2d Fourier Transform

$$F(u,v) = \sum_{x=0}^{M-1} \sum_{y=0}^{N-1} f(x,y) e^{-j2\pi\left[\frac{ux}{M}+\frac{vy}{N}\right]}$$

Fourier Bases

Teases away fast vs. slow changes in the image.

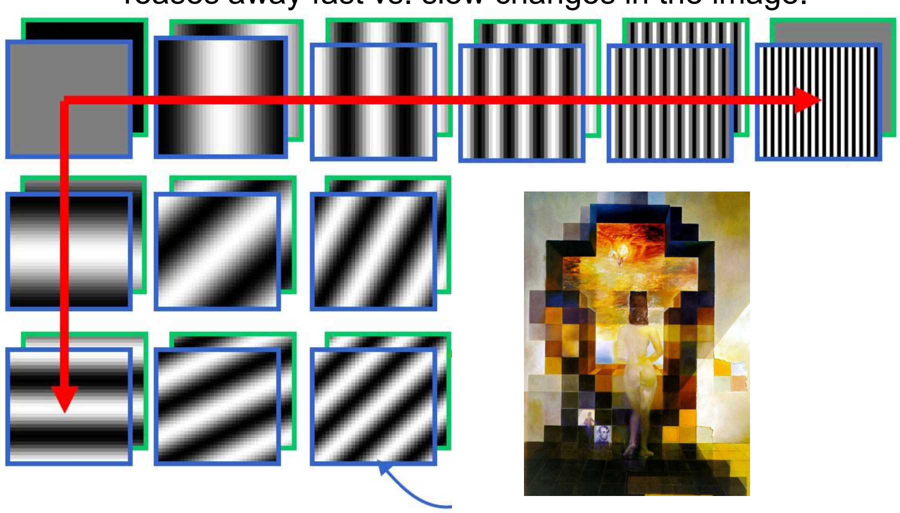

This change of basis is the Fourier Transform

Fourier Bases

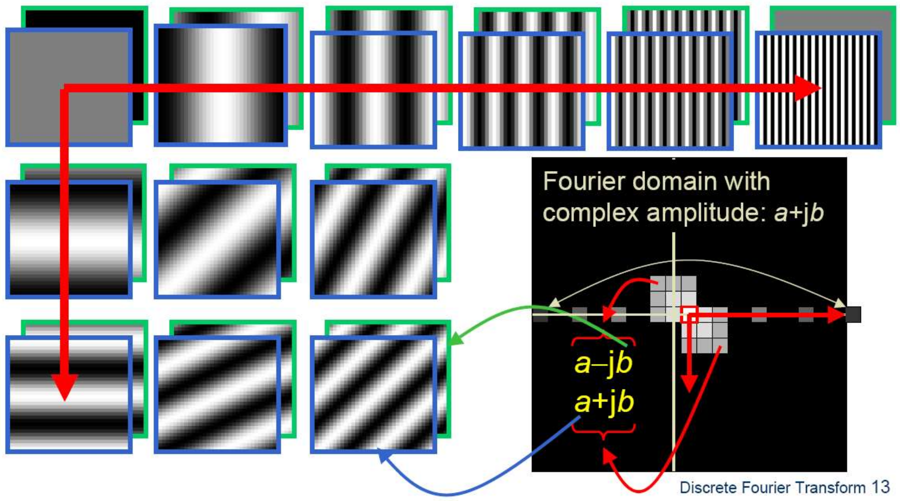

in Matlab, check out: imagesc(log(abs(fftshift(fft2(im)))));

Man-made Scene

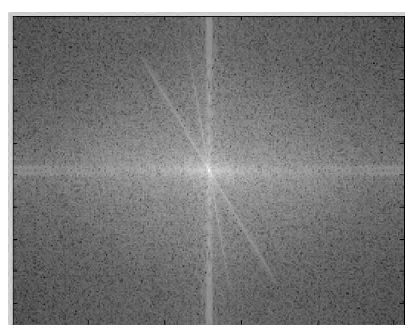

Can change spectrum, then reconstruct

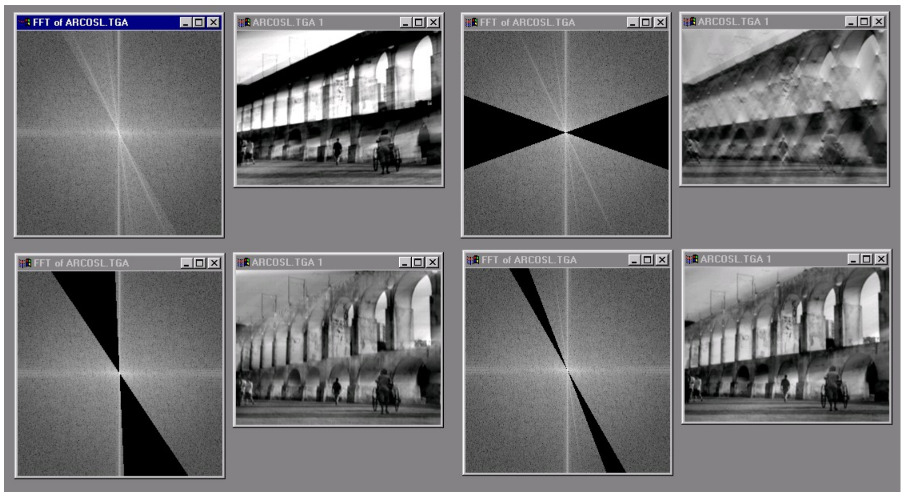

Low and High Pass filtering

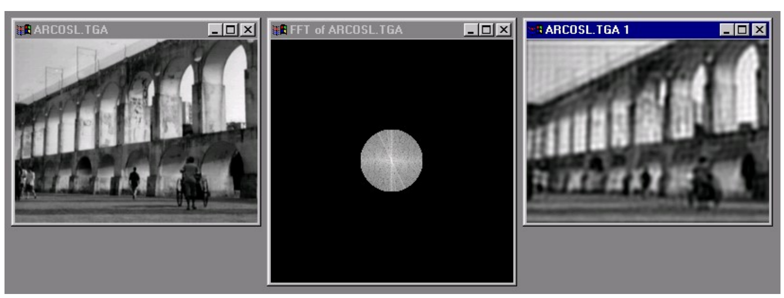

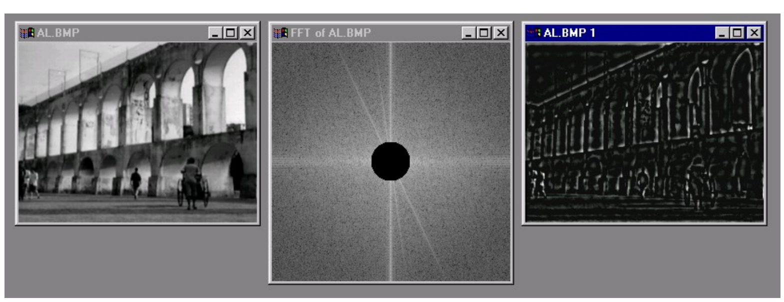

Sinc Filter

- What is the spatial representation of the hard cutoff in the frequency domain?

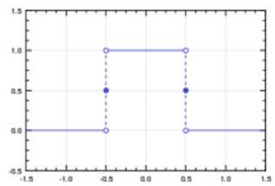

Frequency Domain

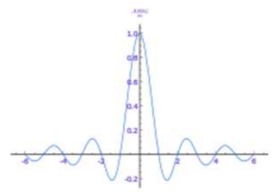

Spatial Domain

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!

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

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)

Filtering in spatial domain

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

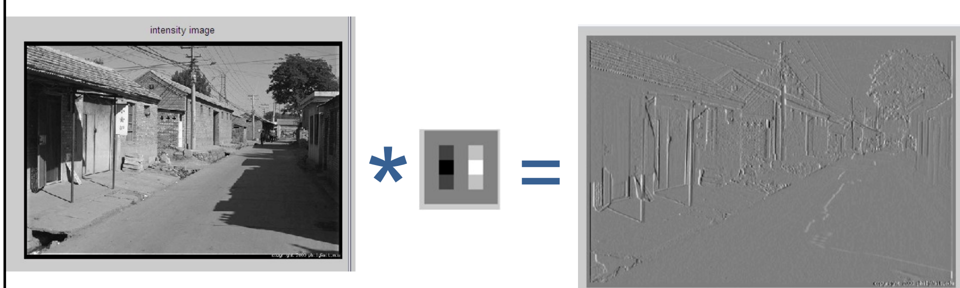

Filtering in frequency domain

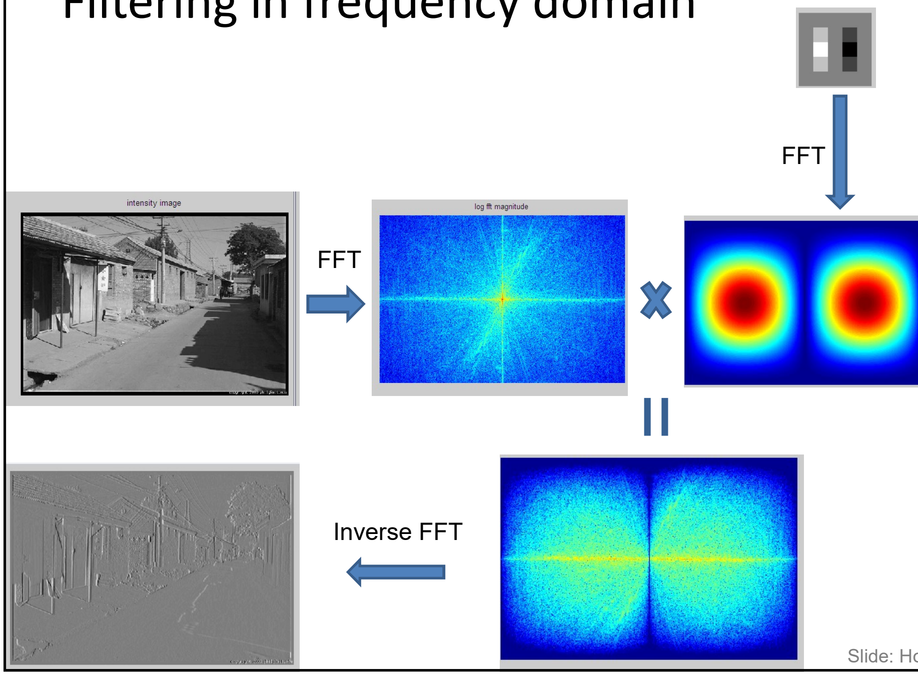

Slide: Hoiem

Filtering

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

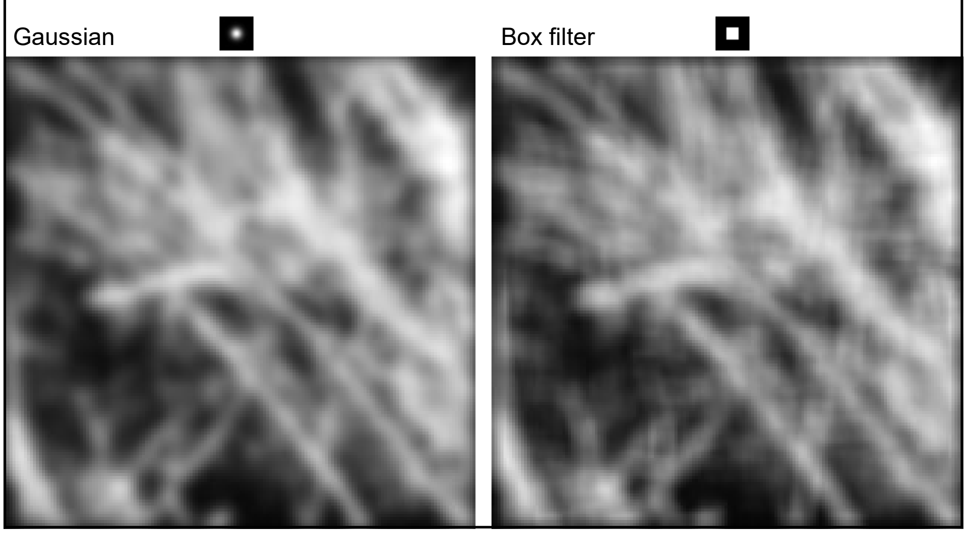

Gaussian

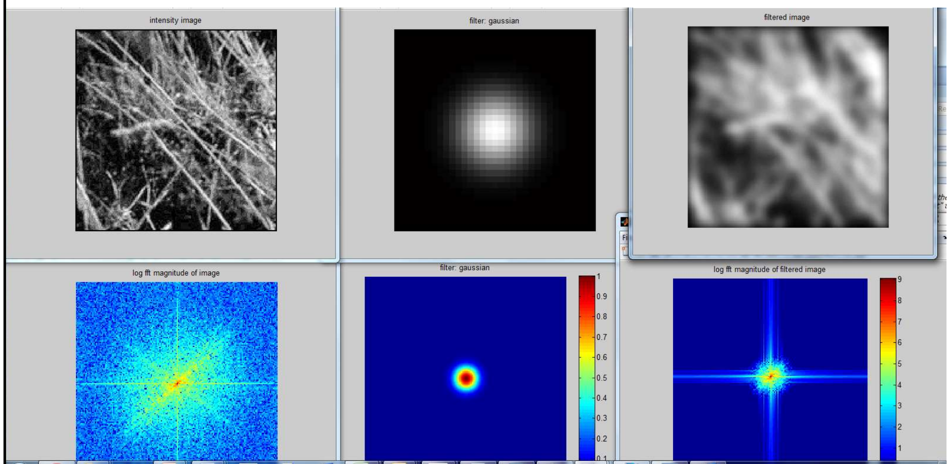

Box Filter

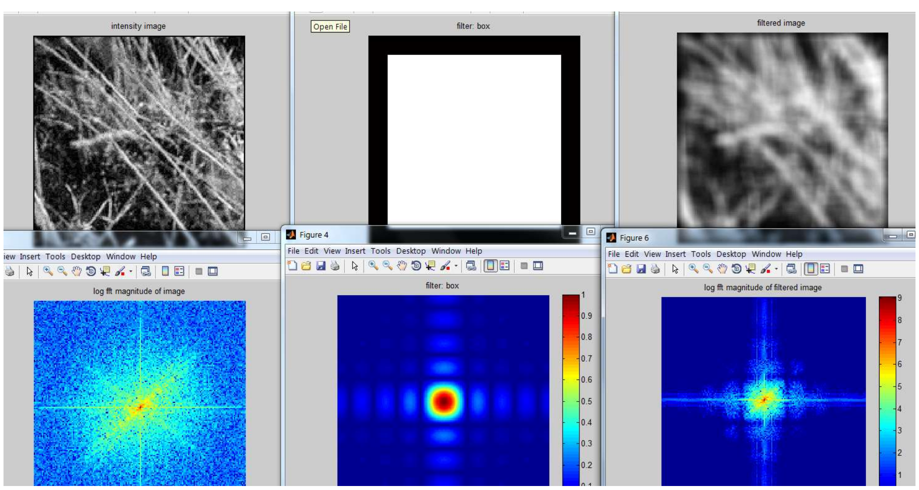

But you can't invert multiplication by 0

- But it's not quite zero, is it…

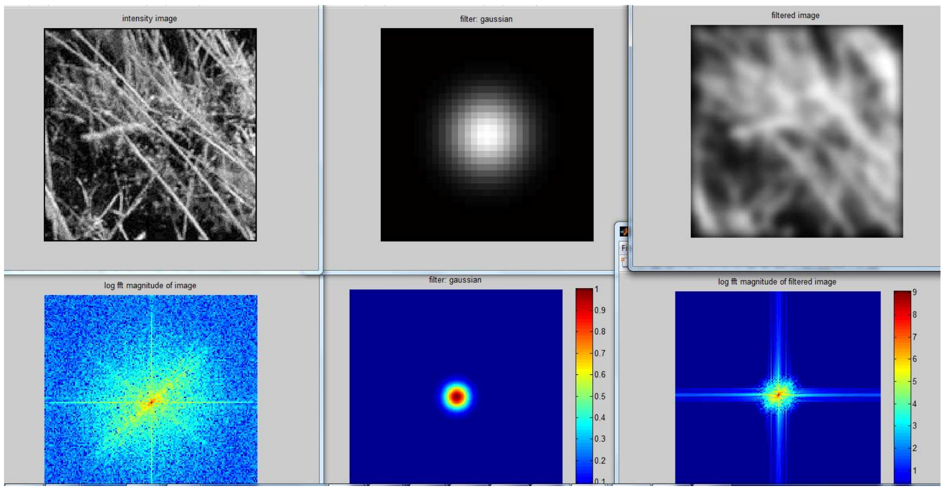

Convolution

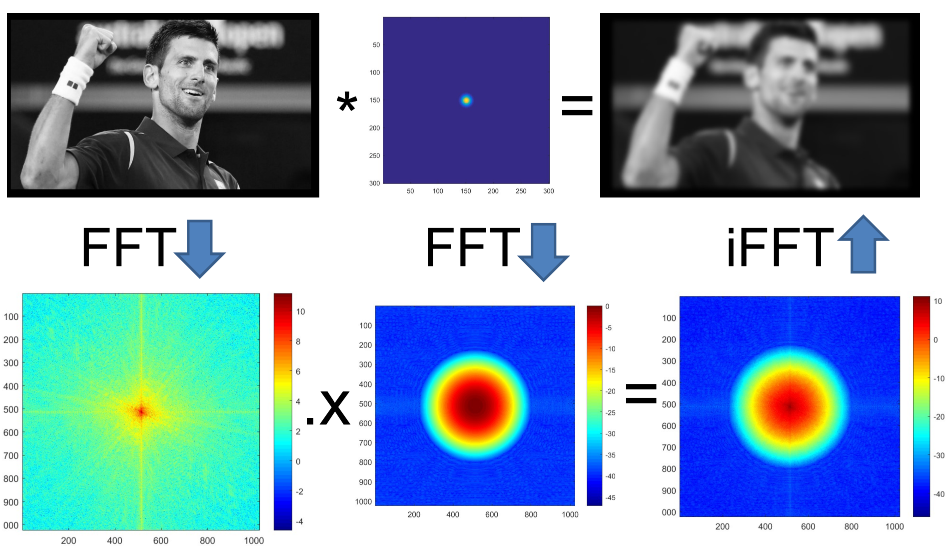

Deconvolution?

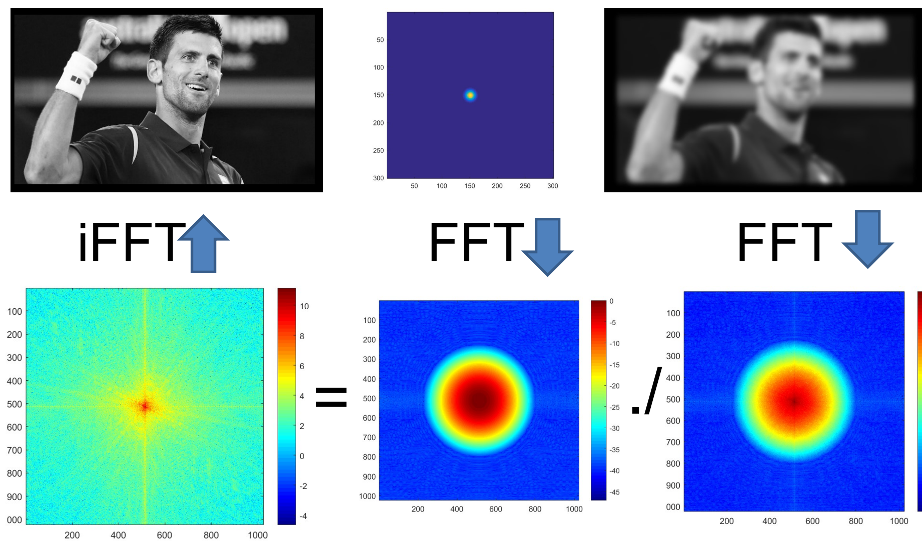

But under more realistic conditions

Random noise, .000001 magnitude

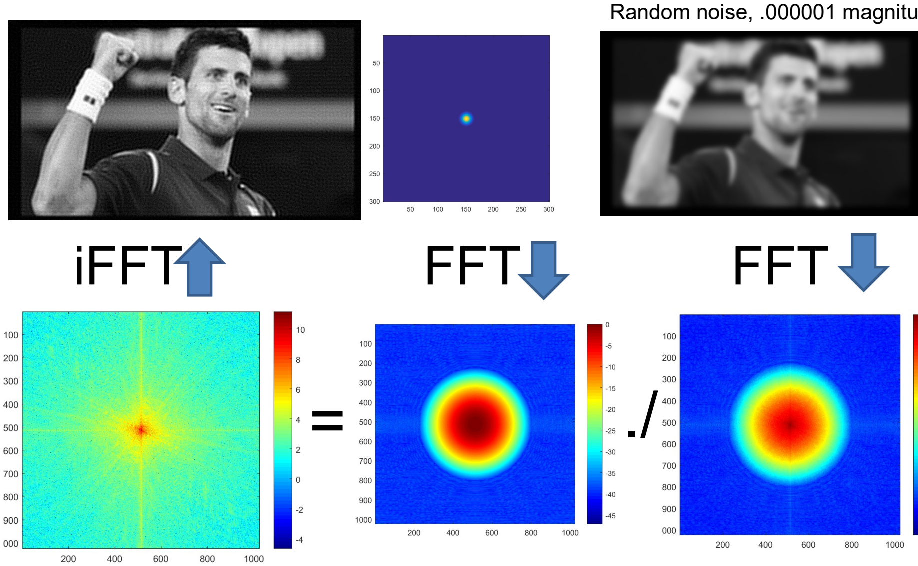

But under more realistic conditions

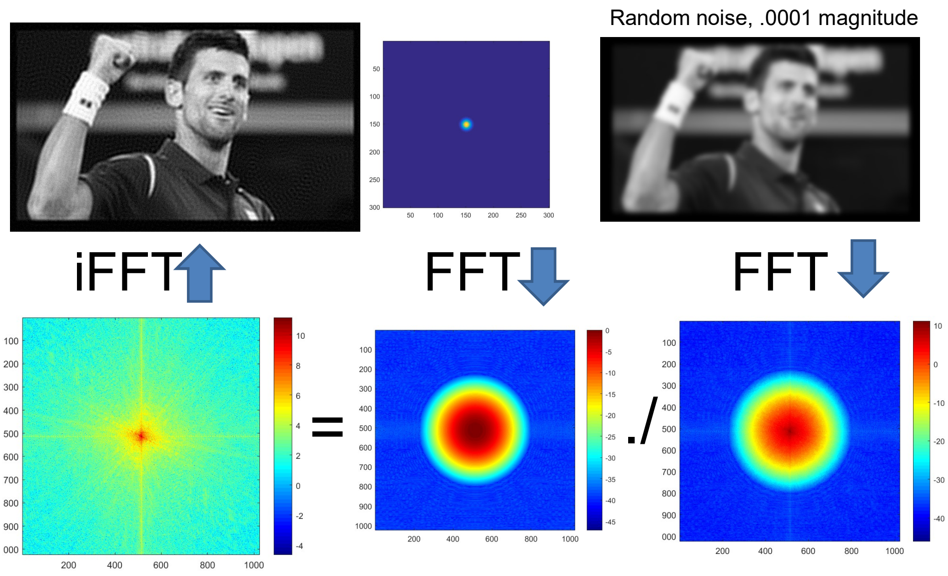

But under more realistic conditions

With a random filter…

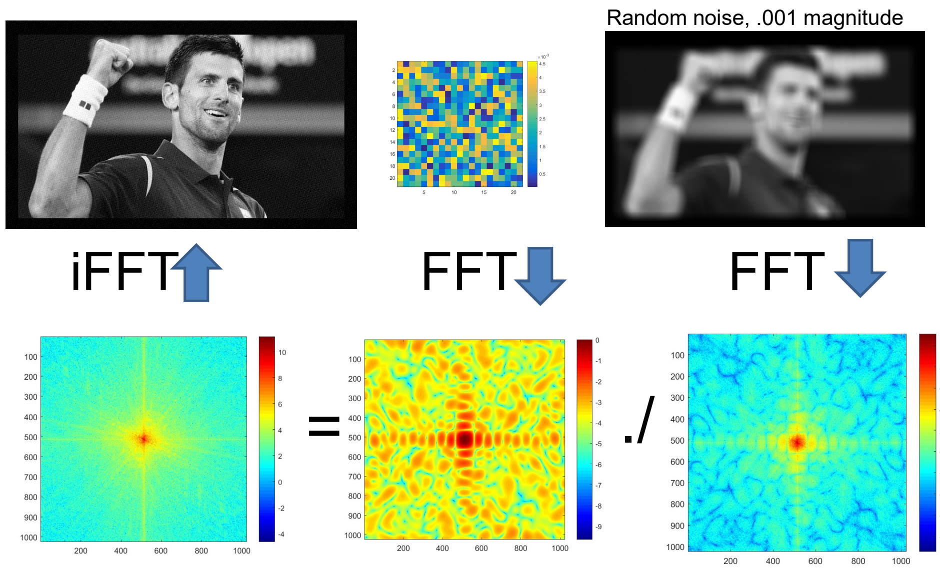

Deconvolution is hard

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

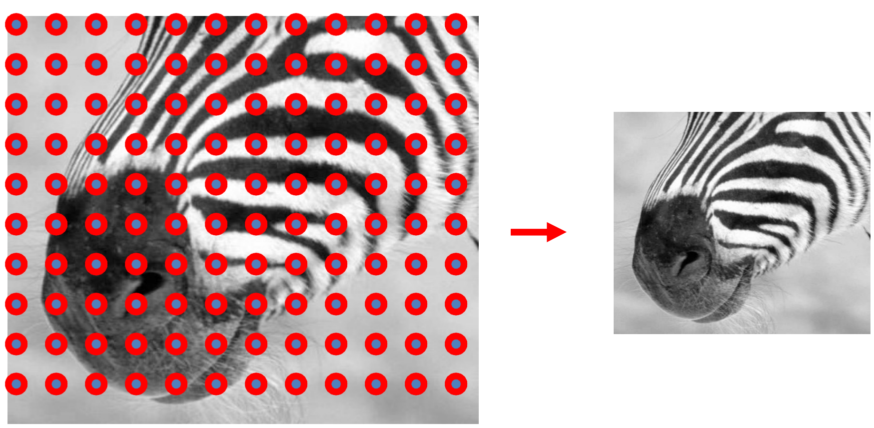

Throw away every other row and column to create a 1/2 size image

Aliasing problem

- 1D example (sinewave):

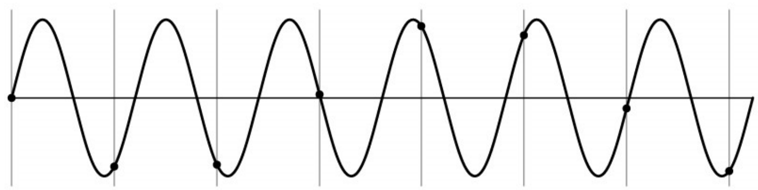

Source: S. Marschner

Aliasing problem

- 1D example (sinewave):

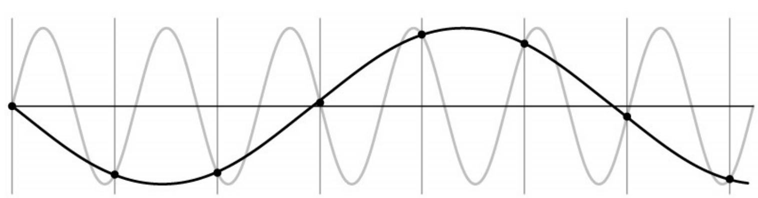

Source: S. Marschner

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”

Source: D. Forsyth

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):

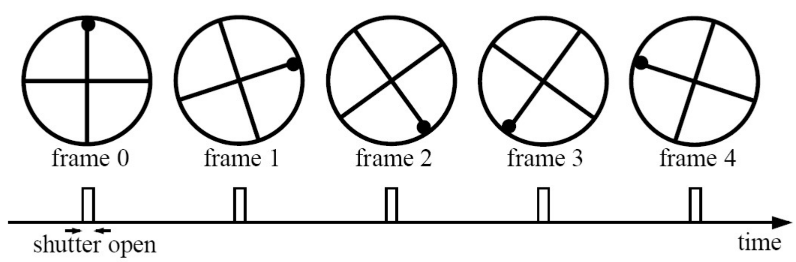

Without dot, wheel appears to be rotating slowly backwards! (counterclockwise)

Slide by Steve Seitz

Aliasing in graphics

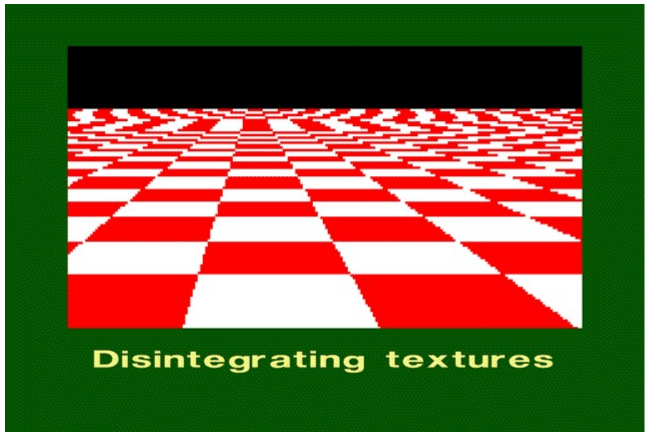

Source: A. Efros

Sampling and aliasing

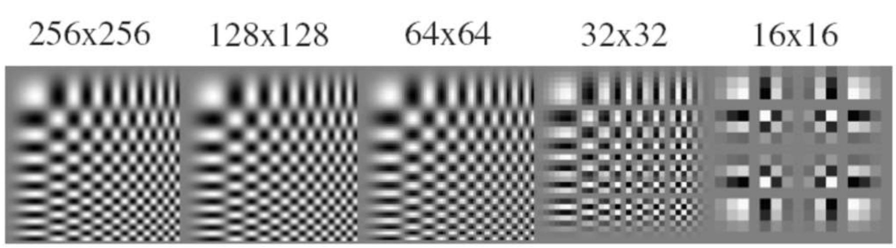

Nyquist-Shannon Sampling Theorem

- 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

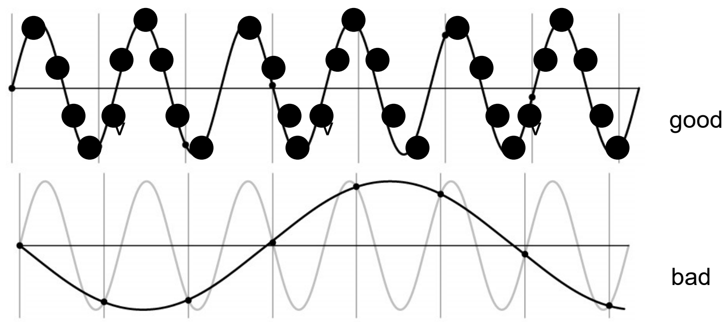

Anti-aliasing

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

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);

Anti-aliasing

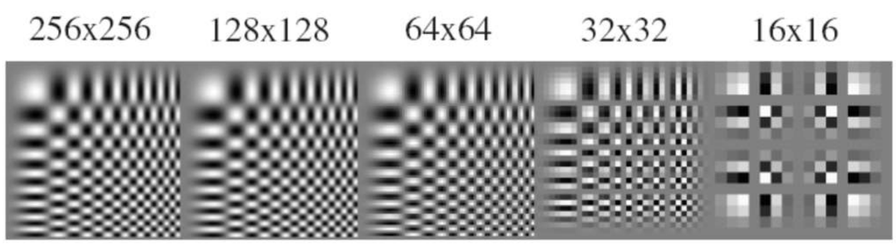

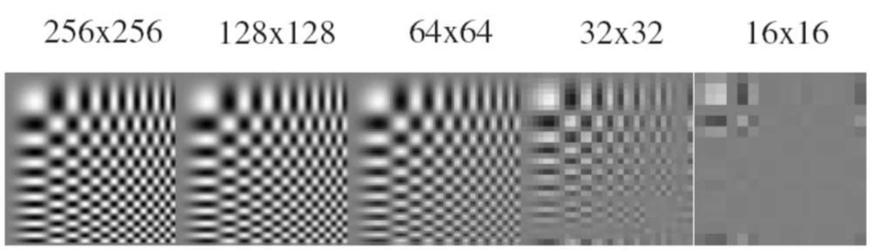

Forsyth and Ponce 2002

Subsampling without pre-filtering

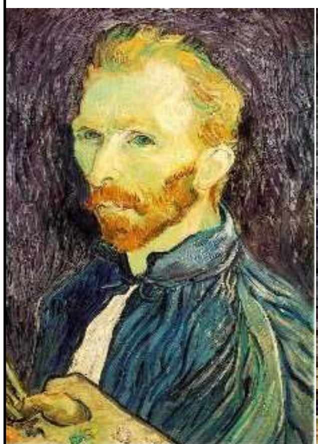

1/2

1/4 (2x zoom)

1/8 (4x zoom)

Slide by Steve Seitz

Subsampling with Gaussian pre-filtering

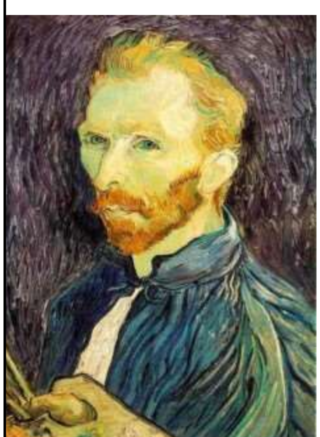

Gaussian 1/2

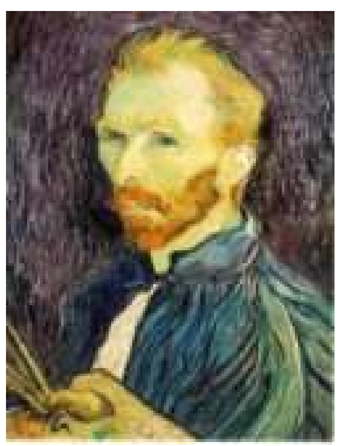

G 1/4

G 1/8

Slide by Steve Seitz