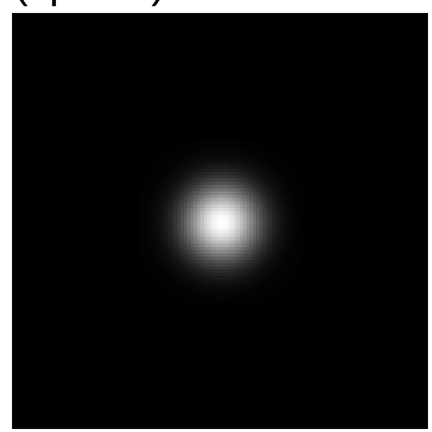

Frequency domain magnitude

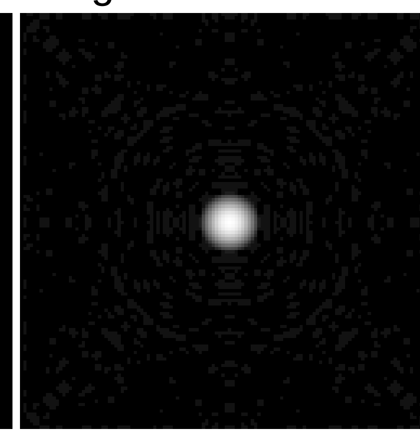

Filtering

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

Gaussian

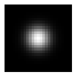

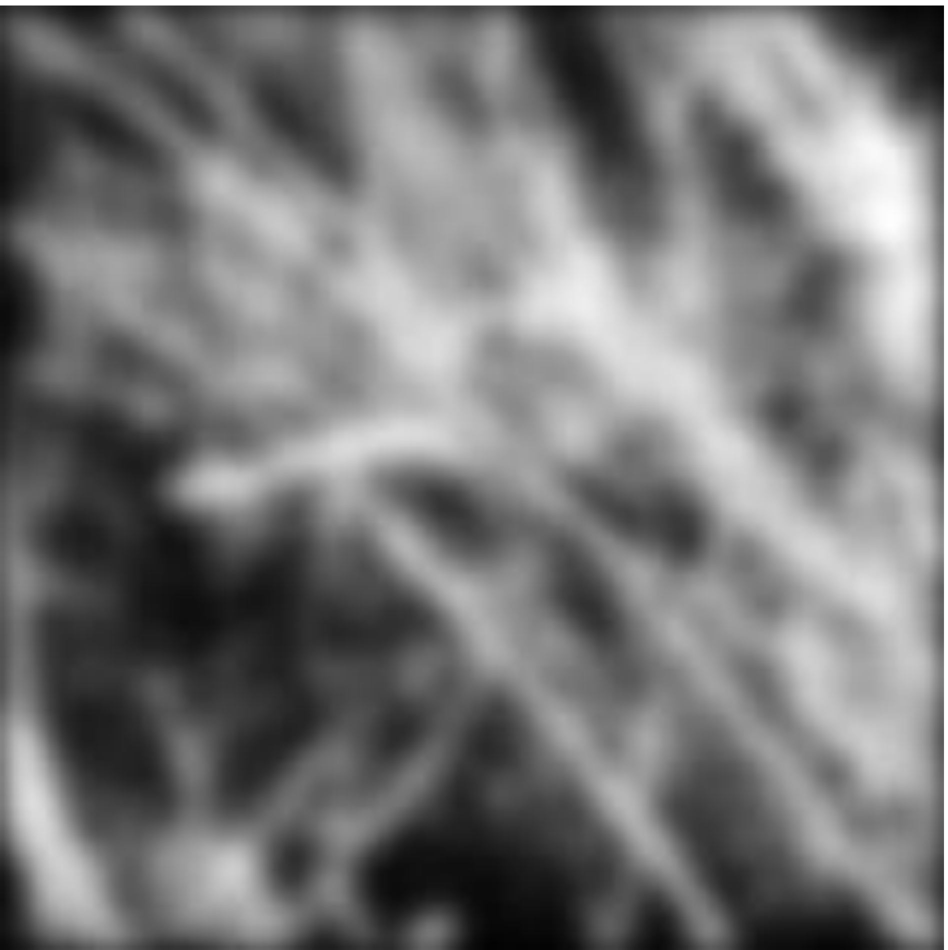

Box filter

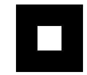

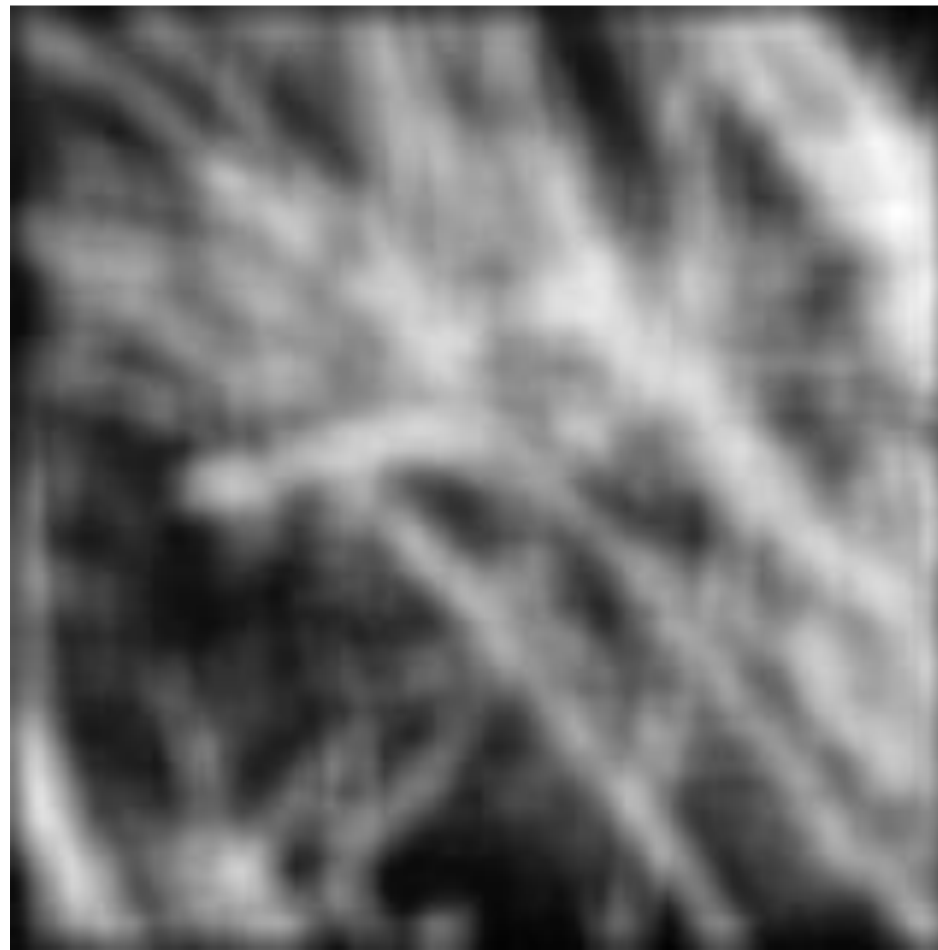

Properties of Fourier Transforms

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

- 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

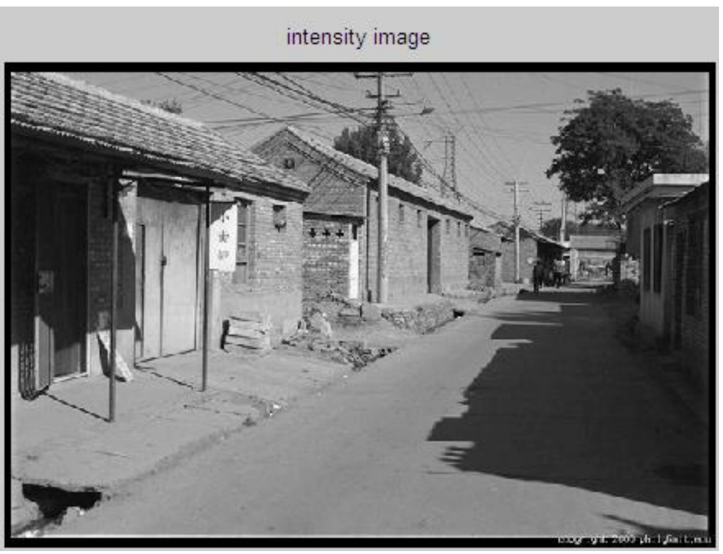

*

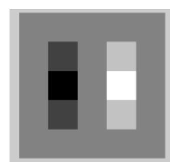

=

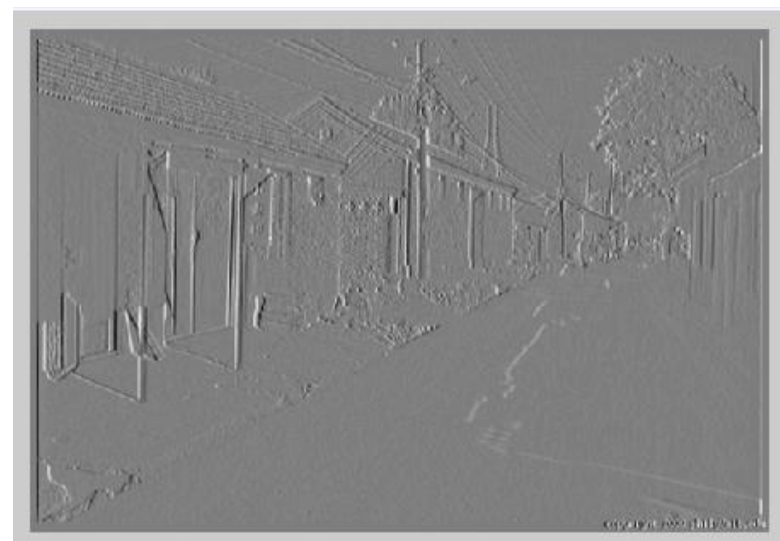

Hays

Filtering in frequency domain

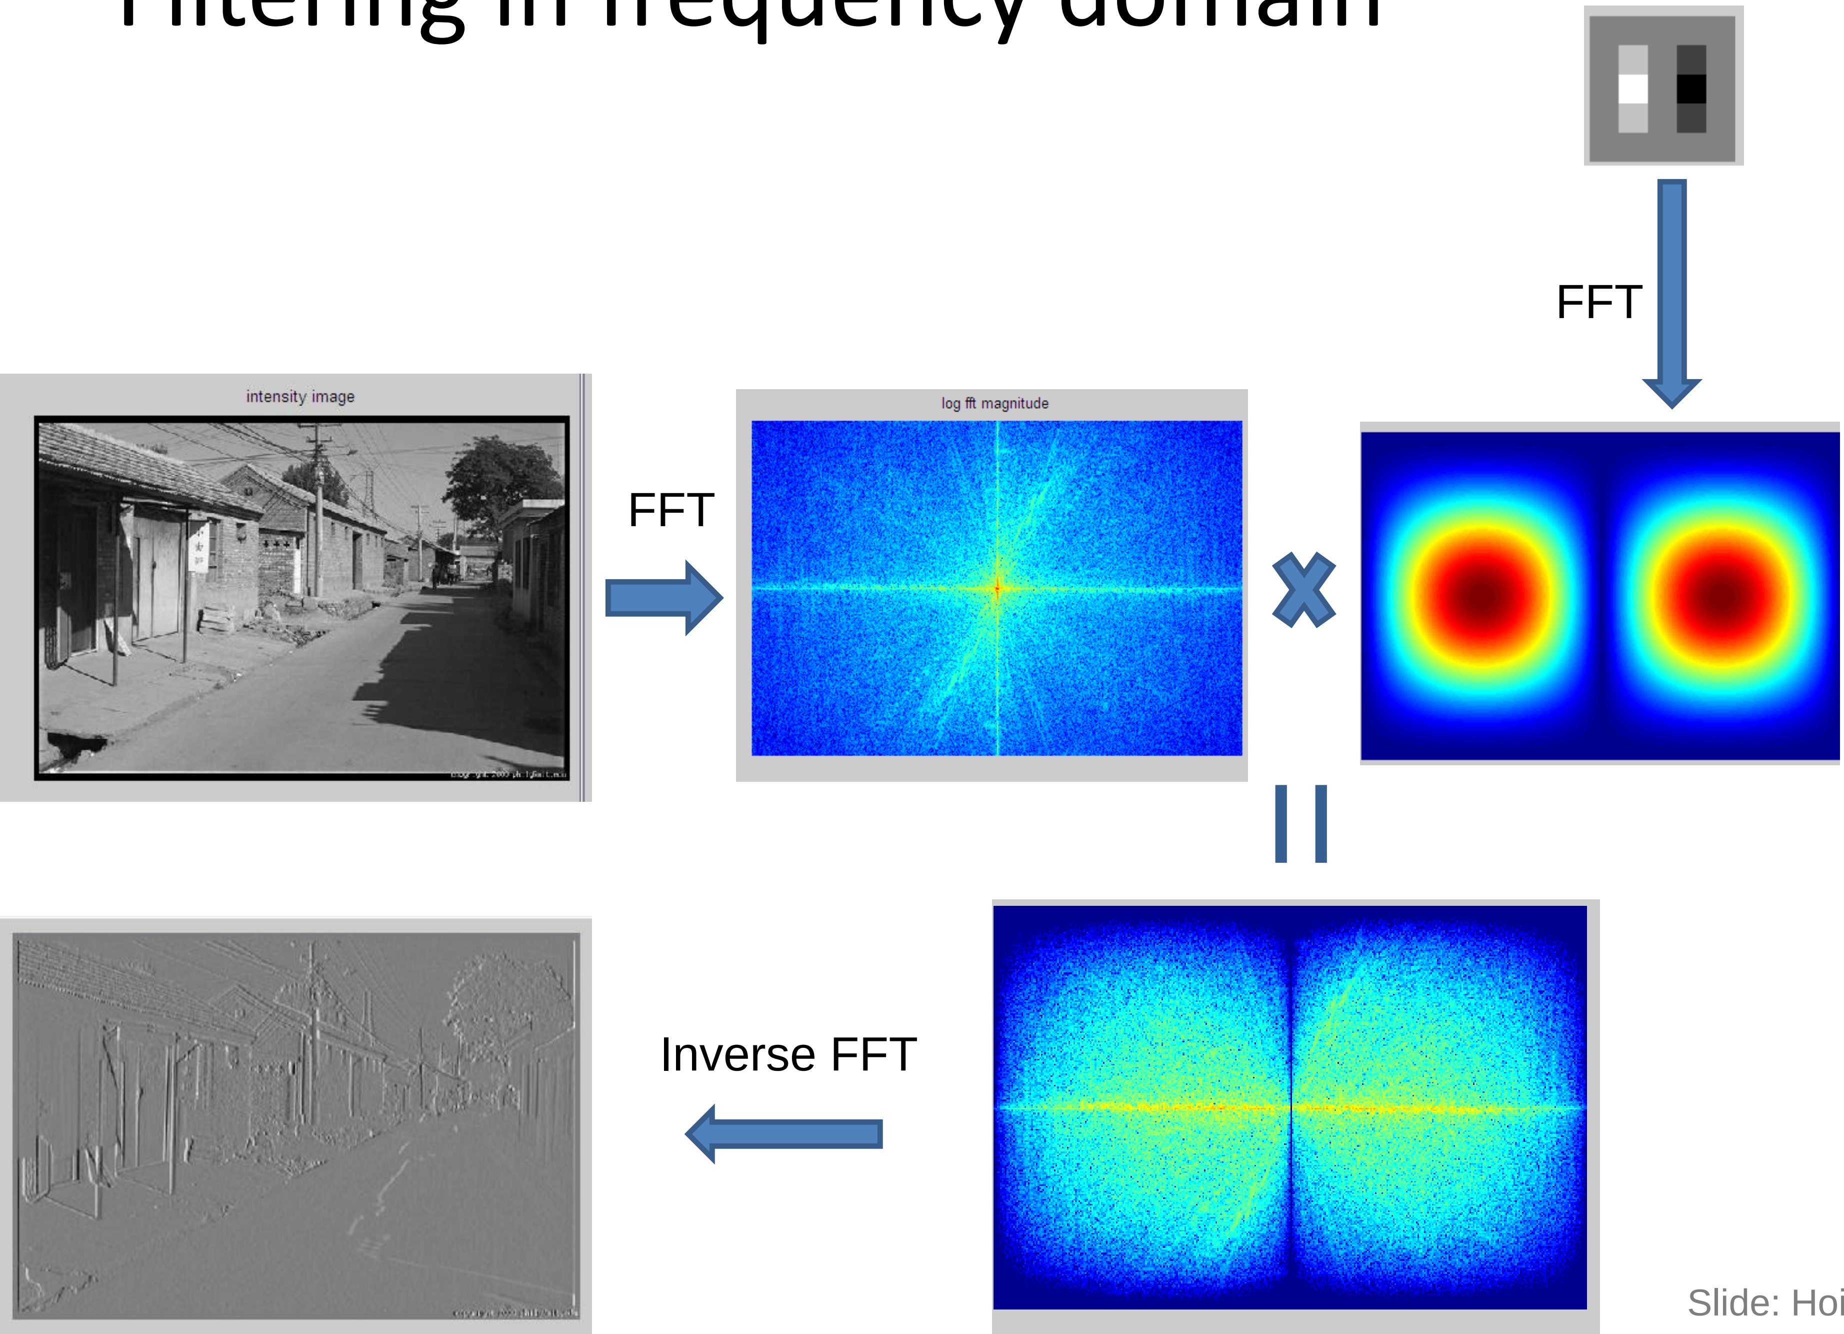

Slide: Hoiem

Think-Pair-Share

Match the spatial domain image to the Fourier magnitude image

1

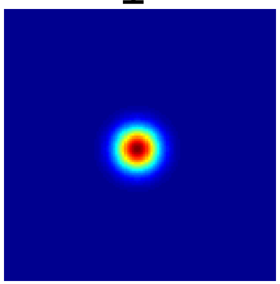

2

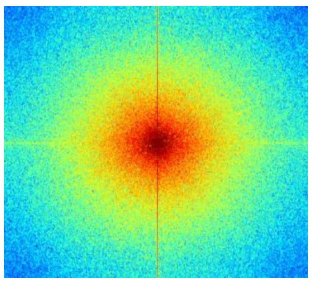

3

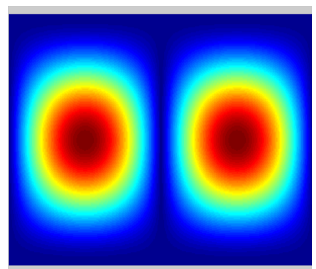

4

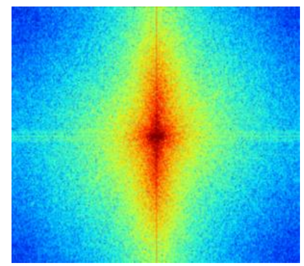

5

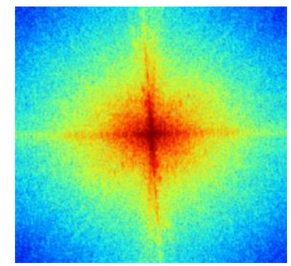

A

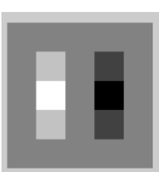

B

C

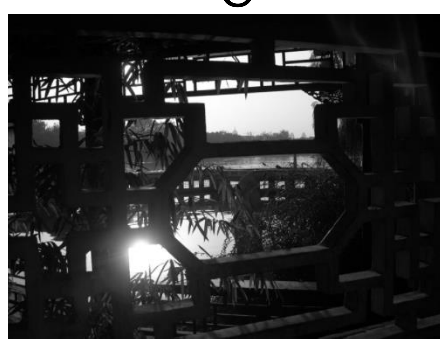

D

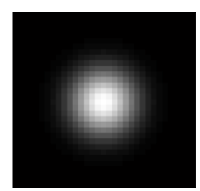

E

Hoiem

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?

Let’s experiment on Novak

Convolution

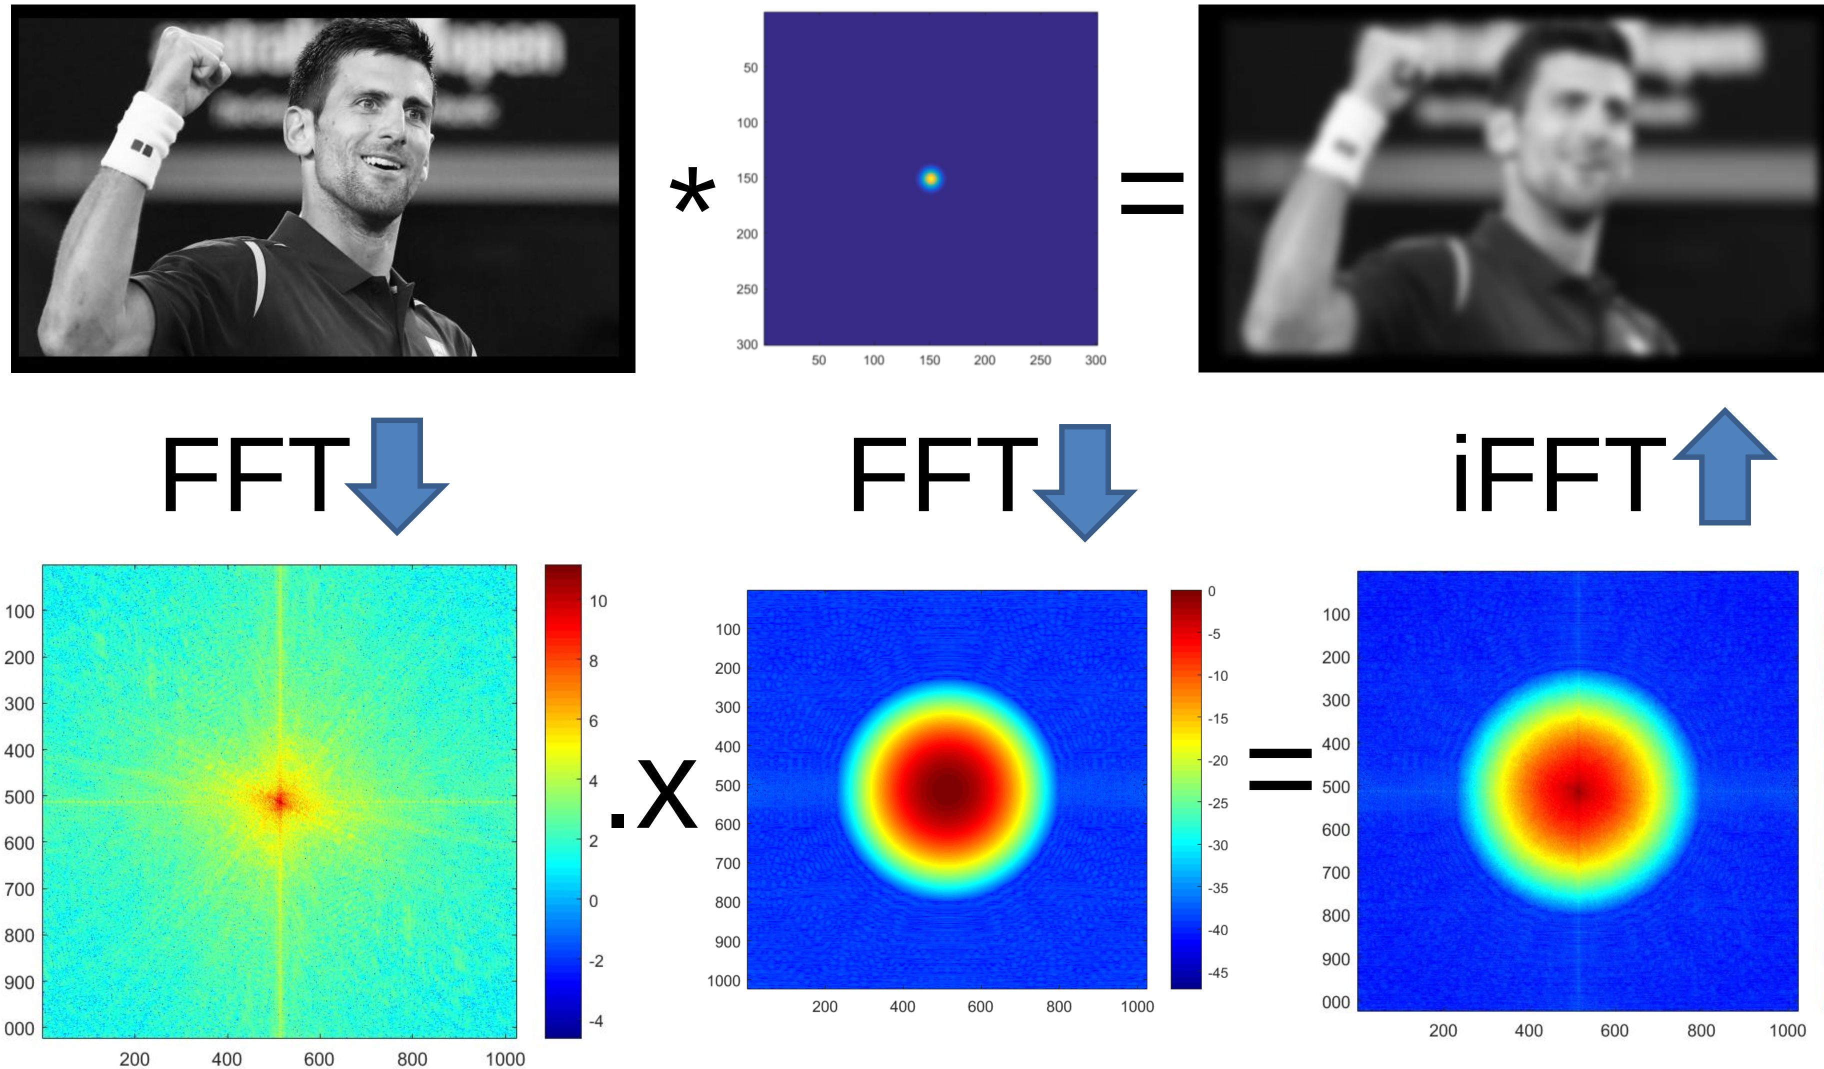

Deconvolution?

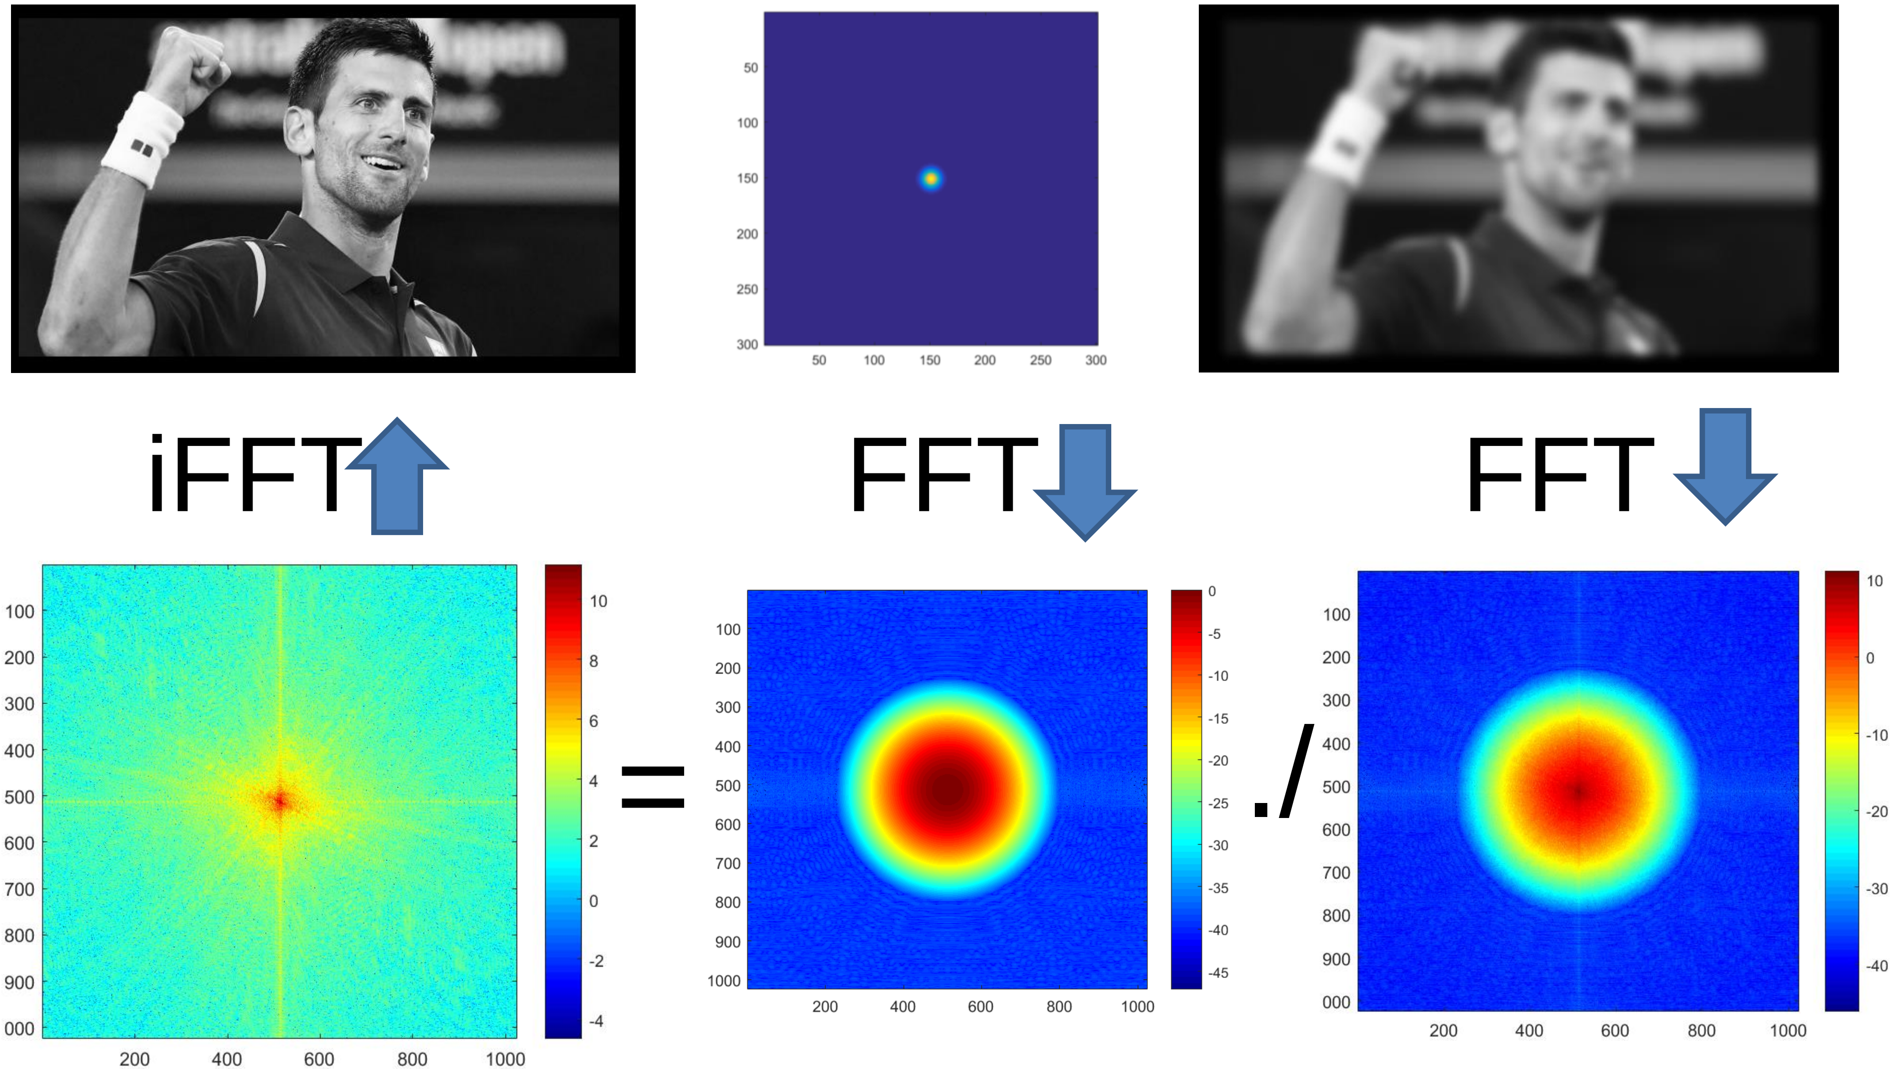

But under more realistic conditions

Random noise, .000001 magnitude

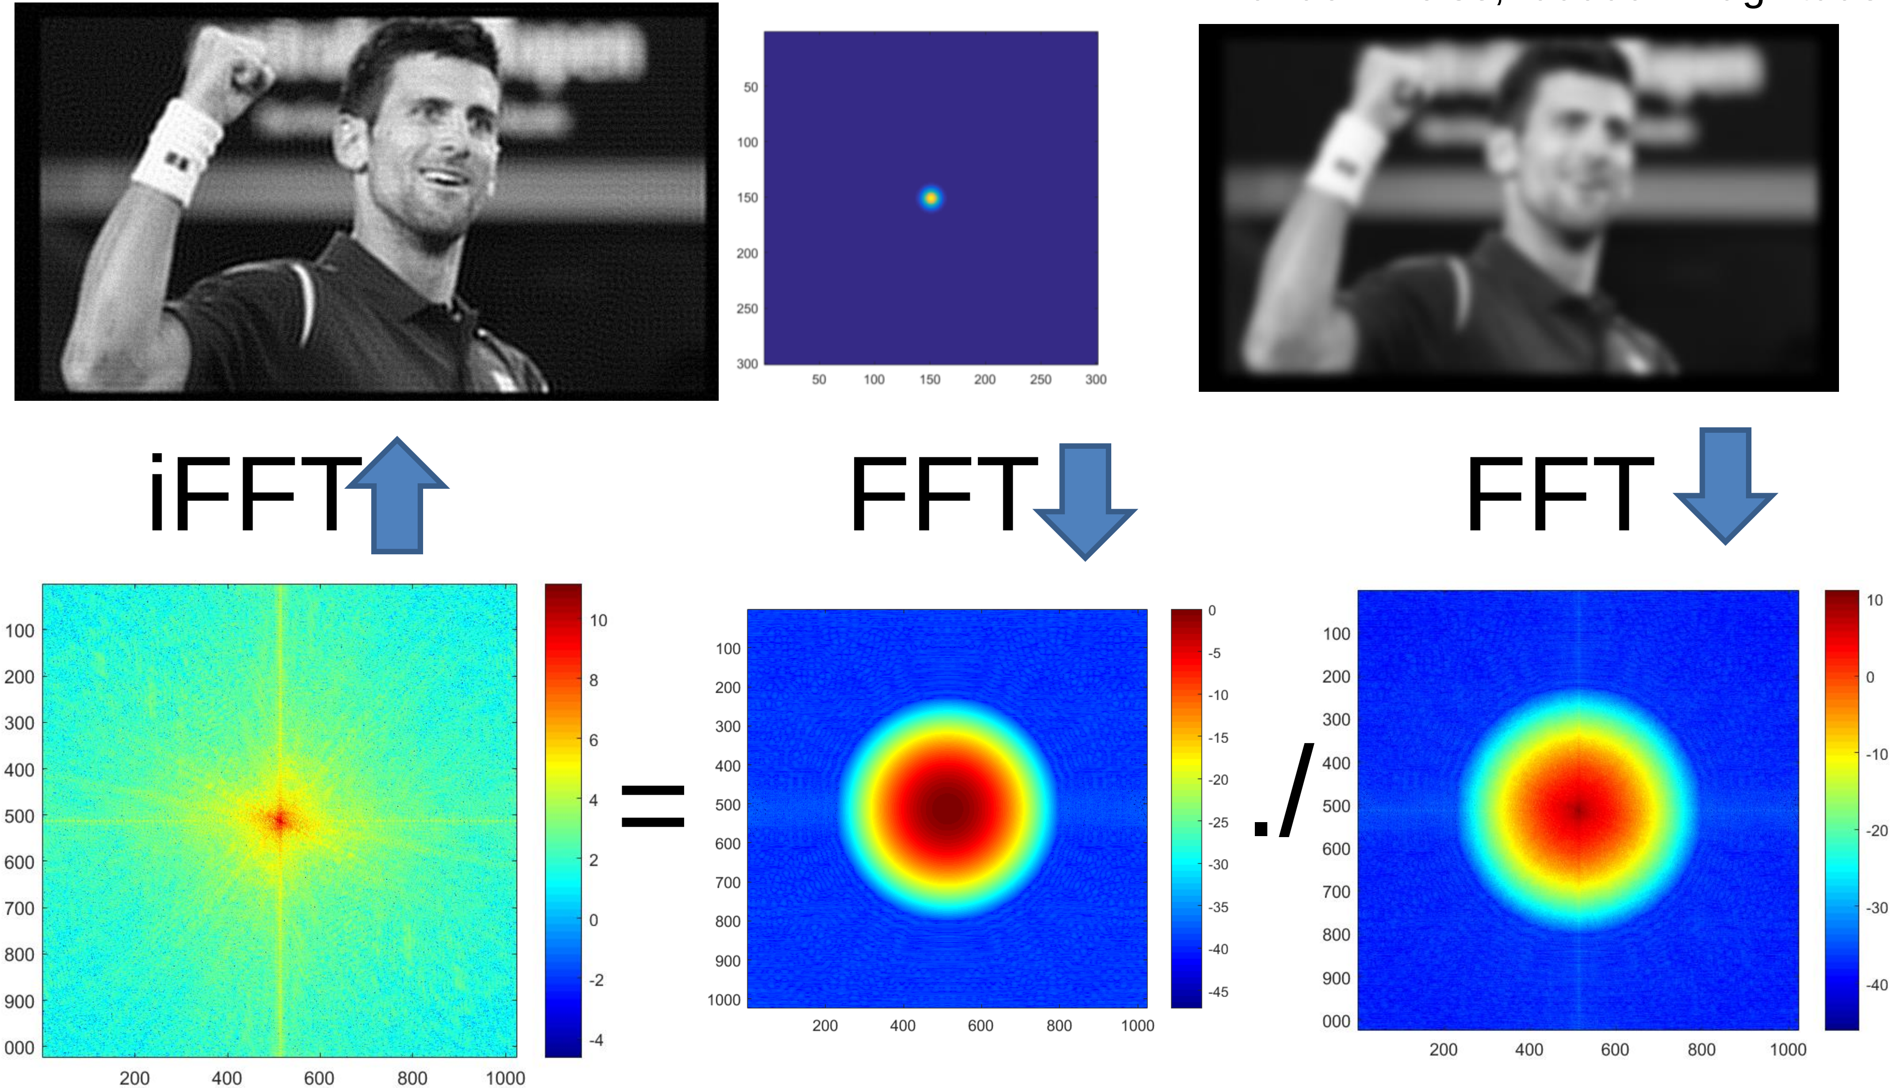

But under more realistic conditions

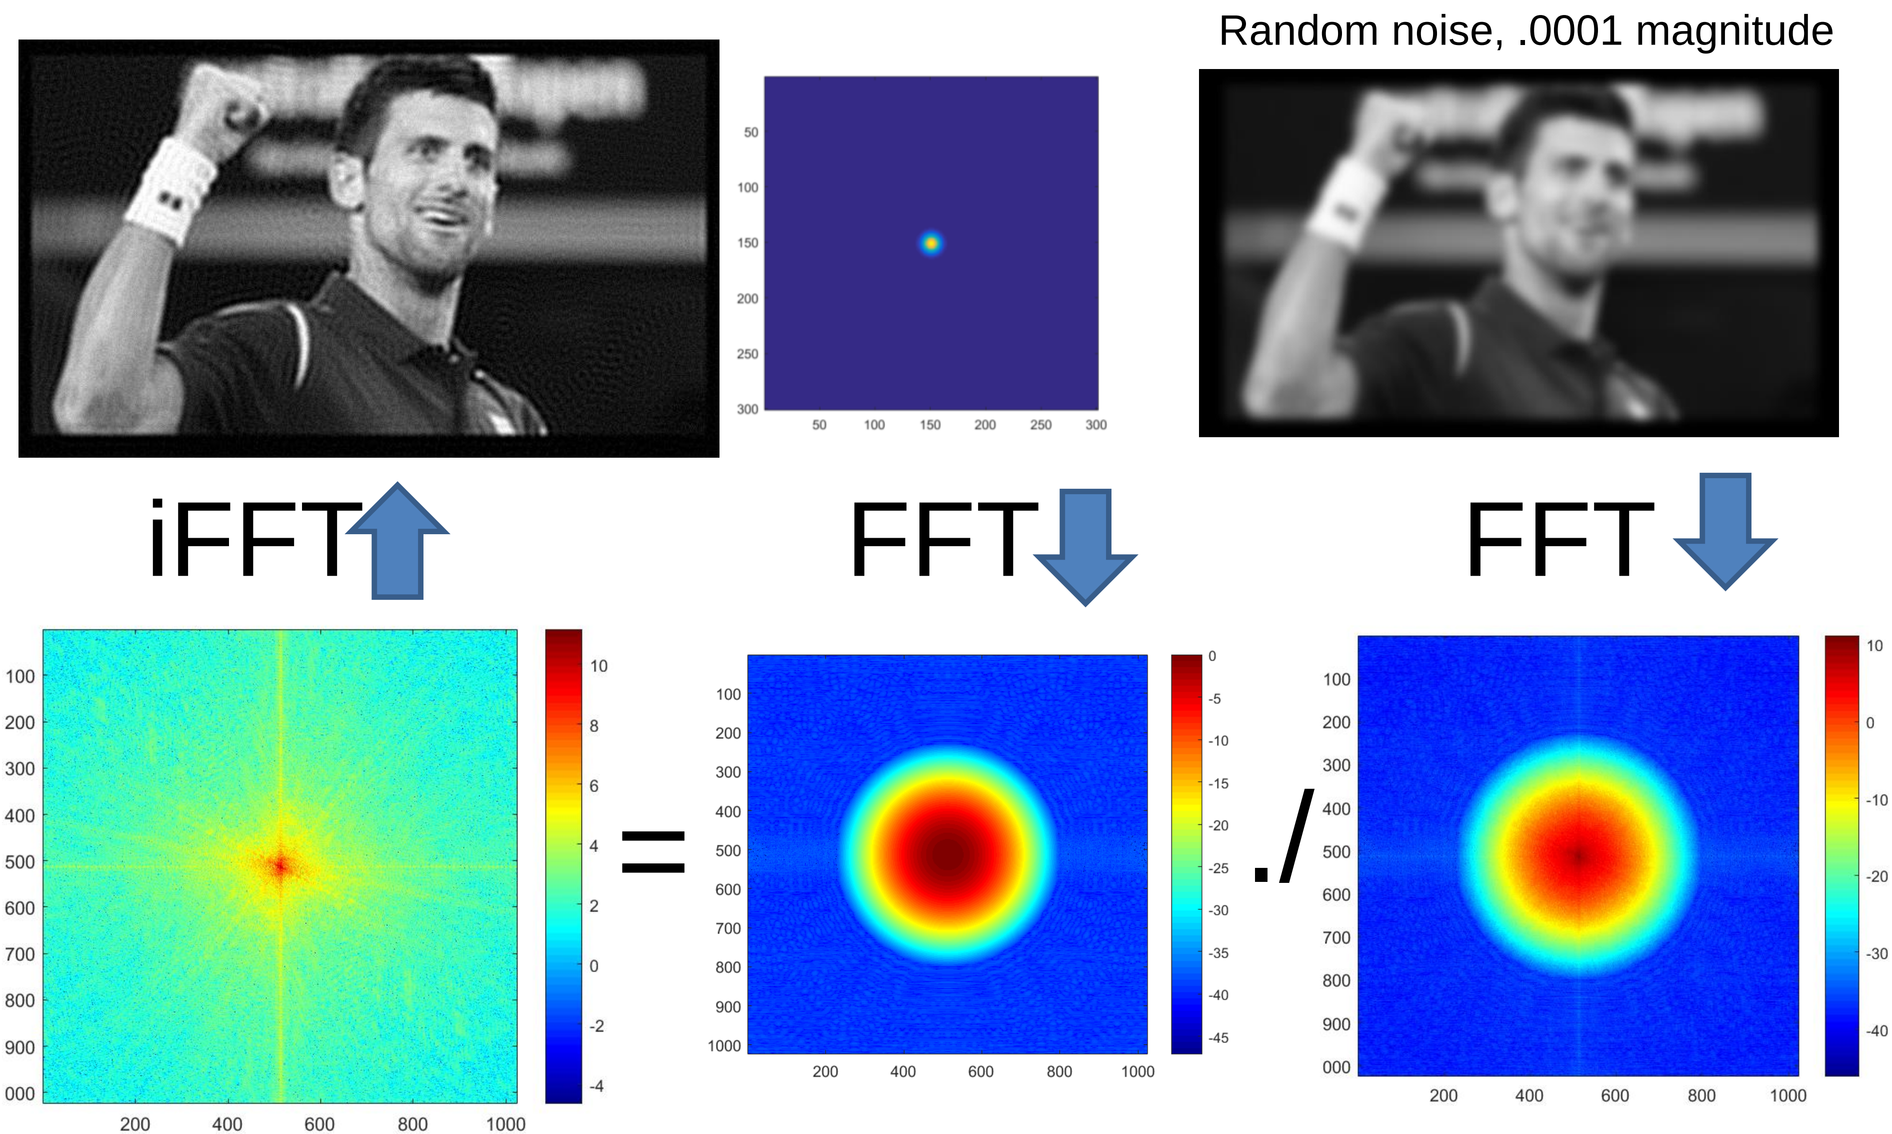

But under more realistic conditions

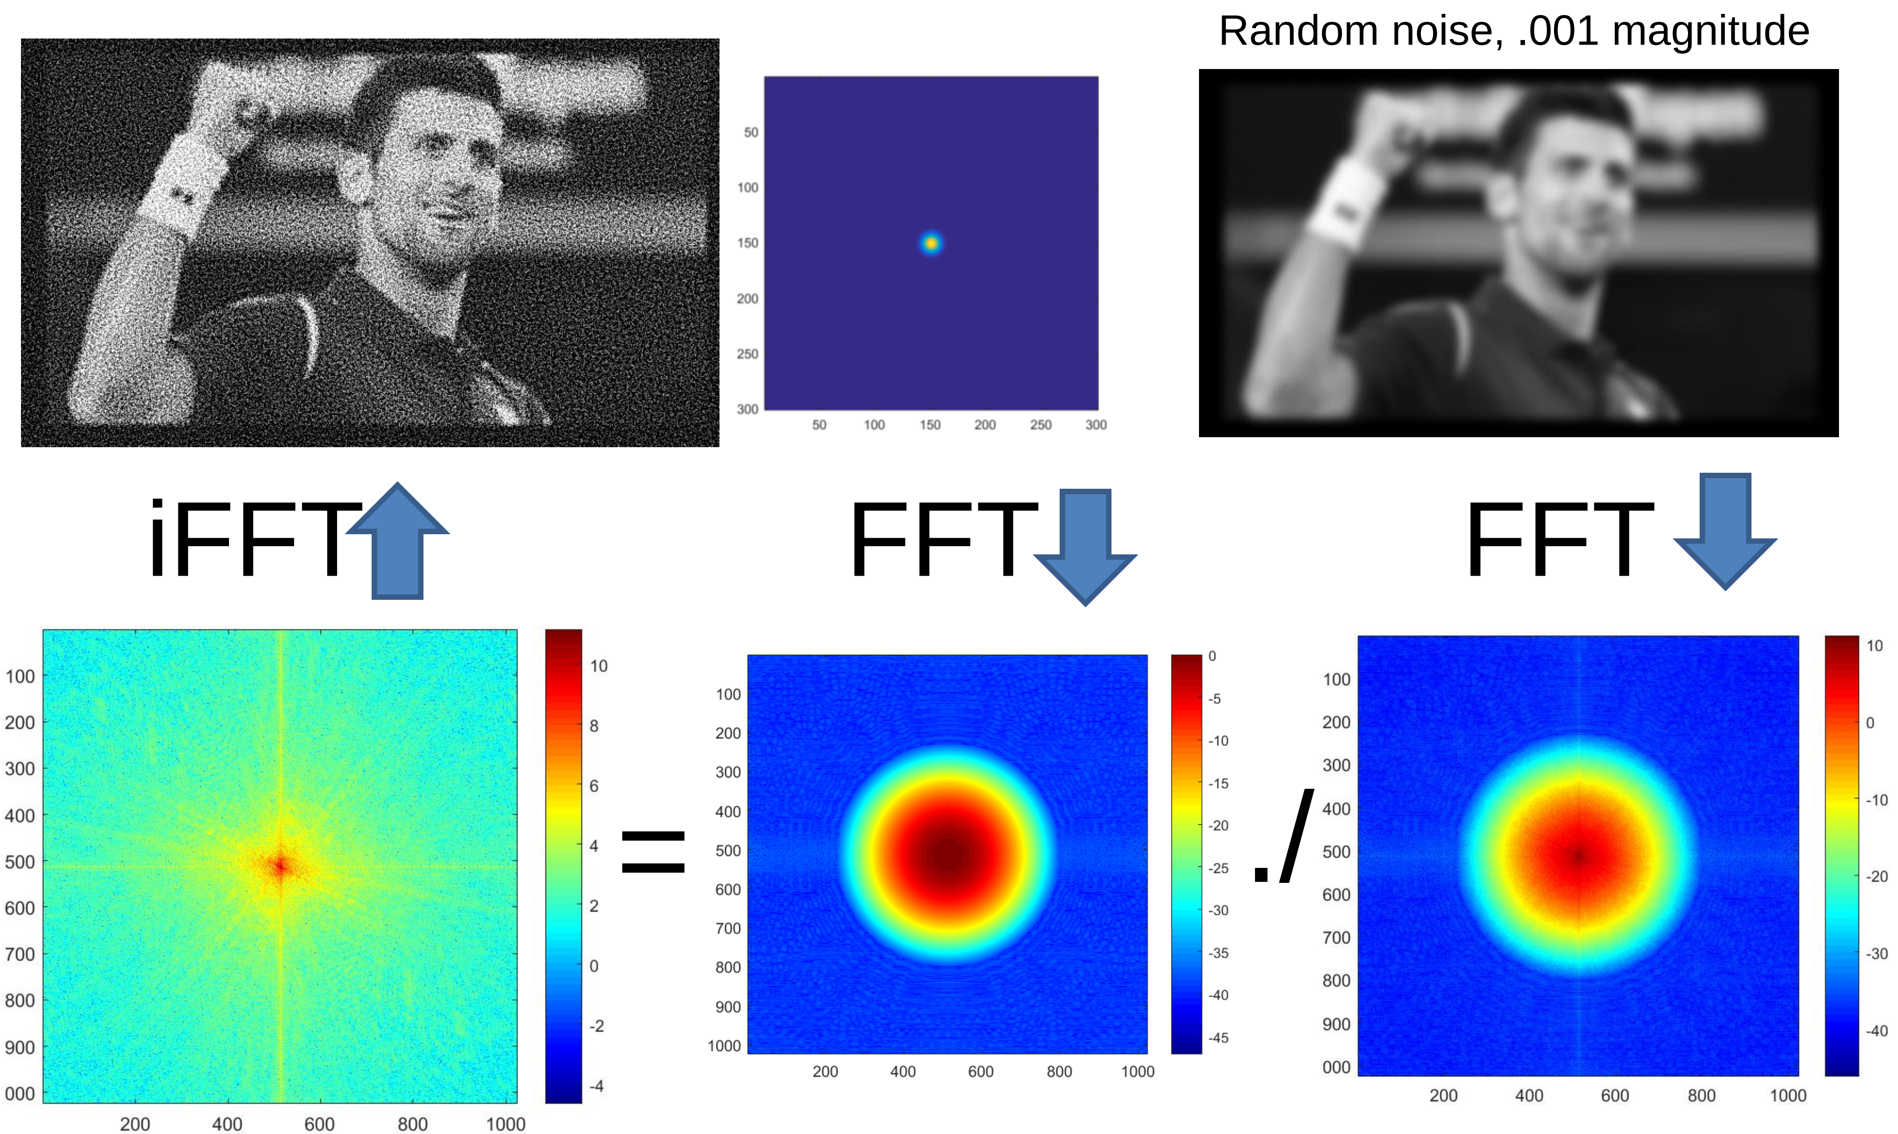

Deconvolution is hard

- Active research area.

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

- If you don’t know the filter (blind deconvolution) it is harder still.

Sampling

Why does a lower resolution image still make sense to us? What do we lose?

Image: http://www.flickr.com/photos/igorms/136916757/

Subsampling by a factor of 2

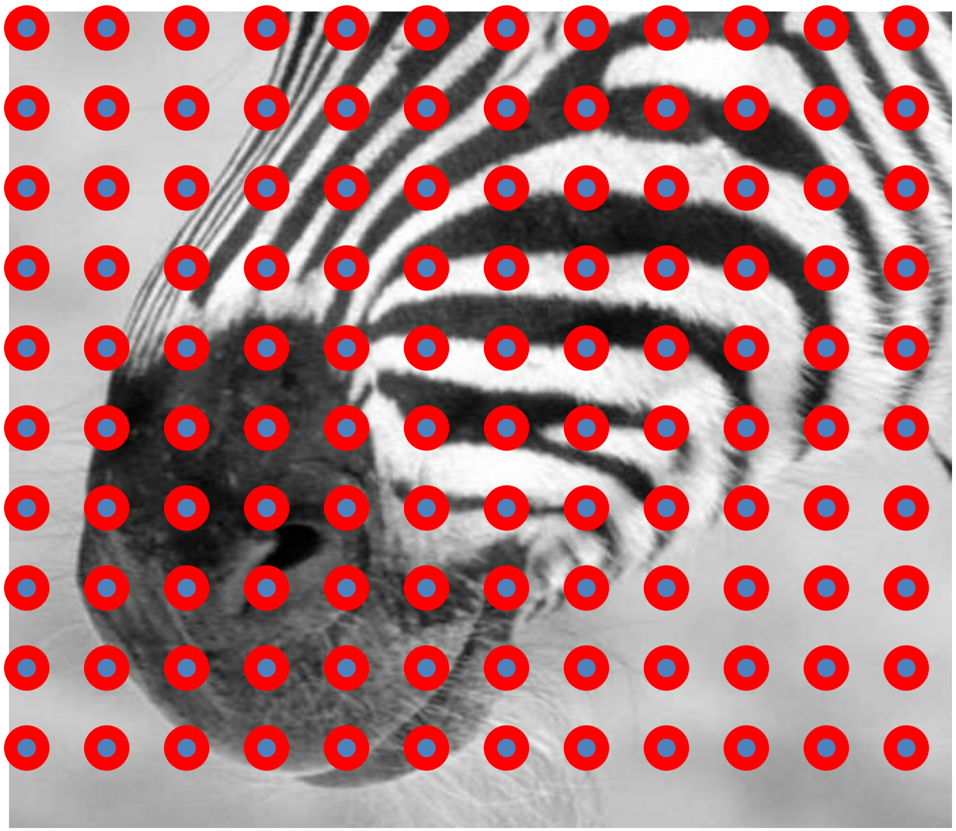

Throw away every other row and column to create a 1/2 size image

Aliasing problem

- 1D example (sinewave):

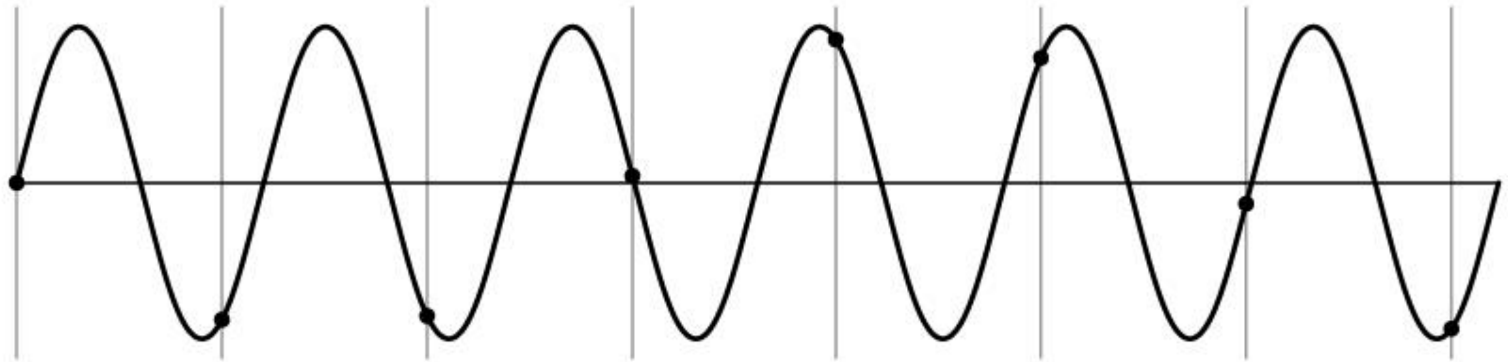

Source: S. Marschner

Aliasing problem

- 1D example (sinewave):

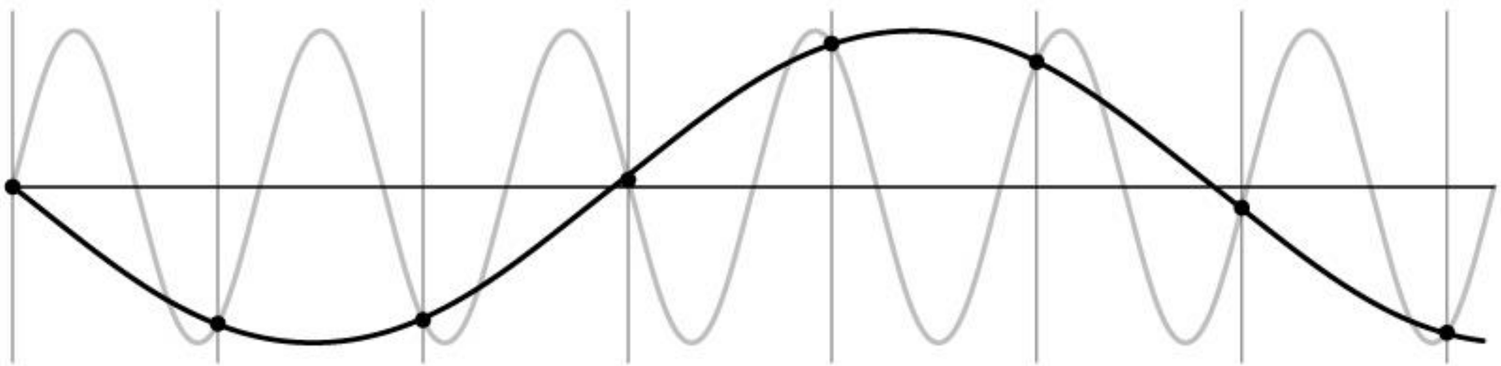

Source: S. Marschner

Aliasing problem

- Sub-sampling may be dangerous….
- Characteristic errors may appear:
  - “car wheels rolling the wrong way in movies”
  - “checkerboards disintegrate in ray tracing”
  - “striped shirts look funny on color television”
    - Moiré patterns

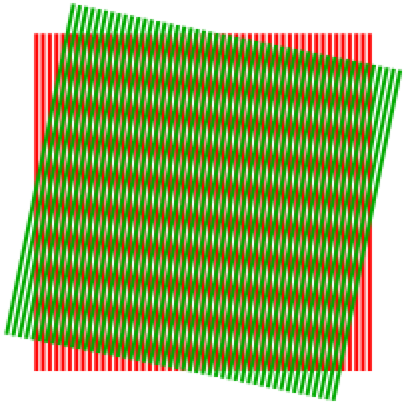

Source: D. Forsyth; Wikipedia

Aliasing in graphics

Source: A. Efros

Aliasing in video

Imagine a spoked wheel moving to the right (rotating clockwise). Mark wheel with dot so we can see what's happening.

If camera shutter is only open for a fraction of a frame time (frame time = 1/30 sec. for video, 1/24 sec. for film):

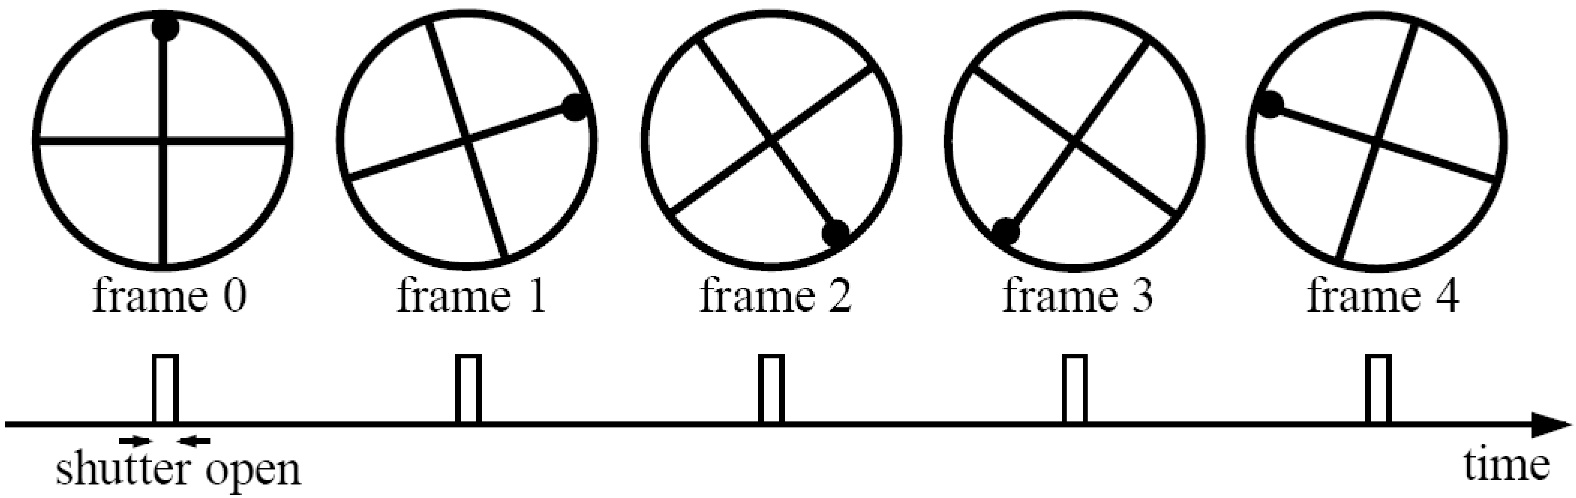

Without dot, wheel appears to be rotating slowly backwards! (counterclockwise)

Slide by Steve Seitz

Sampling and aliasing

256x256 128x128 64x64 32x32 16x16

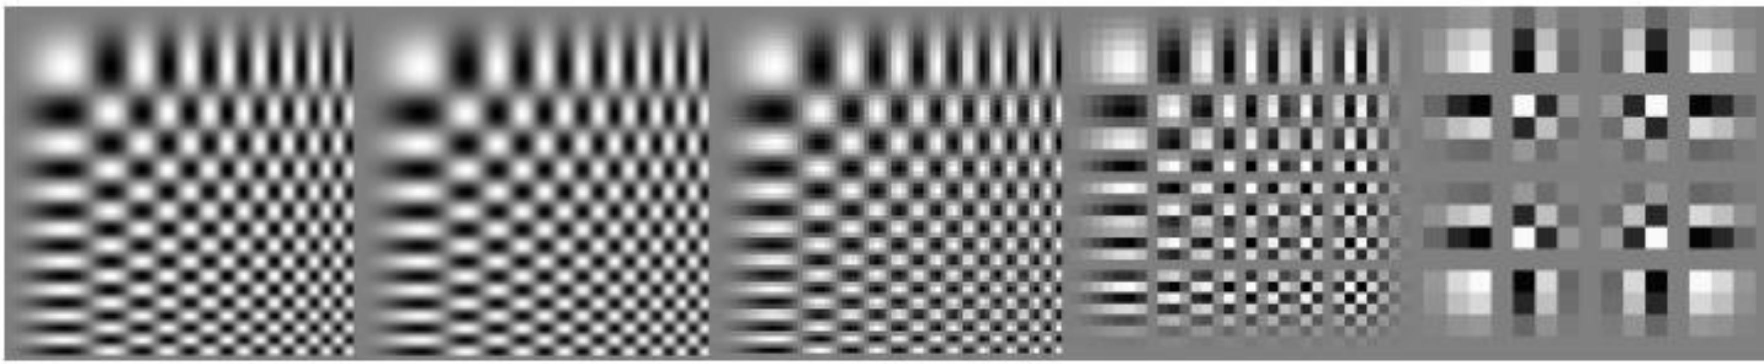

Nyquist-Shannon Sampling Theorem

- When sampling a signal at discrete intervals, the sampling frequency must be $\geq 2 \times f_{max}$
- f_{max} = max frequency of the input signal
- This will allows to reconstruct the original perfectly from the sampled version

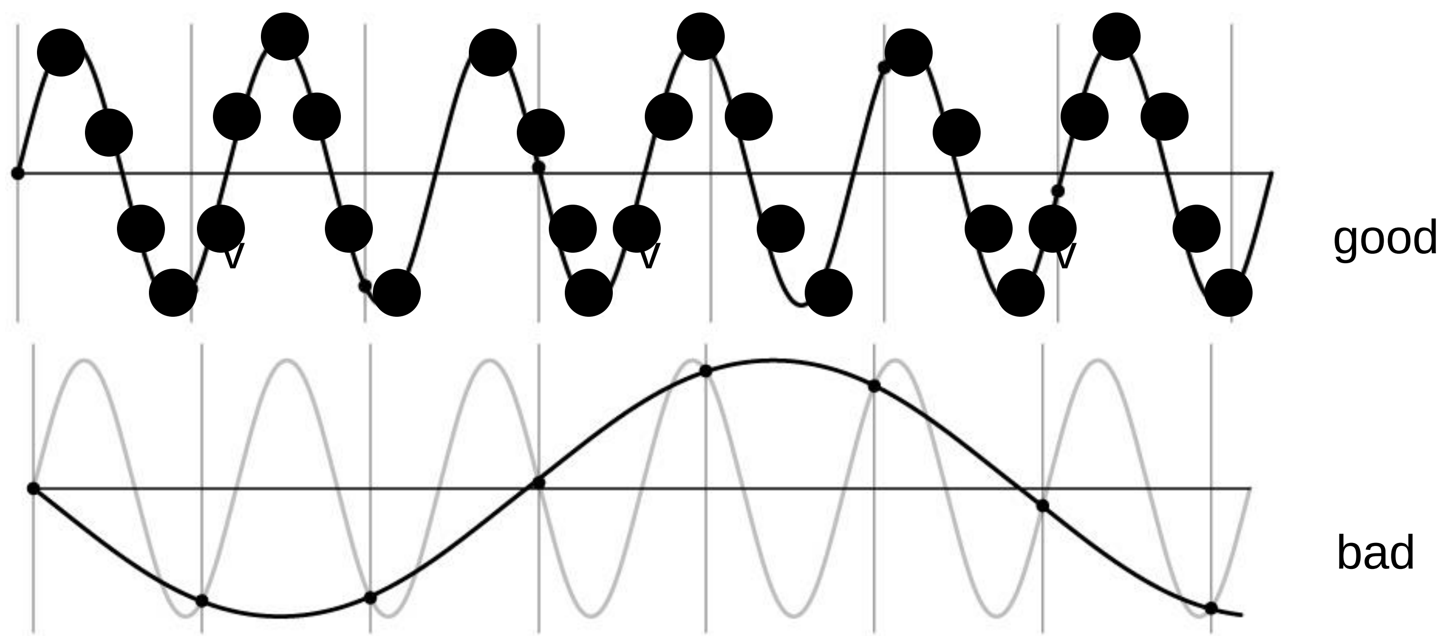

How to fix aliasing?

Solutions?

Better sensors

Solutions:

- Sample more often

Anti-aliasing

Solutions:

- Sample more often

- Get rid of all frequencies that are greater than half the new sampling frequency
  - Will lose information
  - But it's better than aliasing
  - Apply a smoothing filter

Anti-aliasing

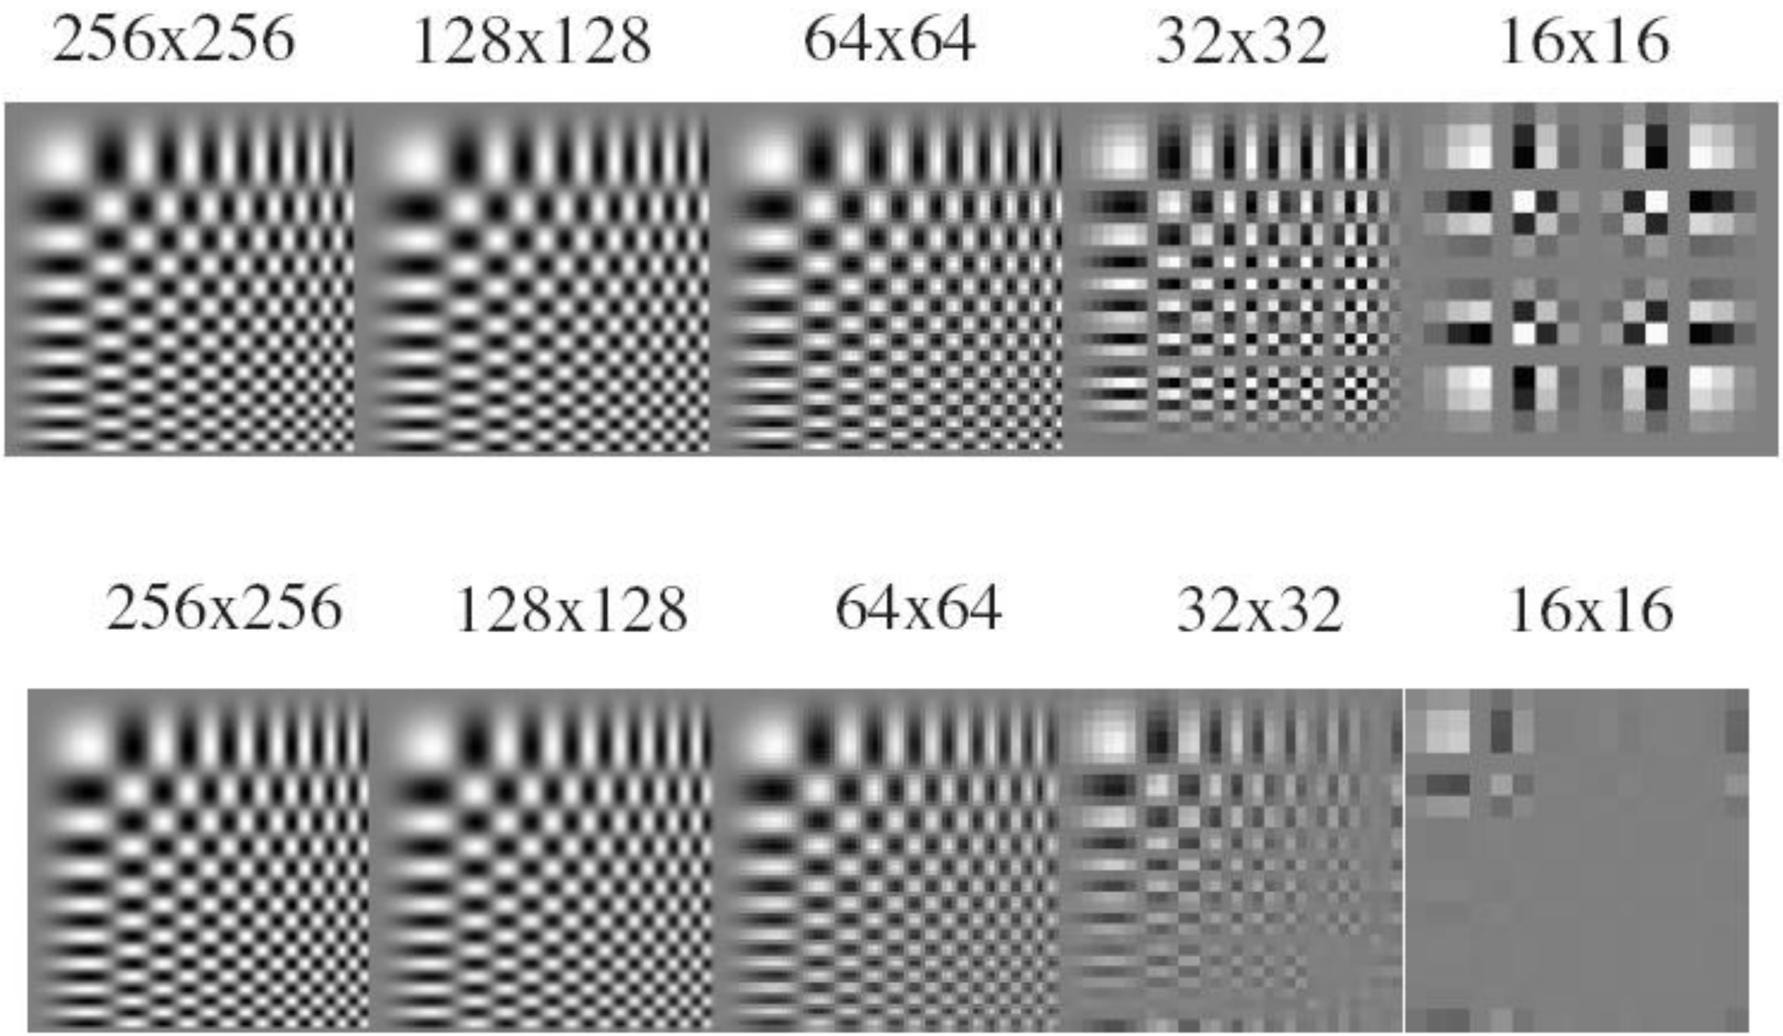

Forsyth and Ponce 2002

Algorithm for downsampling by factor of 2

1. Start with image(h, w)
2. Apply low-pass filter

   im_blur = imfilter(image, fspecial('gaussian', 7, 1))

3. Sample every other pixel

   im_small = im_blur(1:2:end, 1:2:end);

Hays

Subsampling without pre-filtering

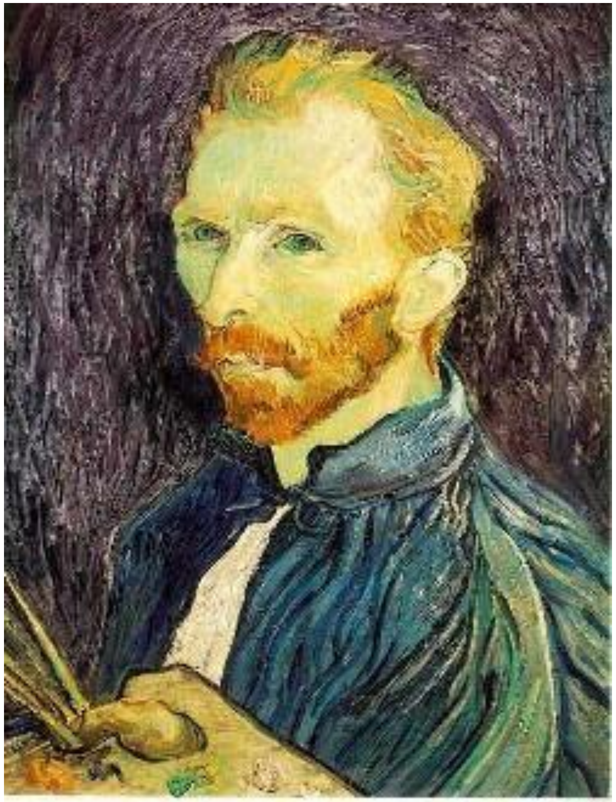

1/2

1/4 (2x zoom)

1/8 (4x zoom)

Slide by Steve Seitz

Subsampling with Gaussian pre-filtering

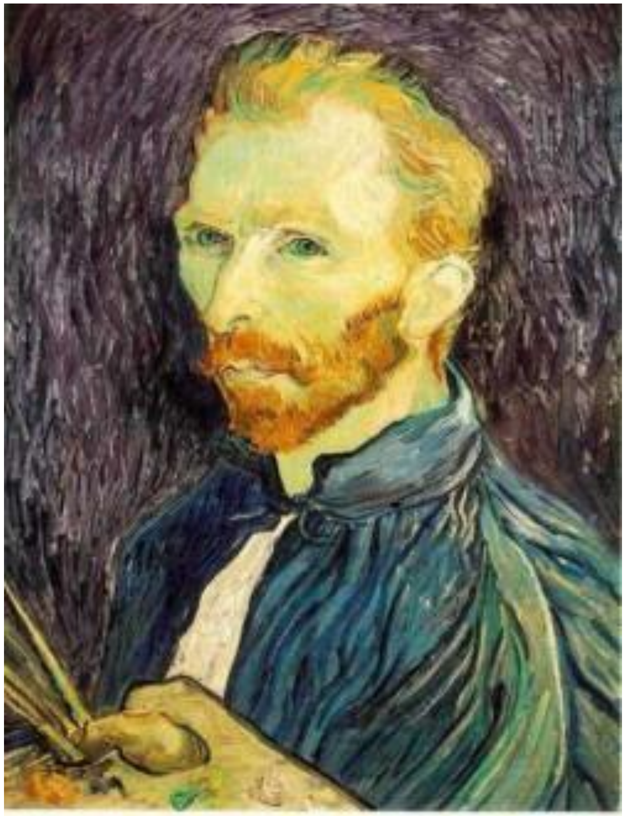

Gaussian 1/2

G 1/4

G 1/8

Slide by Steve Seitz

Image Pyramids

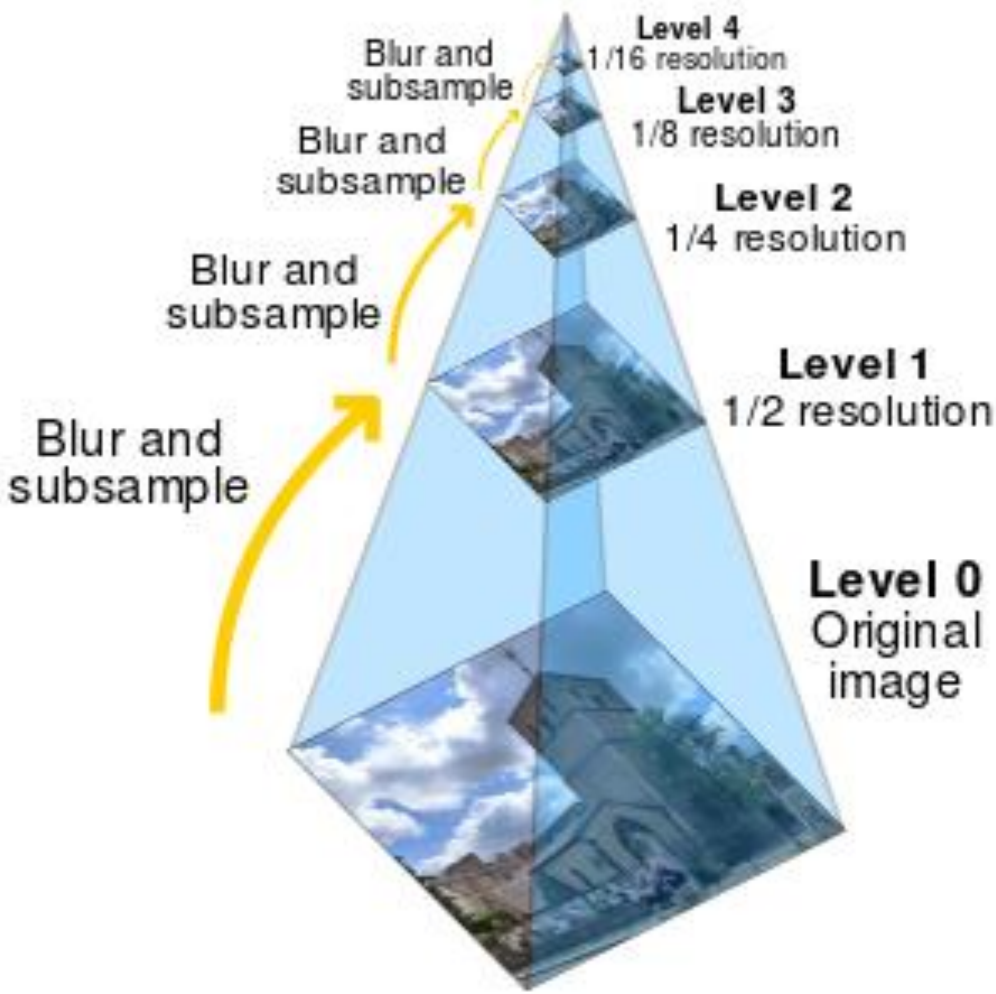

Wikipedia – Image Pyramids

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

Salvador Dali invented Hybrid Images?

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