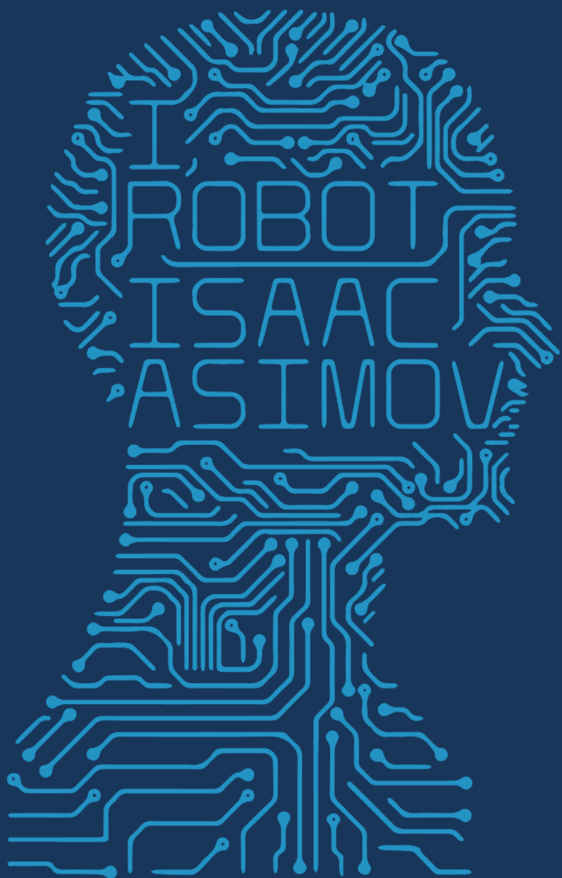

1950

Future Vision

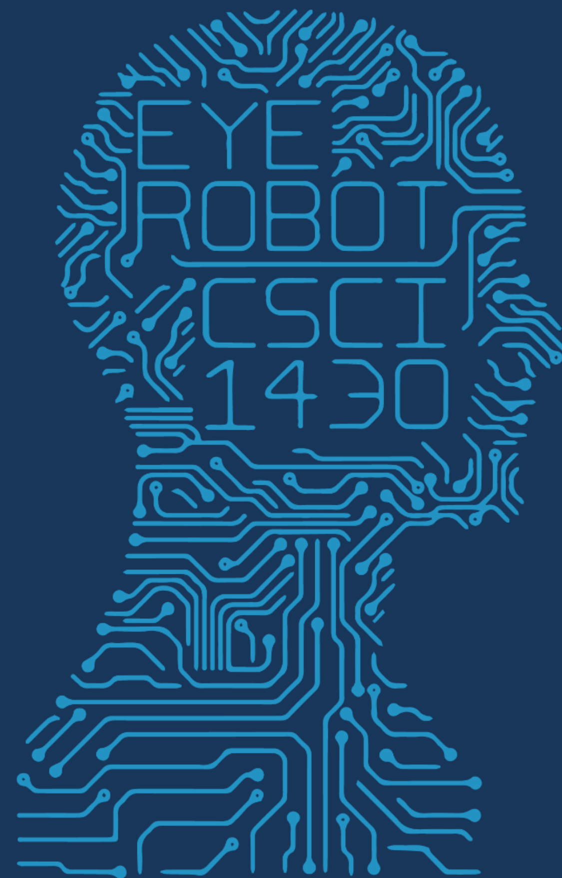

2017 MWF 1pm 368

Computer Vision

Recap of Monday

- Linear filtering

$$h[m,n] = \sum_{k,l} f[k,l]\, I[m+k, n+l]$$

$$\frac{1}{9}\begin{bmatrix}1 & 1 & 1\\ 1 & 1 & 1\\ 1 & 1 & 1\end{bmatrix}$$

  - Not a matrix multiplication
  - Sum over Hadamard product
  - Can smooth, sharpen, translate (among many other uses)

- Be aware of details for filter size, extrapolation, cropping

James Hays

Questions from Monday

- DOUBLE vs UINT8
  - MATLAB coping strategies

Questions from Monday

- DOUBLE vs UINT8
  - MATLAB coping strategies

- What happens to negative numbers?
- Shifting the image +0.5
- Scaling edge response for visualization.

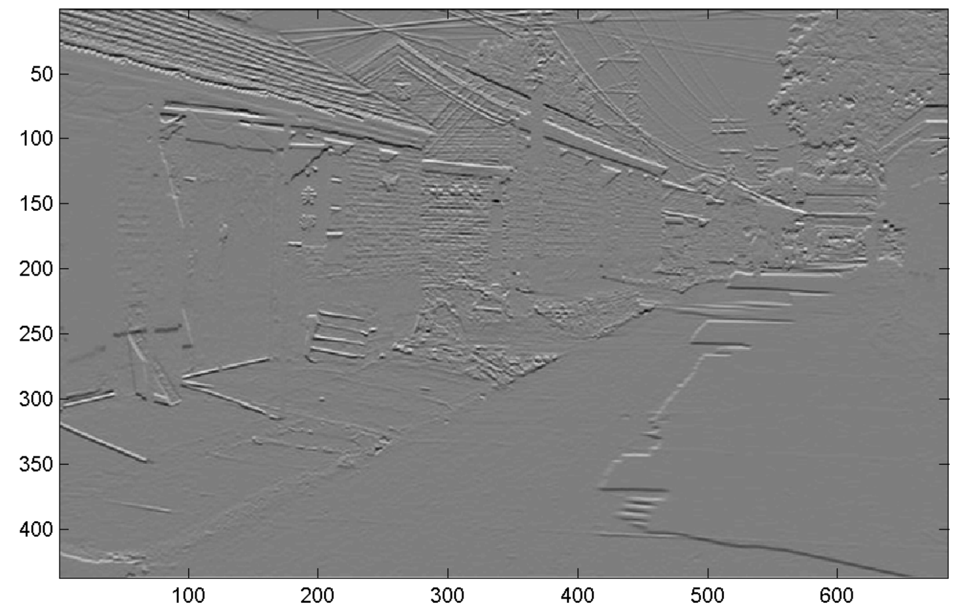

NON-LINEAR FILTERS

Median filters

- Operates over a window by selecting the median intensity in the window.
- ‘Rank’ filter as based on ordering of gray levels
  - E.G., min, max, range filters

Slide by Steve Seitz

Image filtering - mean

$f[\cdot,\cdot]\ \frac{1}{9}$

1	1	1
1	1	1
1	1	1

$I[.,.]$

0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0
0	0	0	90	90	90	90	90	0	0
0	0	0	90	90	90	90	90	0	0
0	0	0	90	90	90	90	90	0	0
0	0	0	90	0	90	90	90	0	0
0	0	0	90	90	90	90	90	0	0
0	0	0	0	0	0	0	0	0	0
0	0	90	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0

$h[.,.]$

	0	10	20	30	30				
				?					

$$h[m,n] = \sum_{k,l} f[k,l]\, I[m+k, n+l]$$

Credit: S. Seitz

Image filtering - mean

$f[\cdot,\cdot]\ \frac{1}{9}$

1	1	1
1	1	1
1	1	1

$I[.,.]$

0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0
0	0	0	90	90	90	90	90	0	0
0	0	0	90	90	90	90	90	0	0
0	0	0	90	90	90	90	90	0	0
0	0	0	90	0	90	90	90	0	0
0	0	0	90	90	90	90	90	0	0
0	0	0	0	0	0	0	0	0	0
0	0	90	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0

$h[.,.]$

	0	10	20	30	30				
				50					

$$h[m,n]=\sum_{k,l} f[k,l]\,I[m+k,n+l]$$

Credit: S. Seitz

Median filter?

$I[.,.]$

0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0
0	0	0	90	90	90	90	90	0	0
0	0	0	90	90	90	90	90	0	0
0	0	0	90	90	90	90	90	0	0
0	0	0	90	0	90	90	90	0	0
0	0	0	90	90	90	90	90	0	0
0	0	0	0	0	0	0	0	0	0
0	0	90	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0

$h[.,.]$

				?					

Credit: S. Seitz

Median filters

- Operates over a window by selecting the median intensity in the window.
- What advantage does a median filter have over a mean filter?

Slide by Steve Seitz

Noisy Jack – Salt and Pepper

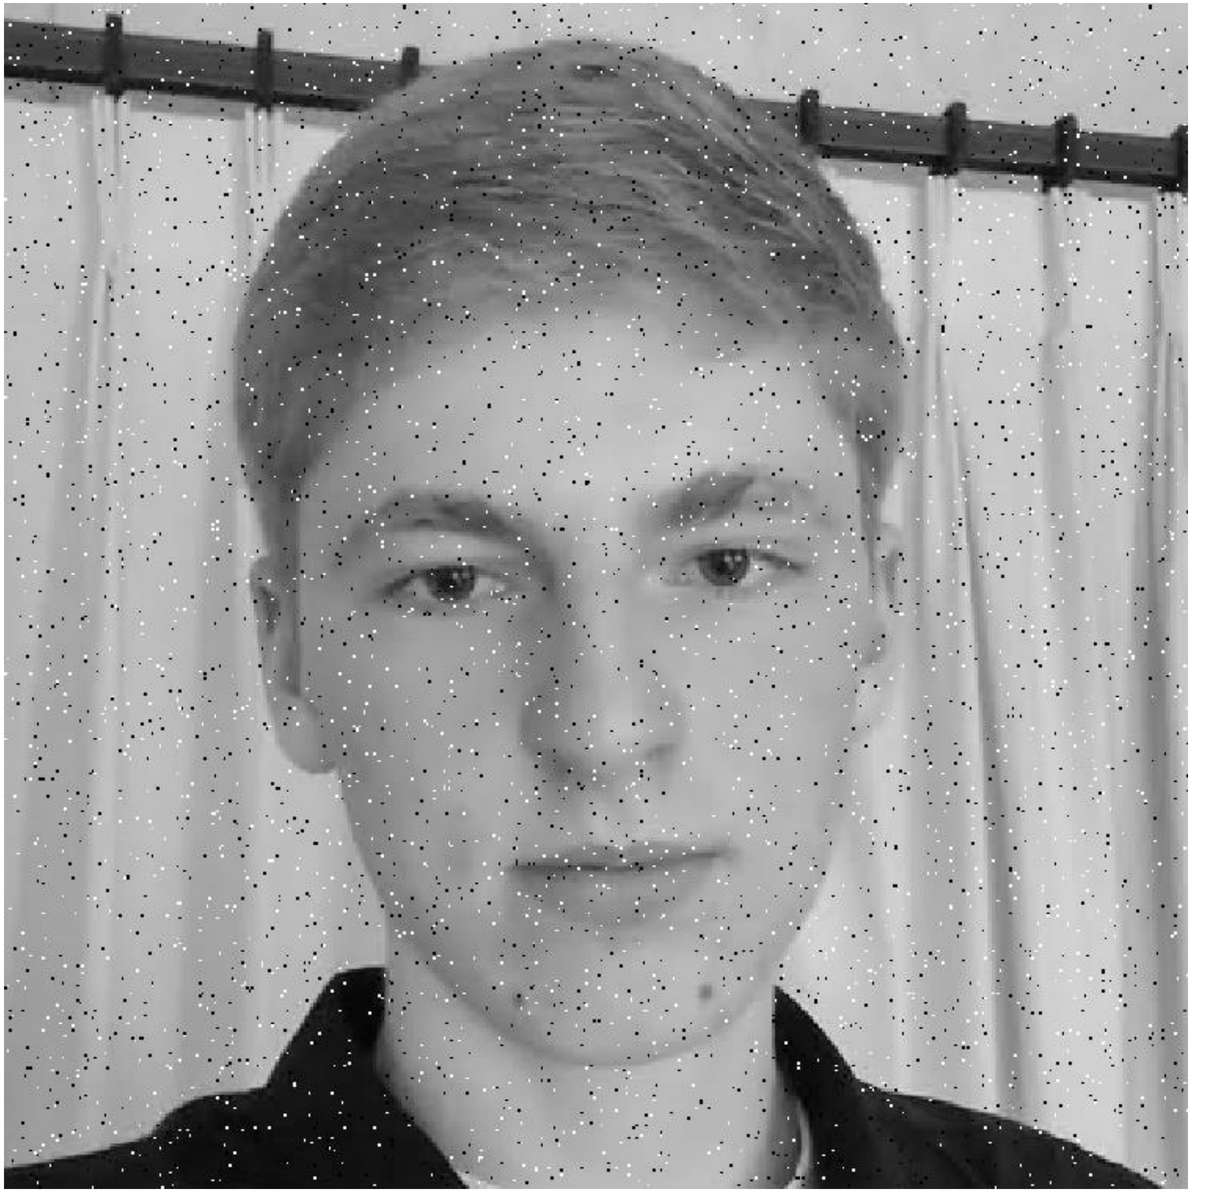

Mean Jack – 3 x 3 filter

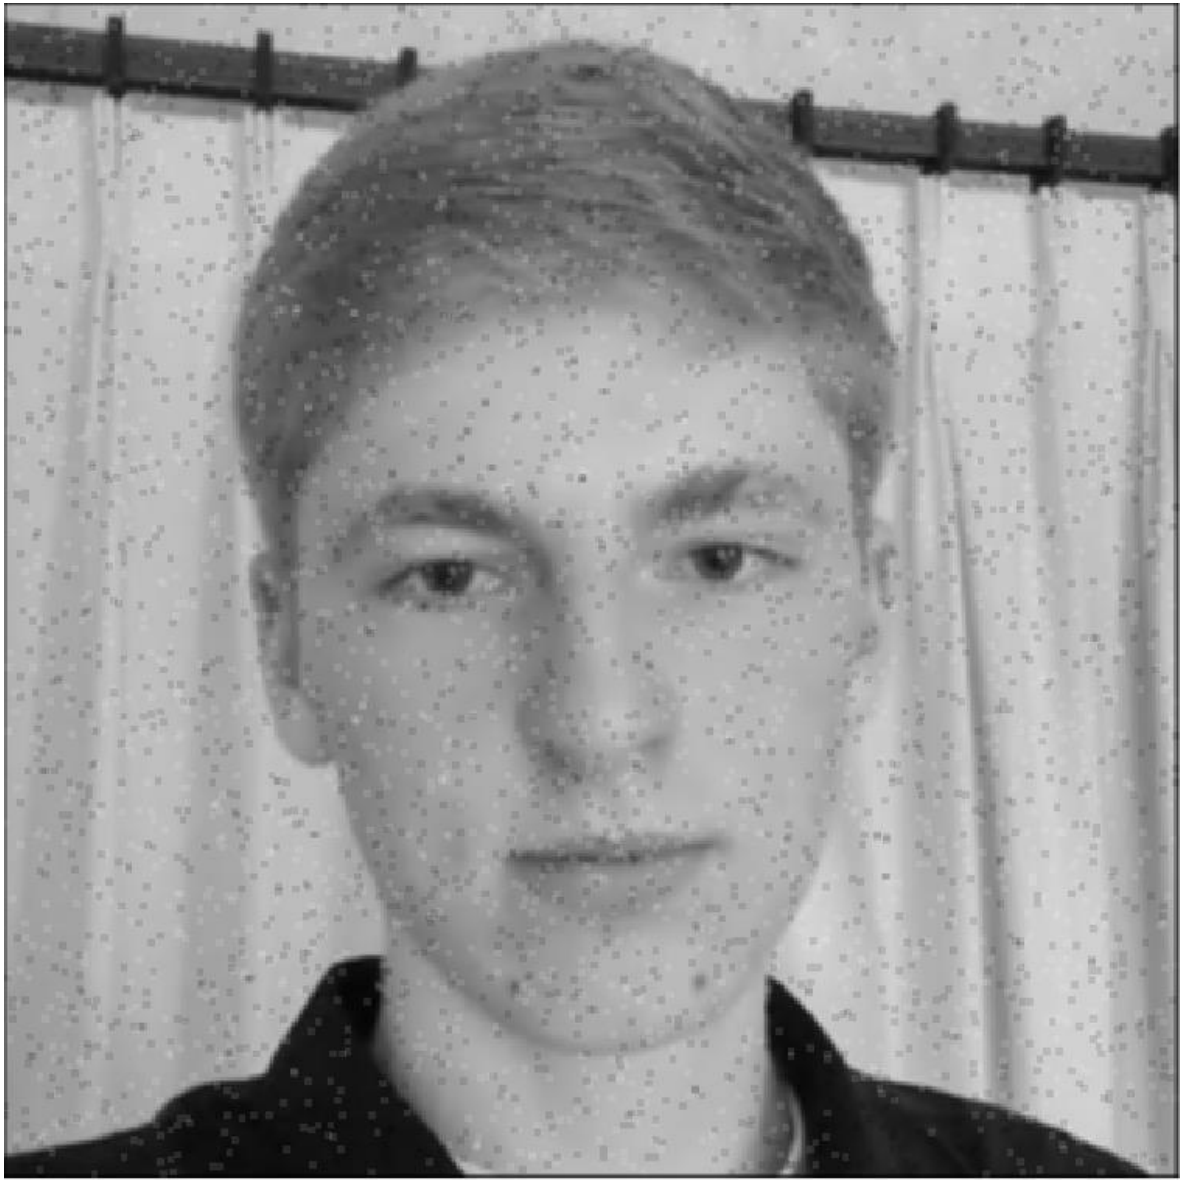

Very Mean Jack – 11 x 11 filter

Noisy Jack – Salt and Pepper

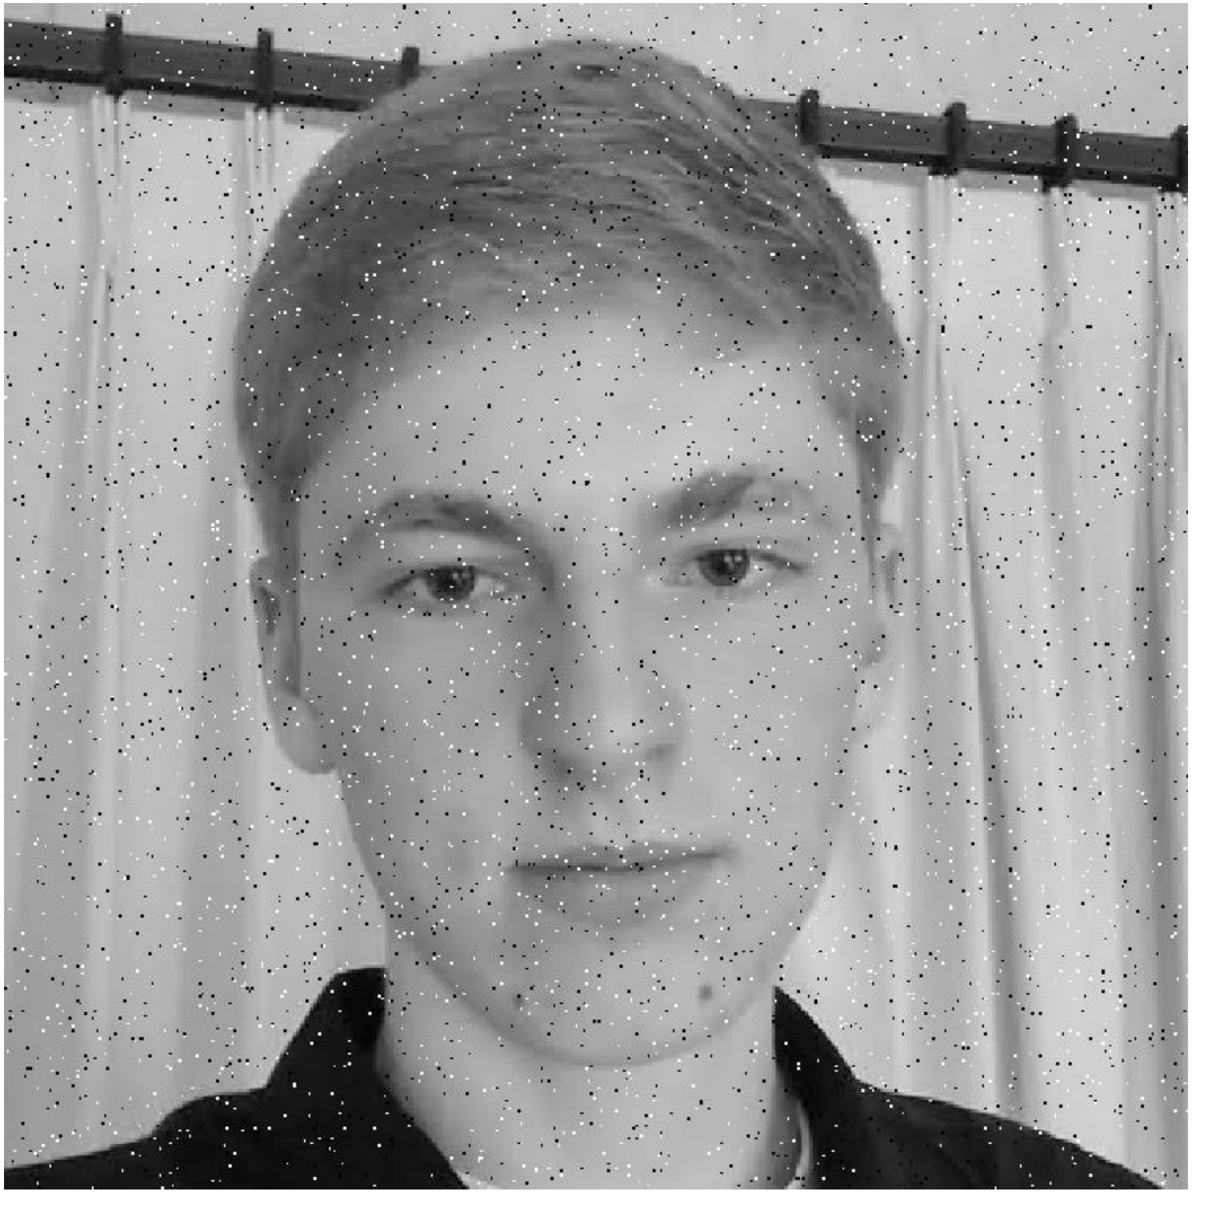

Median Jack – 3 x 3

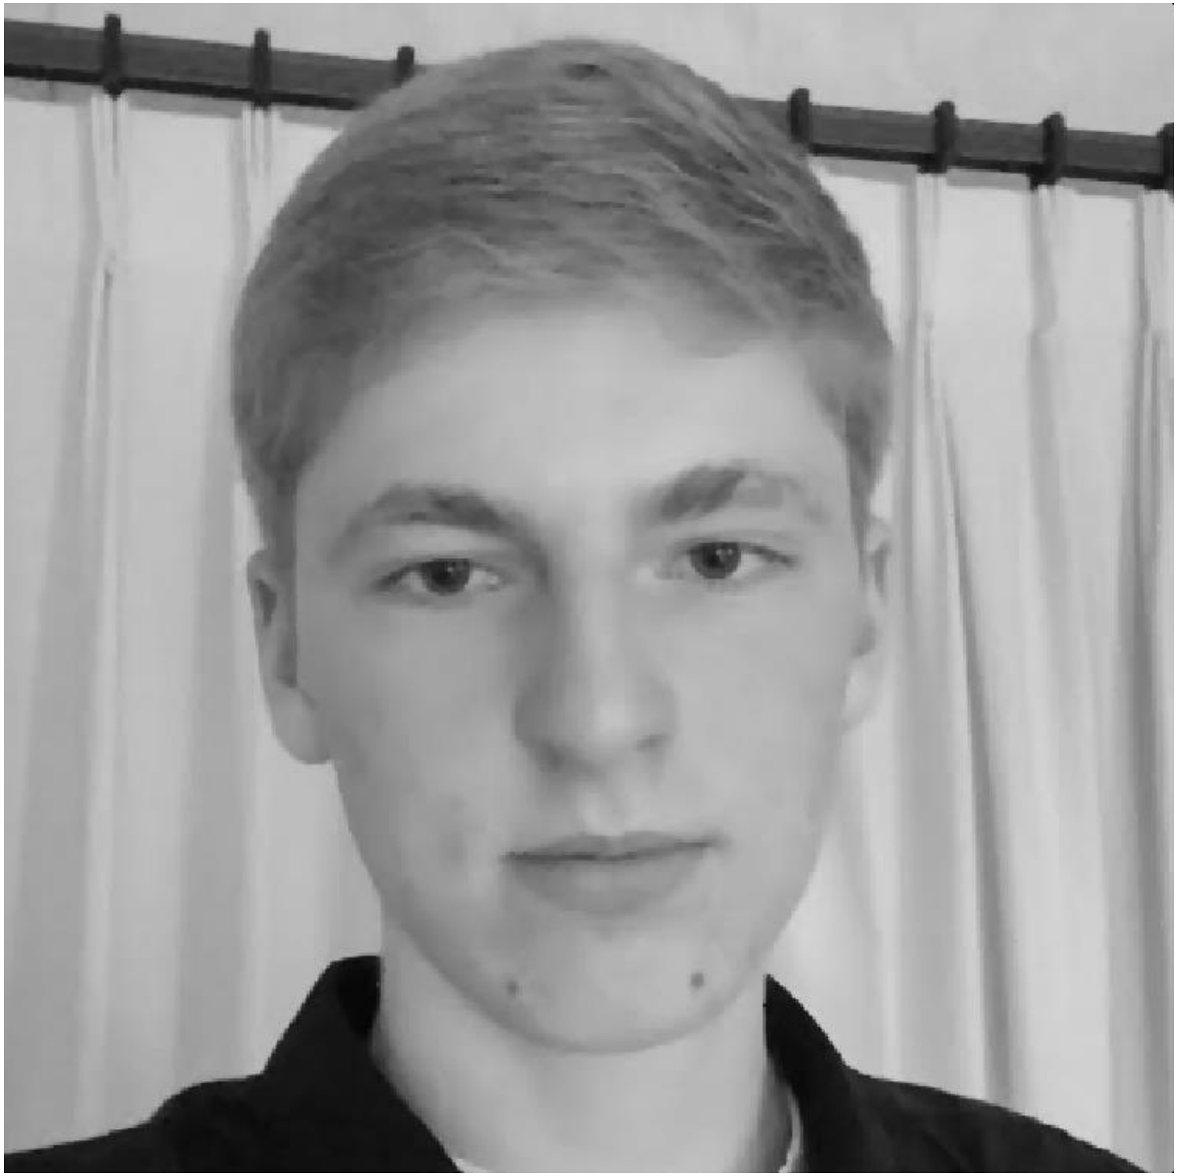

Very Median Jack – 11 x 11

Median filters

- Operates over a window by selecting the median intensity in the window.
- What advantage does a median filter have over a mean filter?
- Is a median filter a kind of convolution?

Slide by Steve Seitz

Think-Pair-Share

* = Convolution operator

a) _ = D * B
b) A = _ * _
c) F = D * _
d) _ = D * D

D

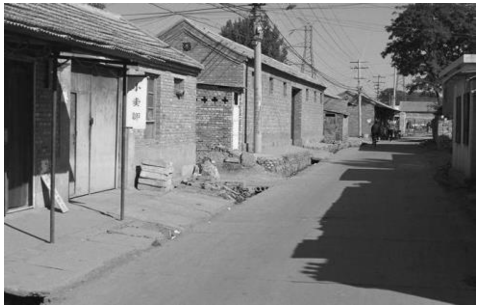

H

F

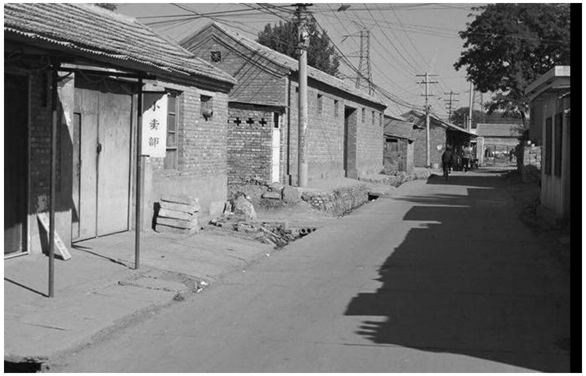

I

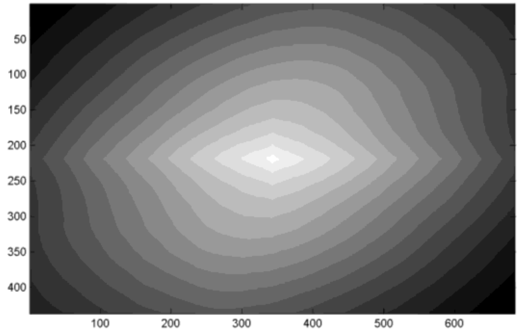

A

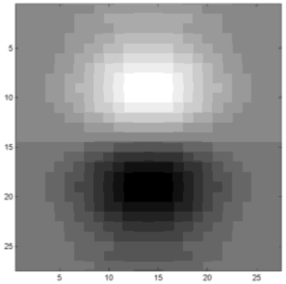

B

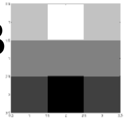

C

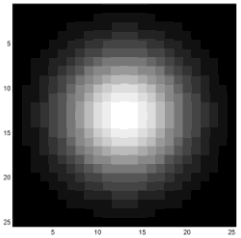

E

G

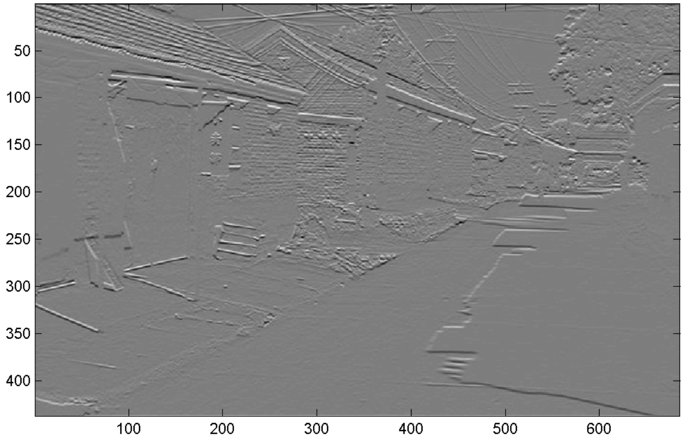

Slide: Hoiem

Salvador Dali, 1976

Today's Class

- Fourier transform and frequency domain
  - Frequency view of filtering
  - Hybrid images
  - Sampling
- Reminder: Textbook
  - Today's lecture covers material in 3.4

Slide: Hoiem

Hays

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

Gaussian 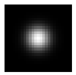

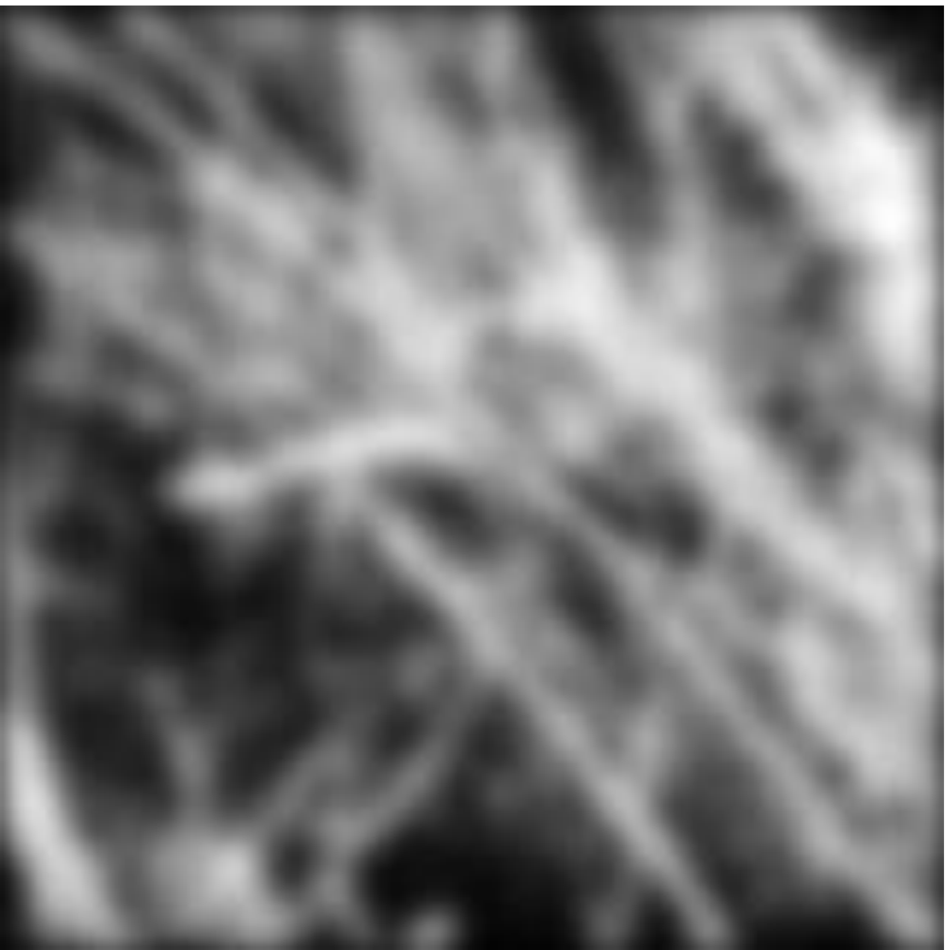

Box filter 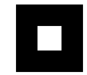

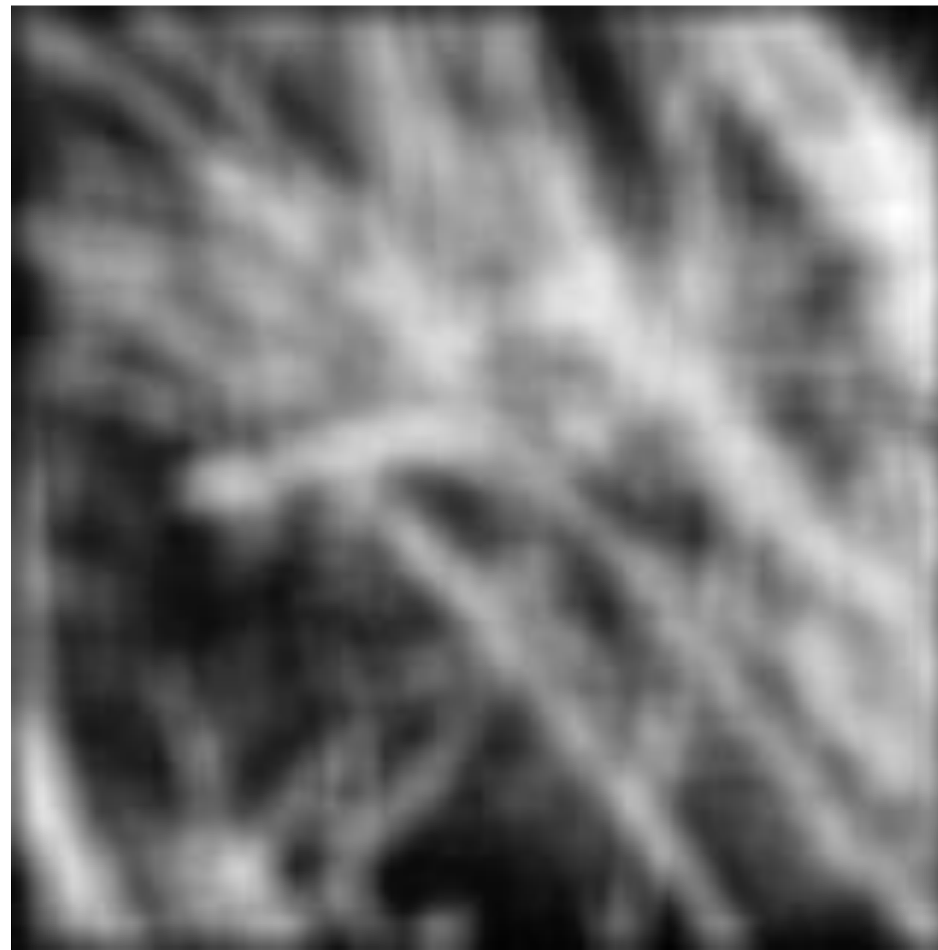

Why does a lower resolution image still make sense to us? What information do we lose?

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

Hybrid Images

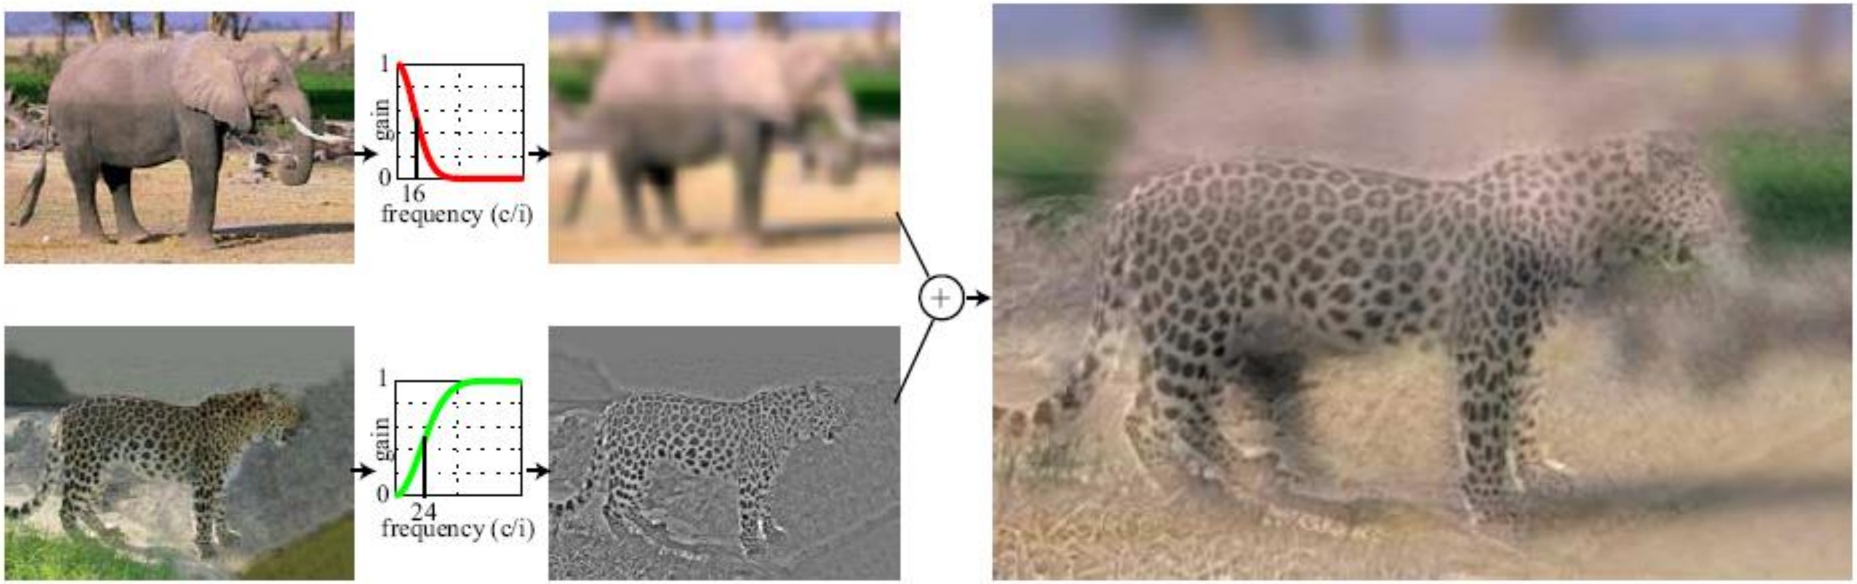

- A. Oliva, A. Torralba, P.G. Schyns, "Hybrid Images," SIGGRAPH 2006

Why do we get different, distance-dependent interpretations of hybrid images?

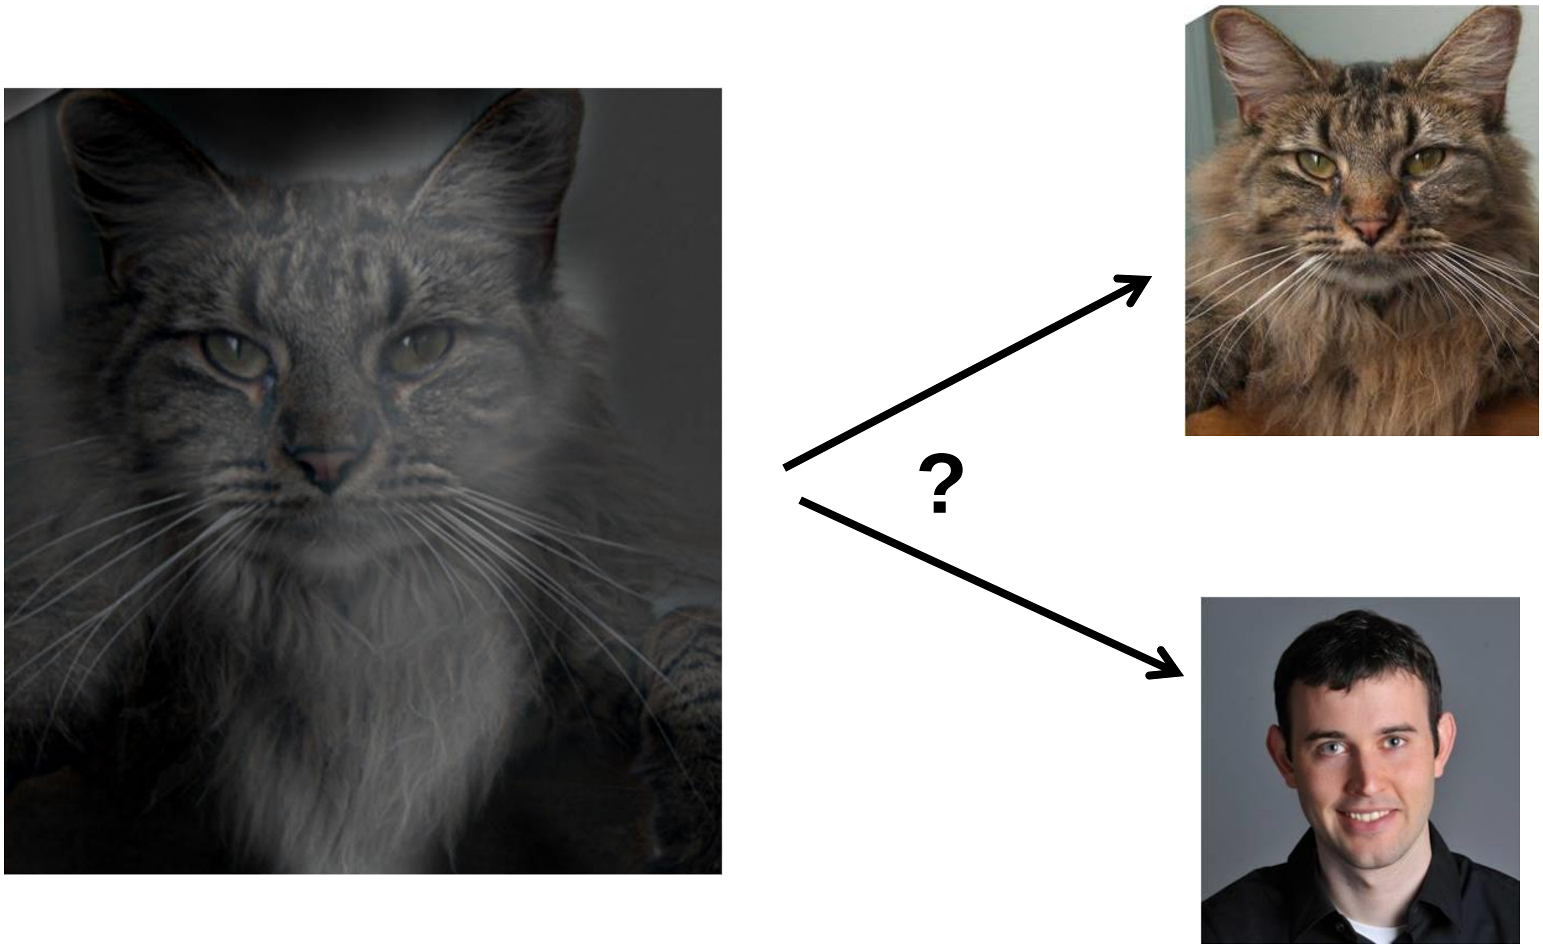

Slide: Hoiem

Jean Baptiste Joseph Fourier (1768-1830)

A bold idea (1807):

Any *univariate function can rewritten as a weighted sum 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

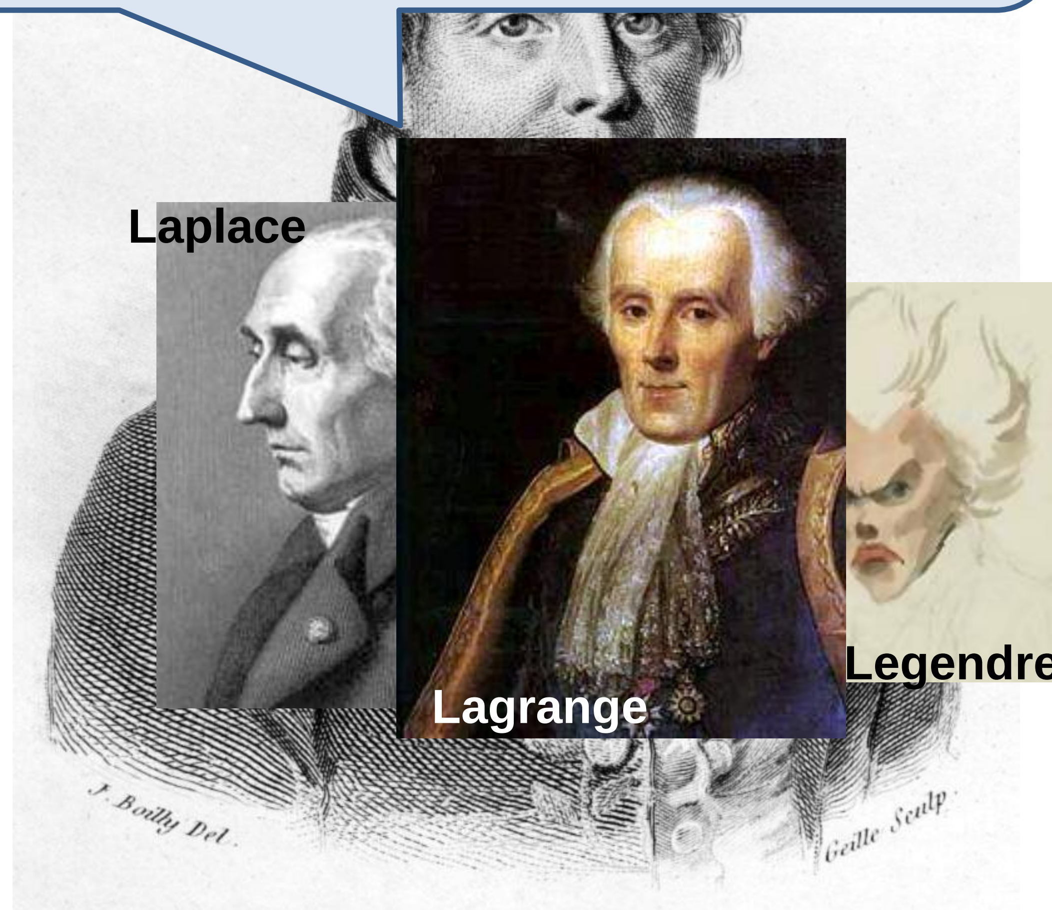

Hays

A sum of sines and cosines

Our building block:

$$A\sin(\omega x) + B\cos(\omega x)$$

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

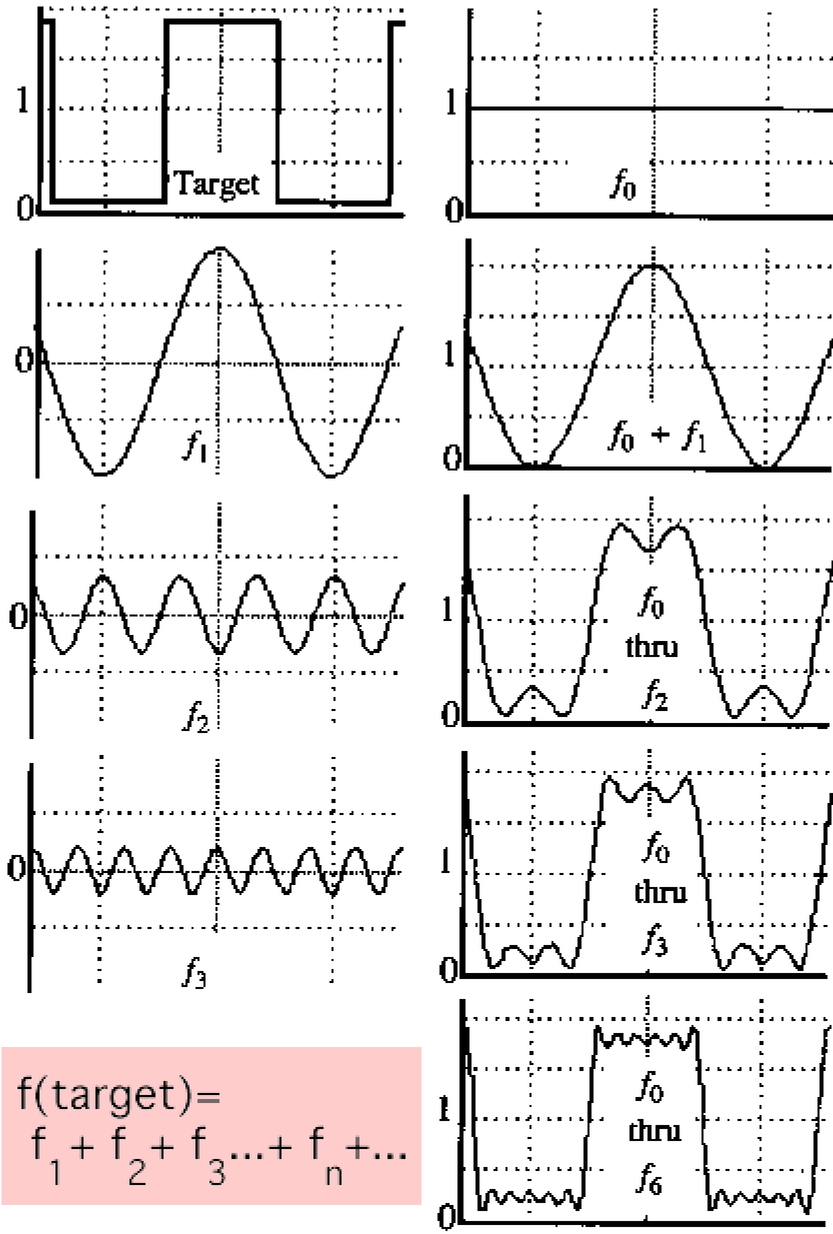

Hays

Frequency Spectra

- Example : $g(t) = \sin(2\pi f\, t) + (1/3)\sin(2\pi(3f)\, t)$

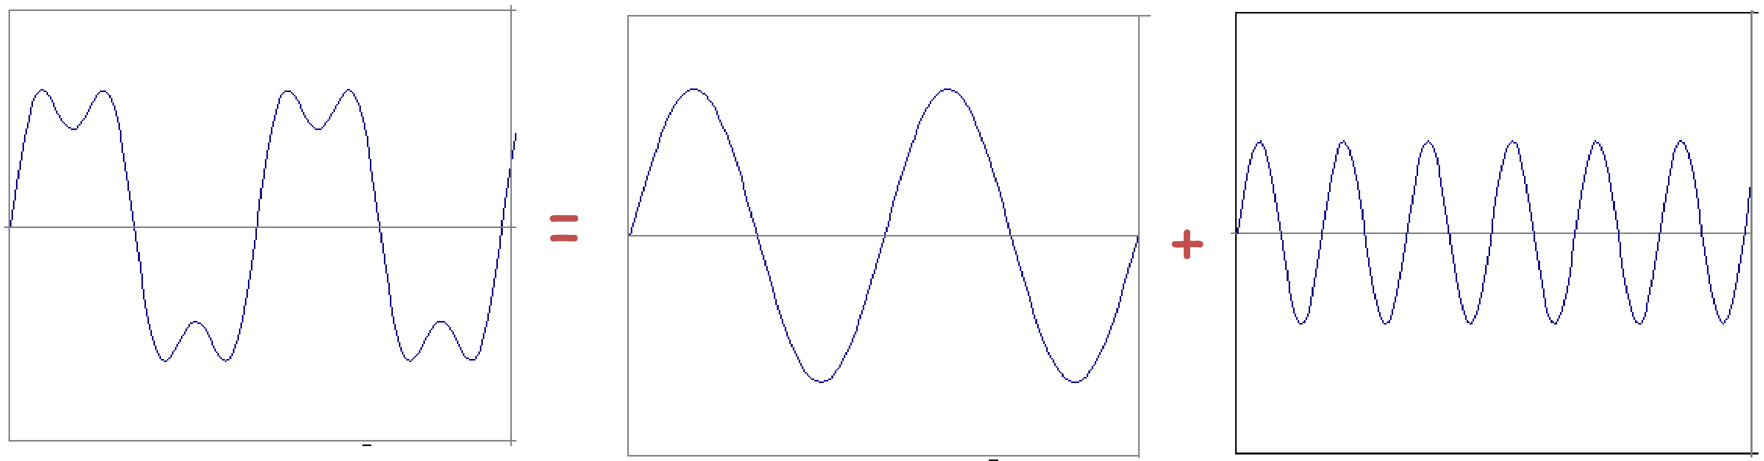

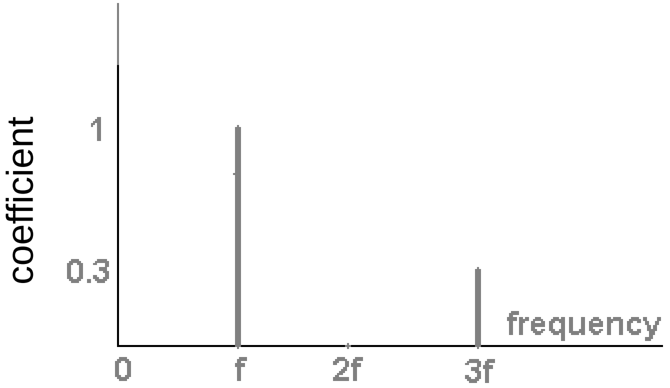

Slides: Efros

Frequency Spectra

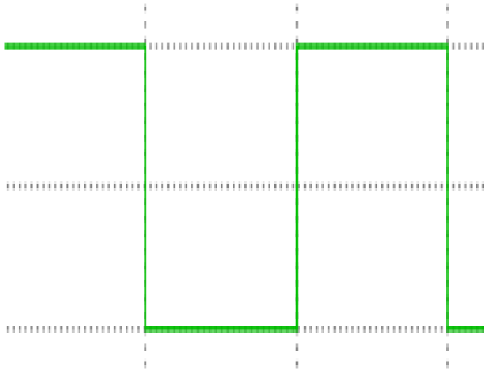

Frequency Spectra

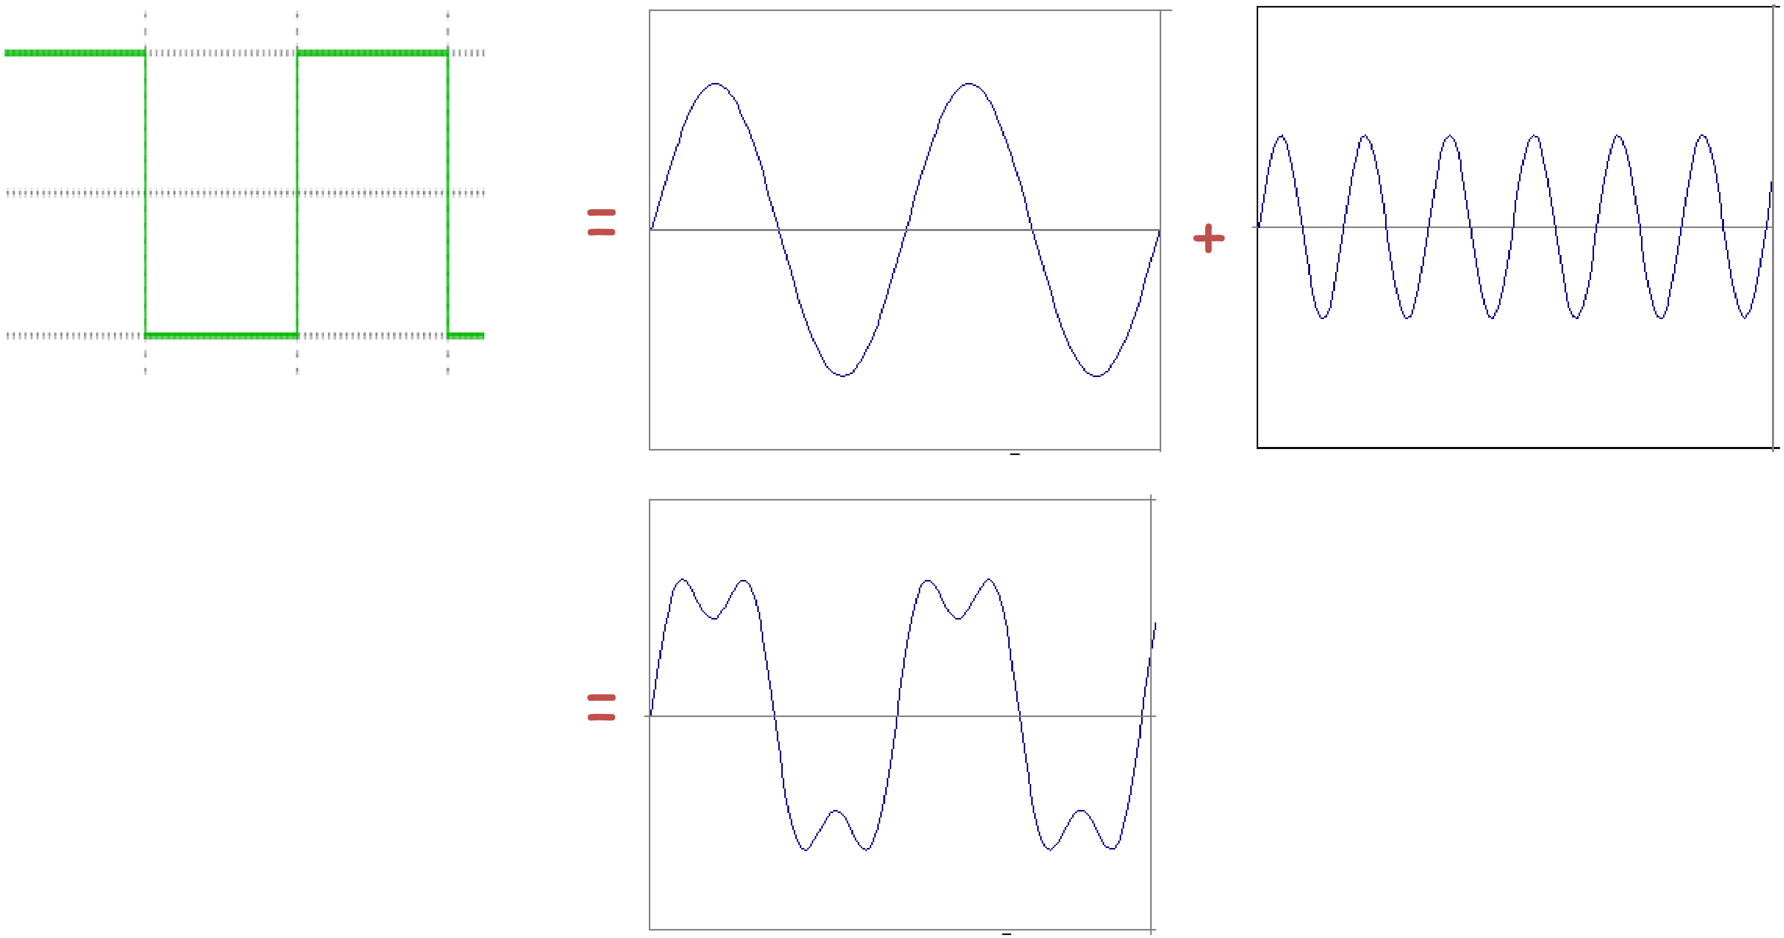

Frequency Spectra

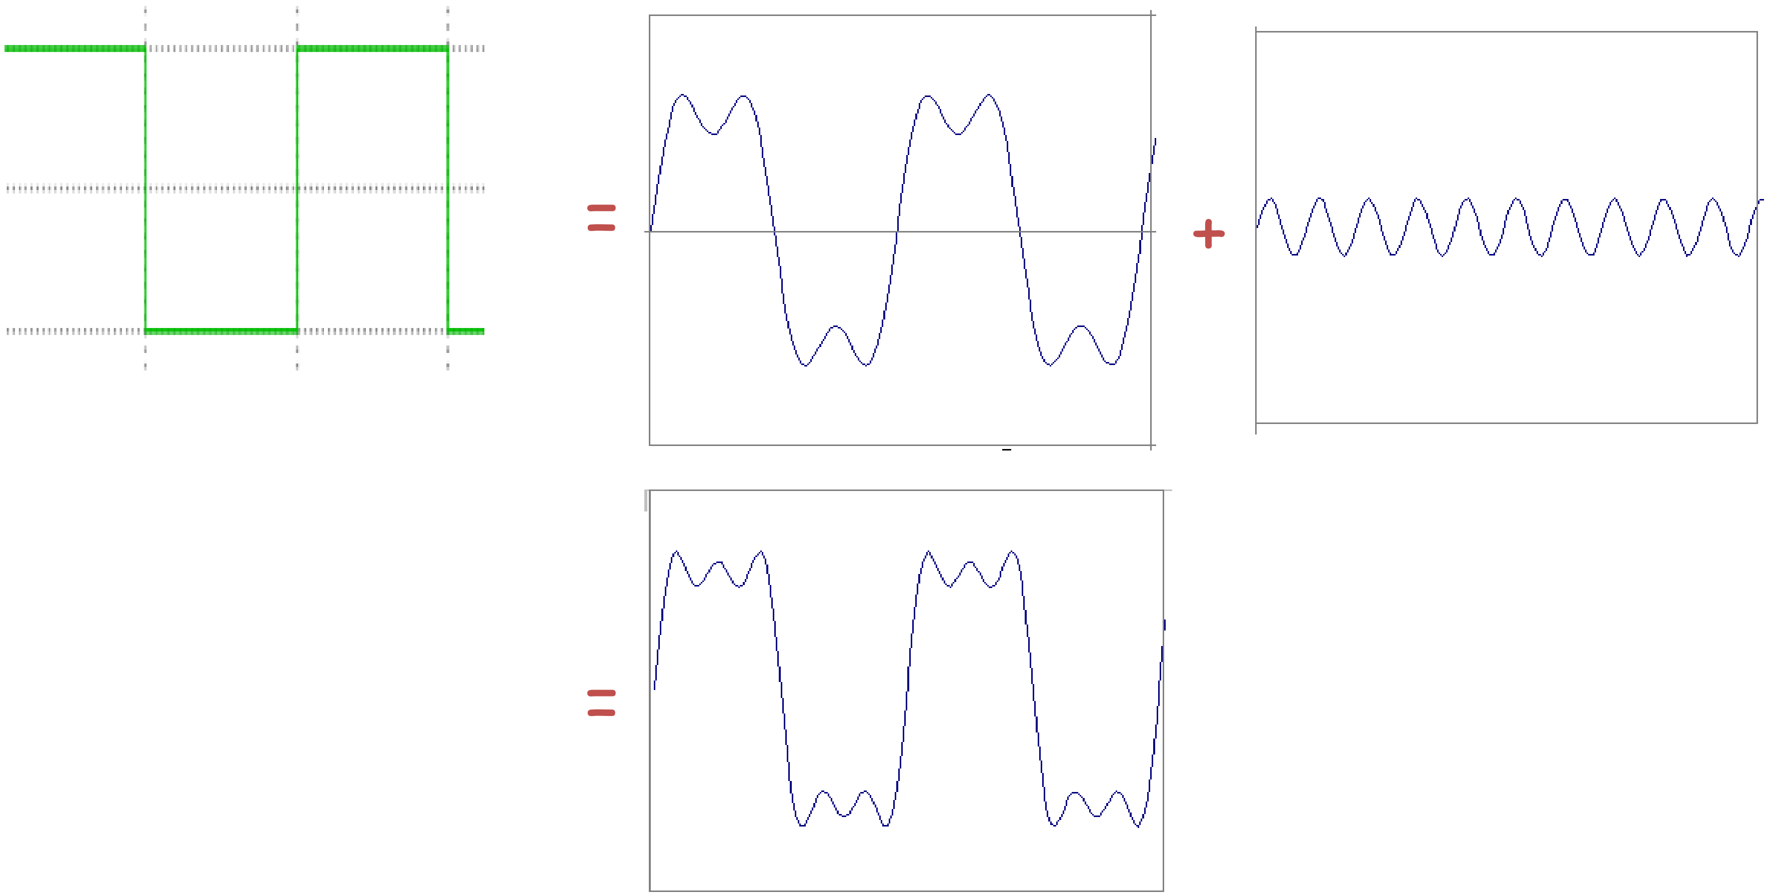

Frequency Spectra

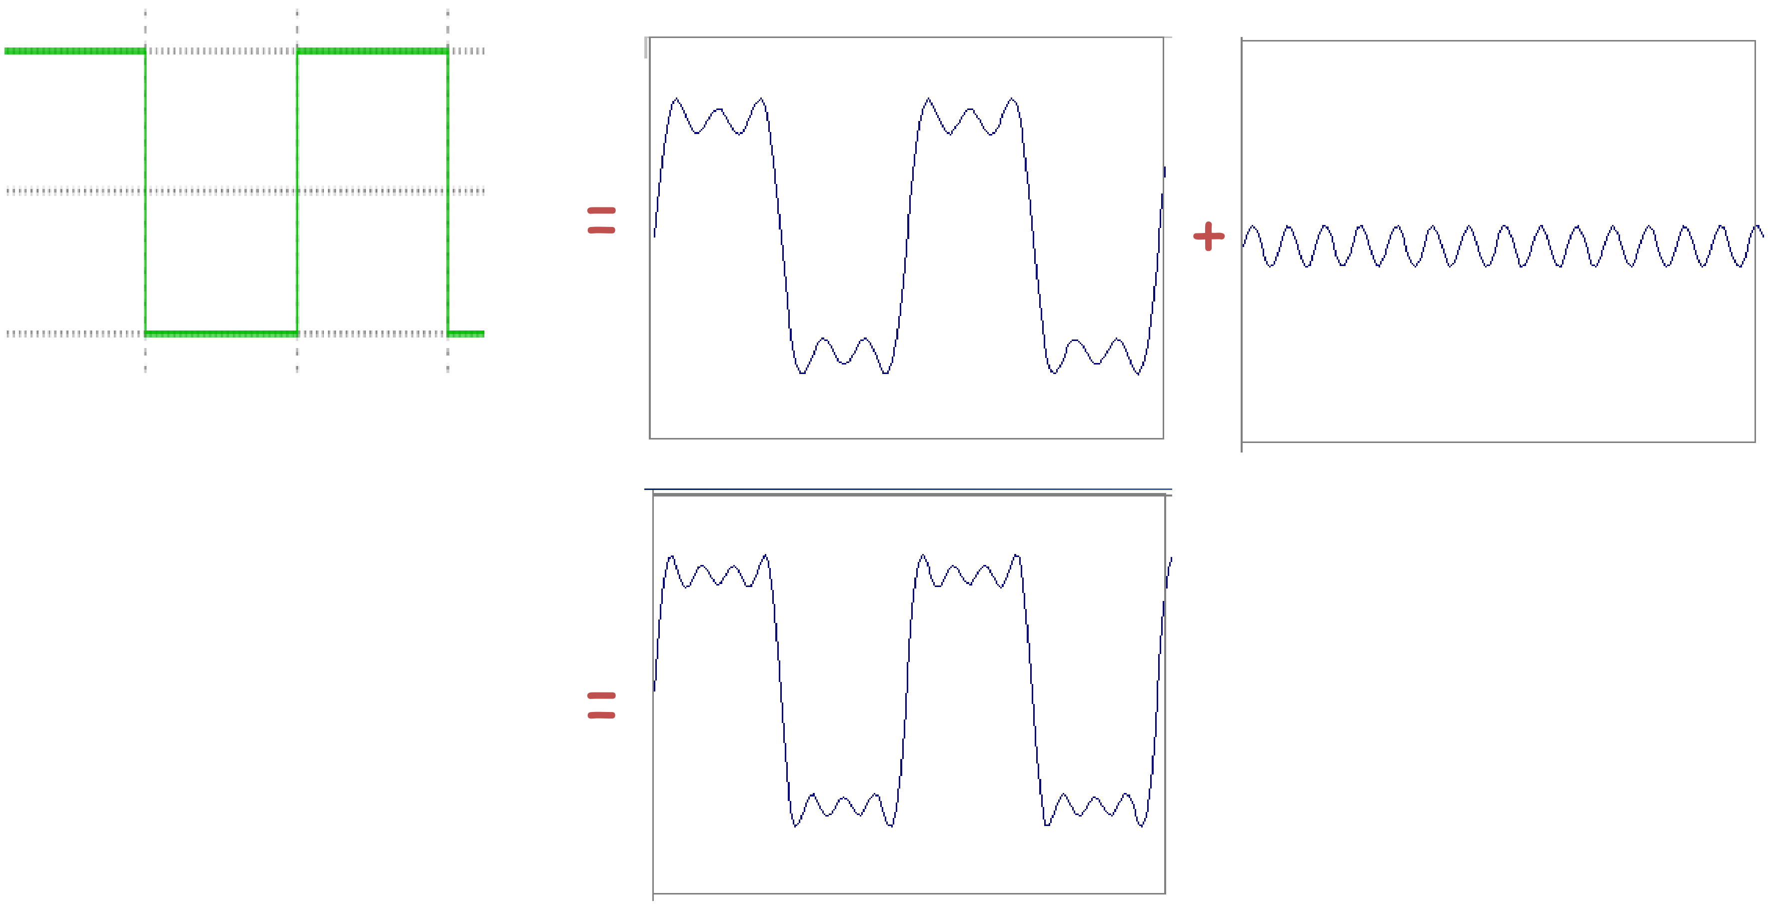

Frequency Spectra

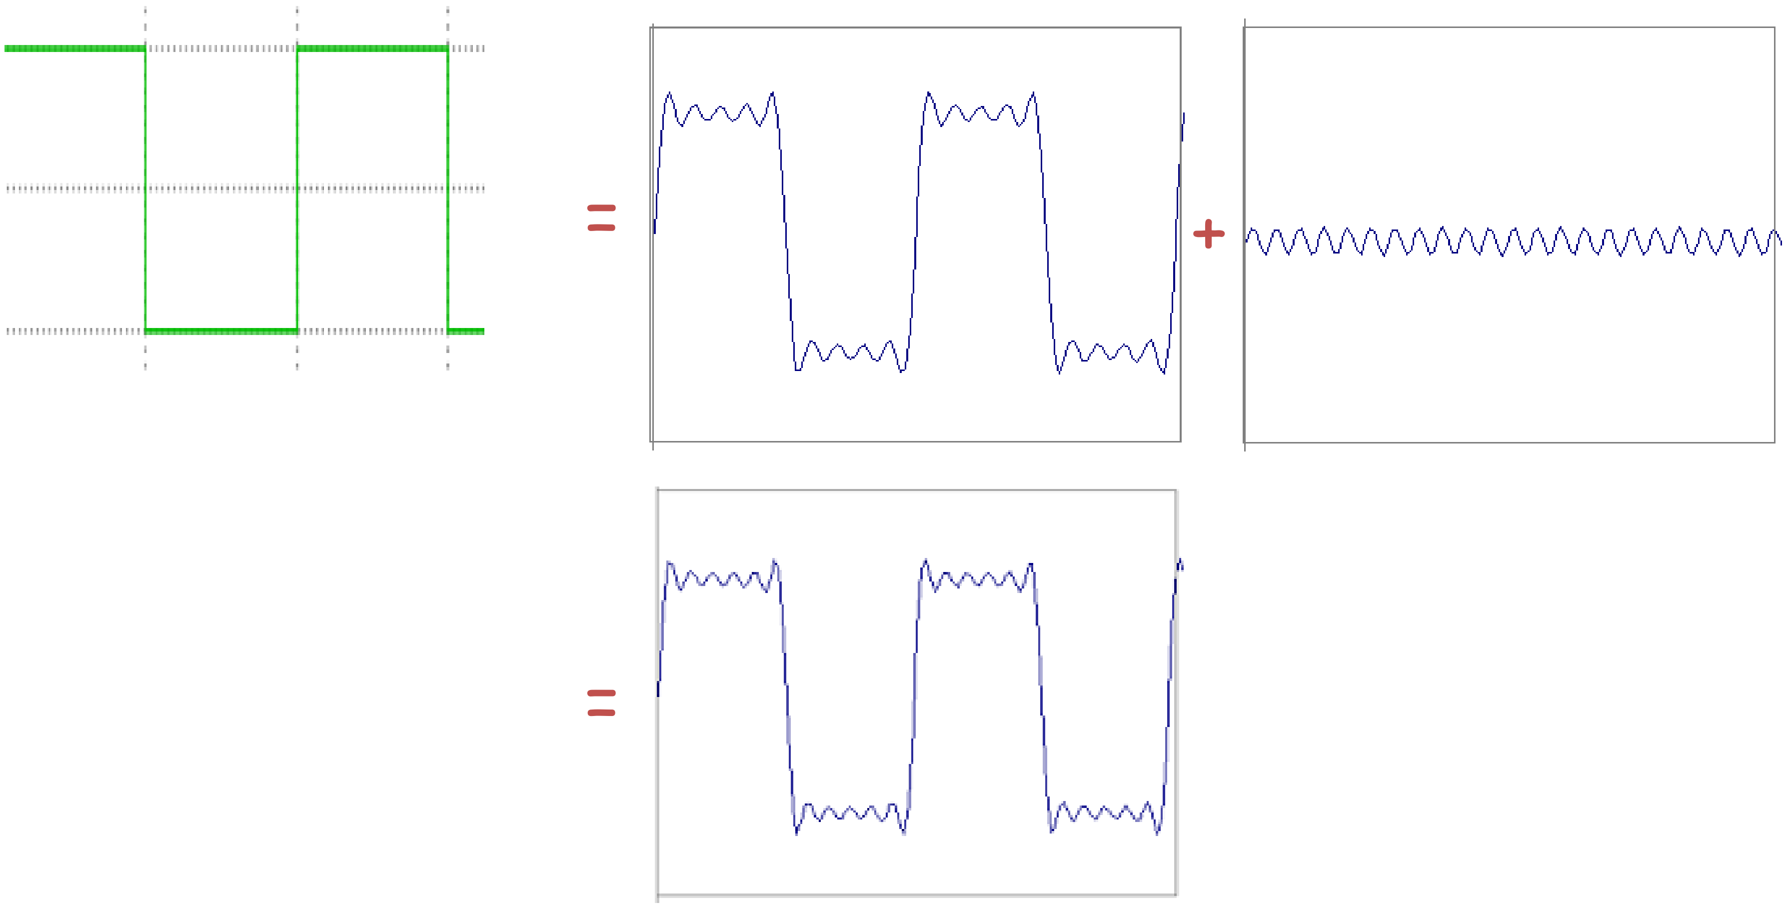

Frequency Spectra

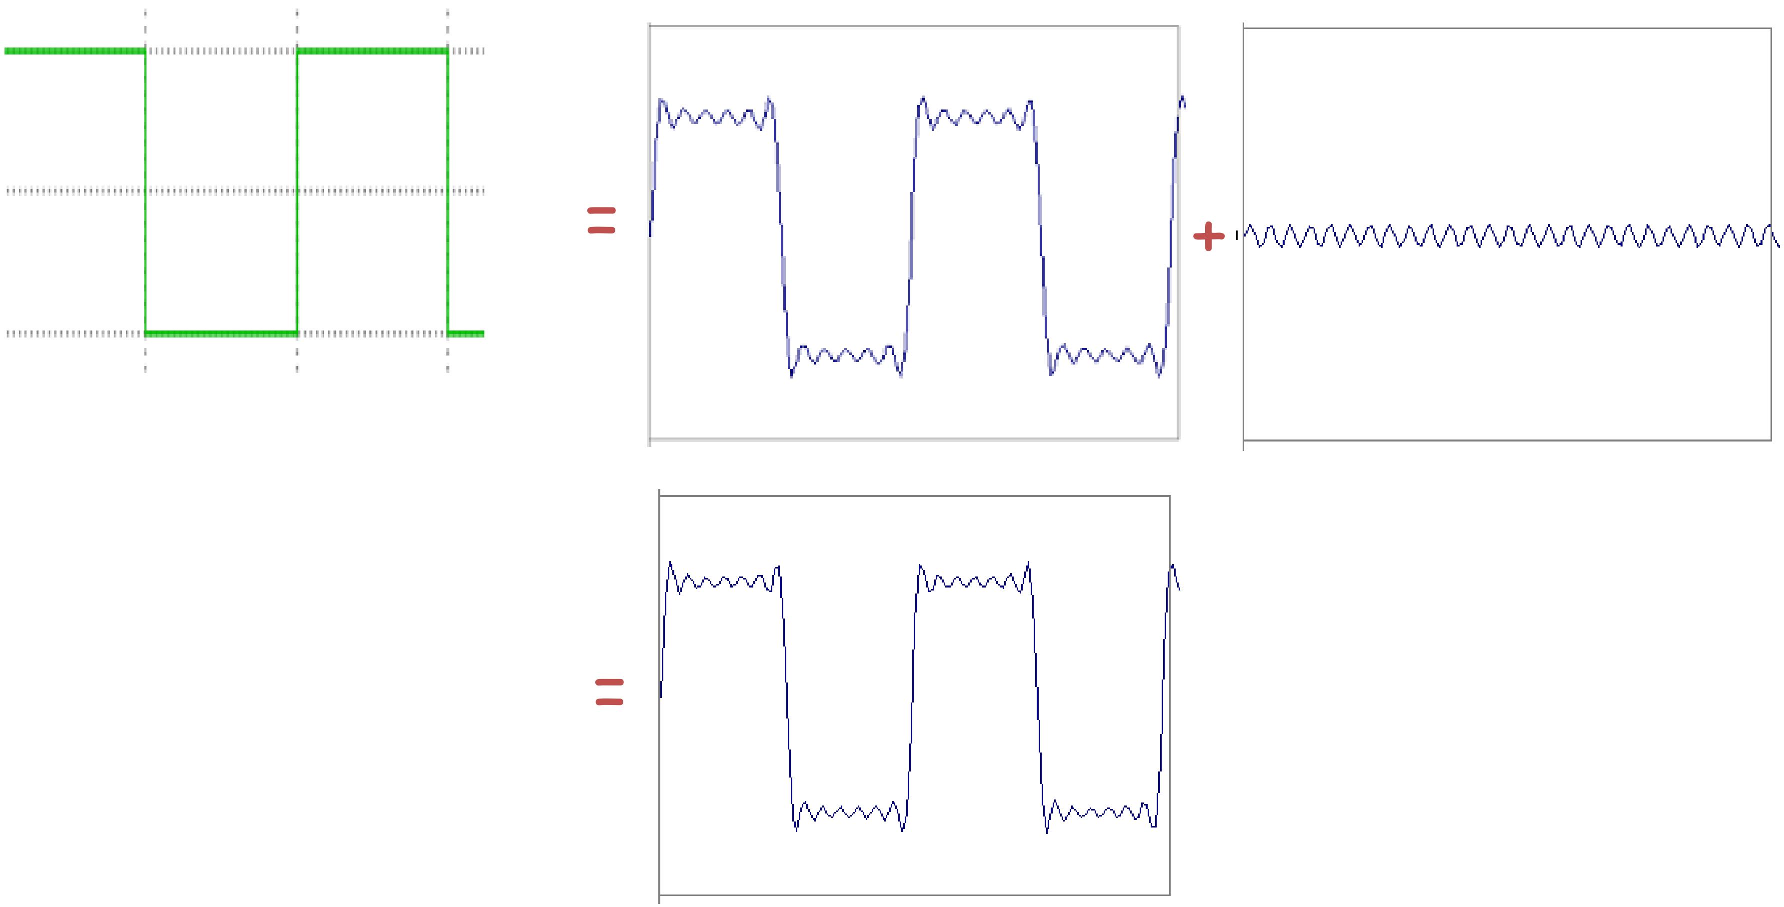

Frequency Spectra

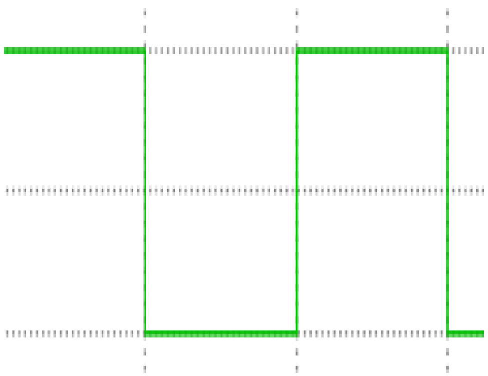

$$= \quad A\sum_{k=1}^{\infty}\frac{1}{k}\sin(2\pi kt)$$

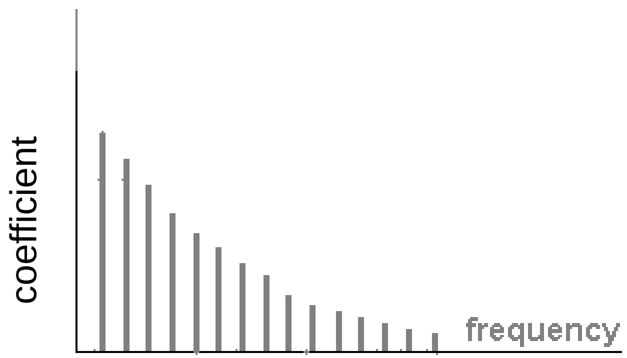

Example: Music

- We think of music in terms of frequencies at different magnitudes

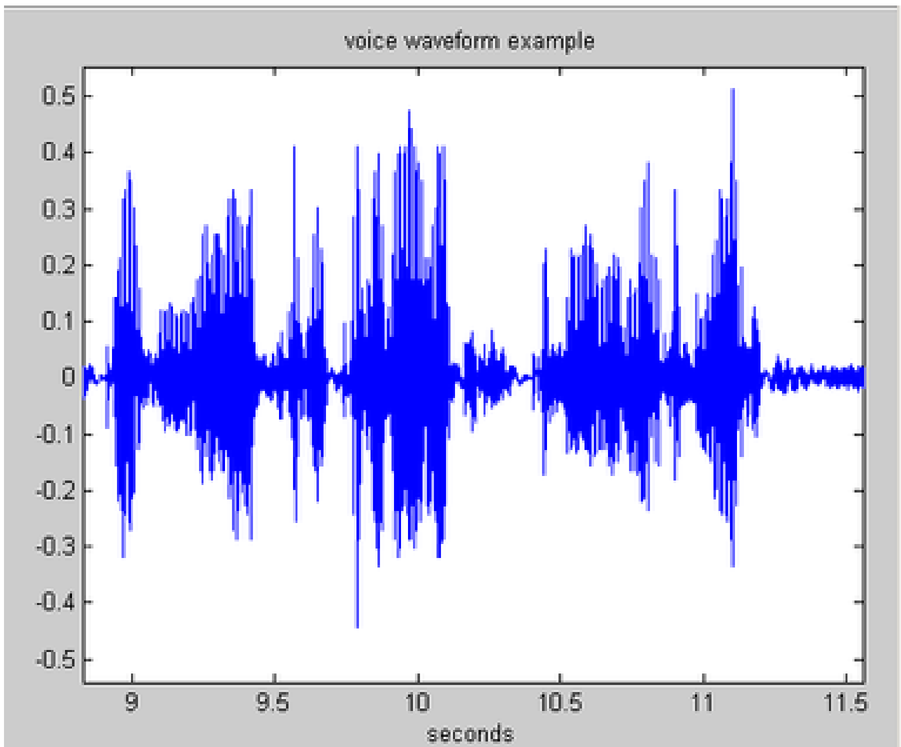

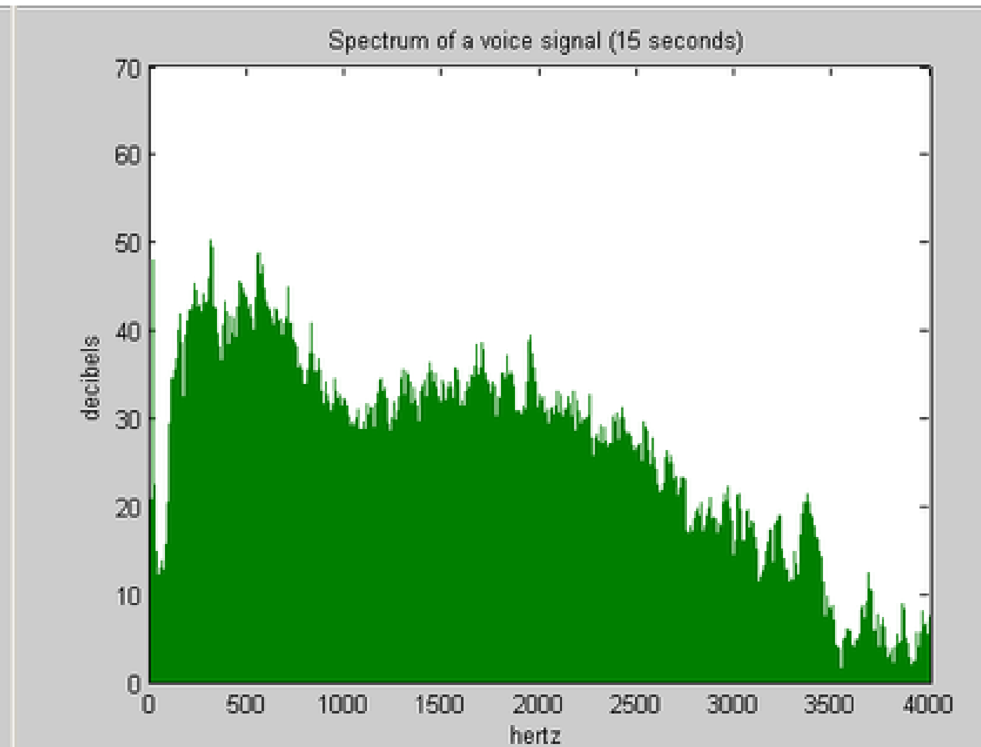

Slide: Hoiem

Evan Wallace demo

- Made for CS123
- 1D example
- Forbes 30 under 30
  - Figma (collaborative design tools)

- http://madebyevan.com/dft/

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?!?!
: (

Other signals

- We can also think of all kinds of other signals the same way

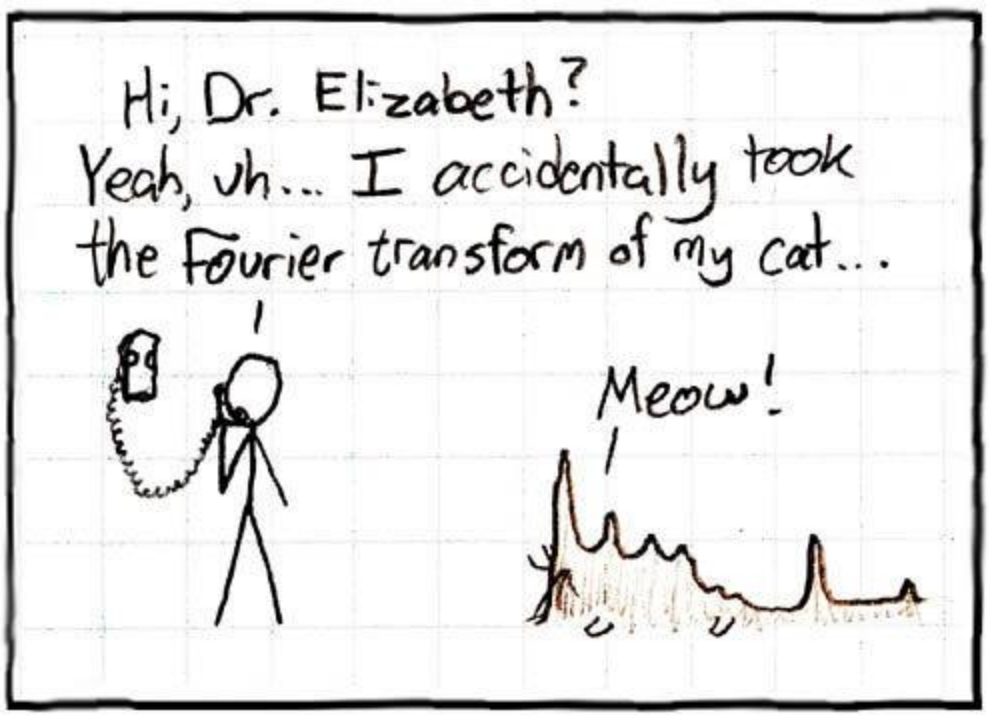

xkcd.com

Fourier analysis in images

Intensity images

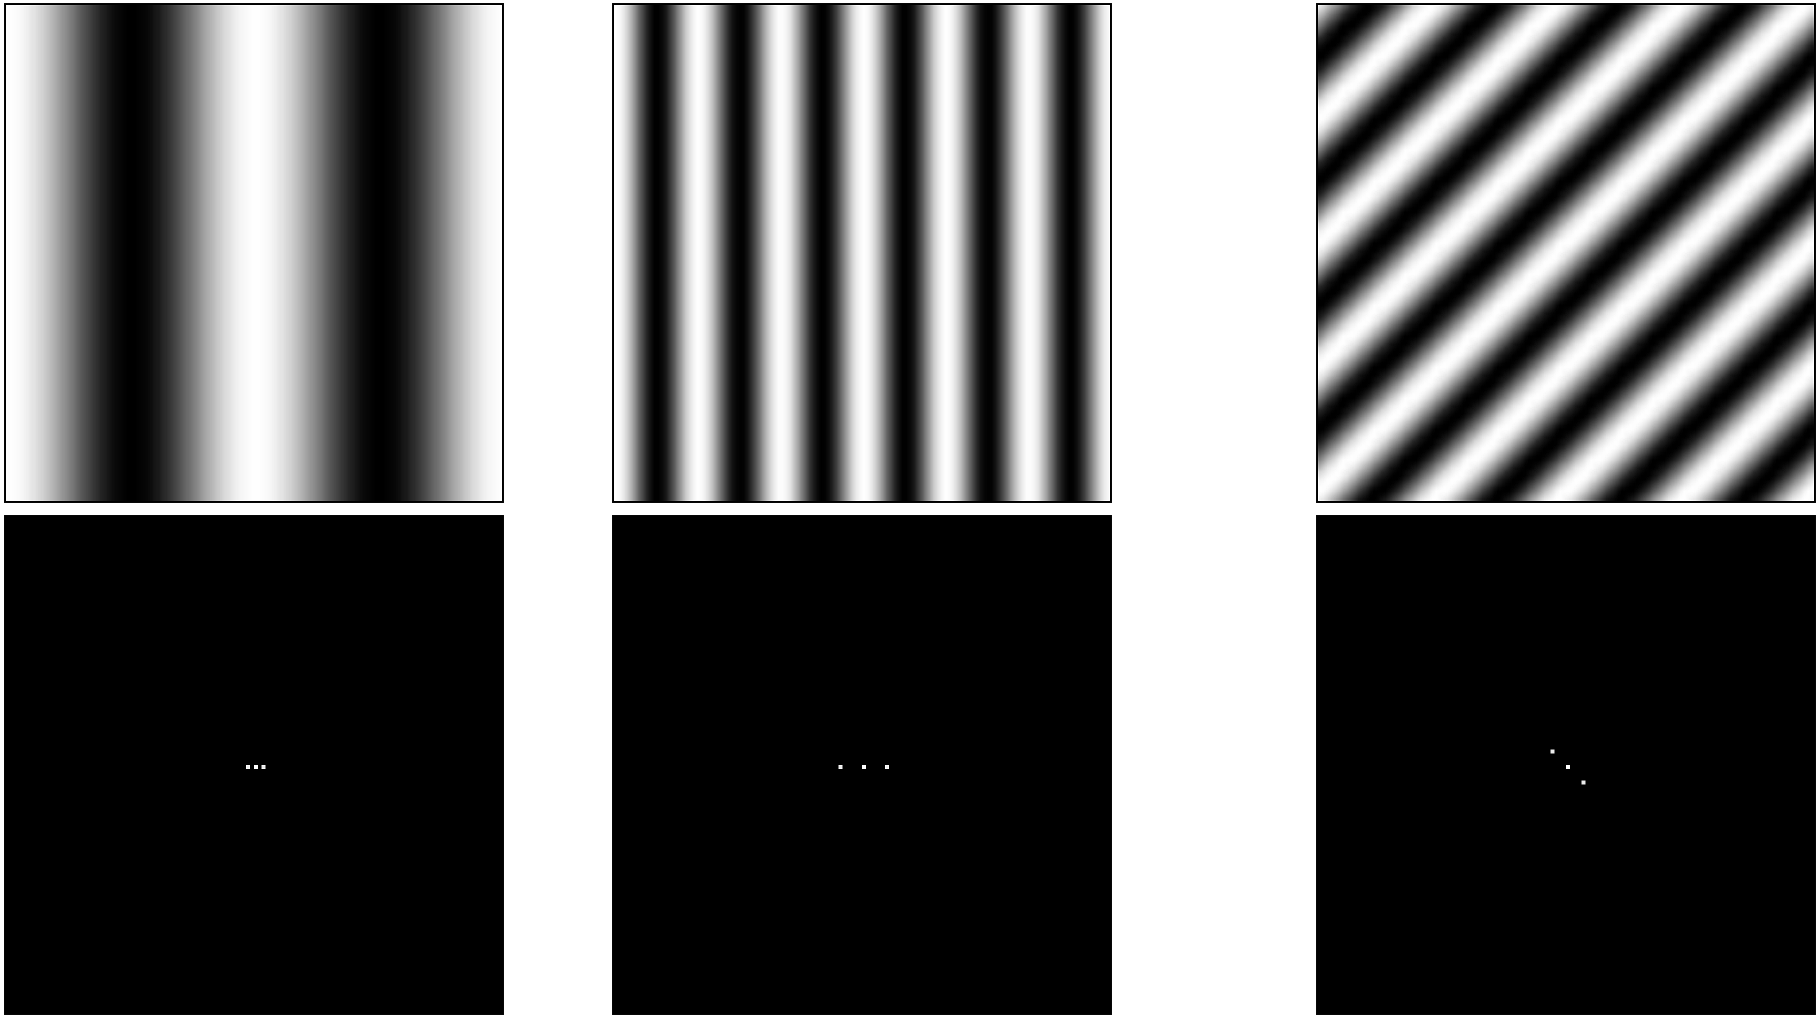

Fourier decomposition images

http://sharp.bu.edu/~slehar/fourier/fourier.html#filtering

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

Euler's formula

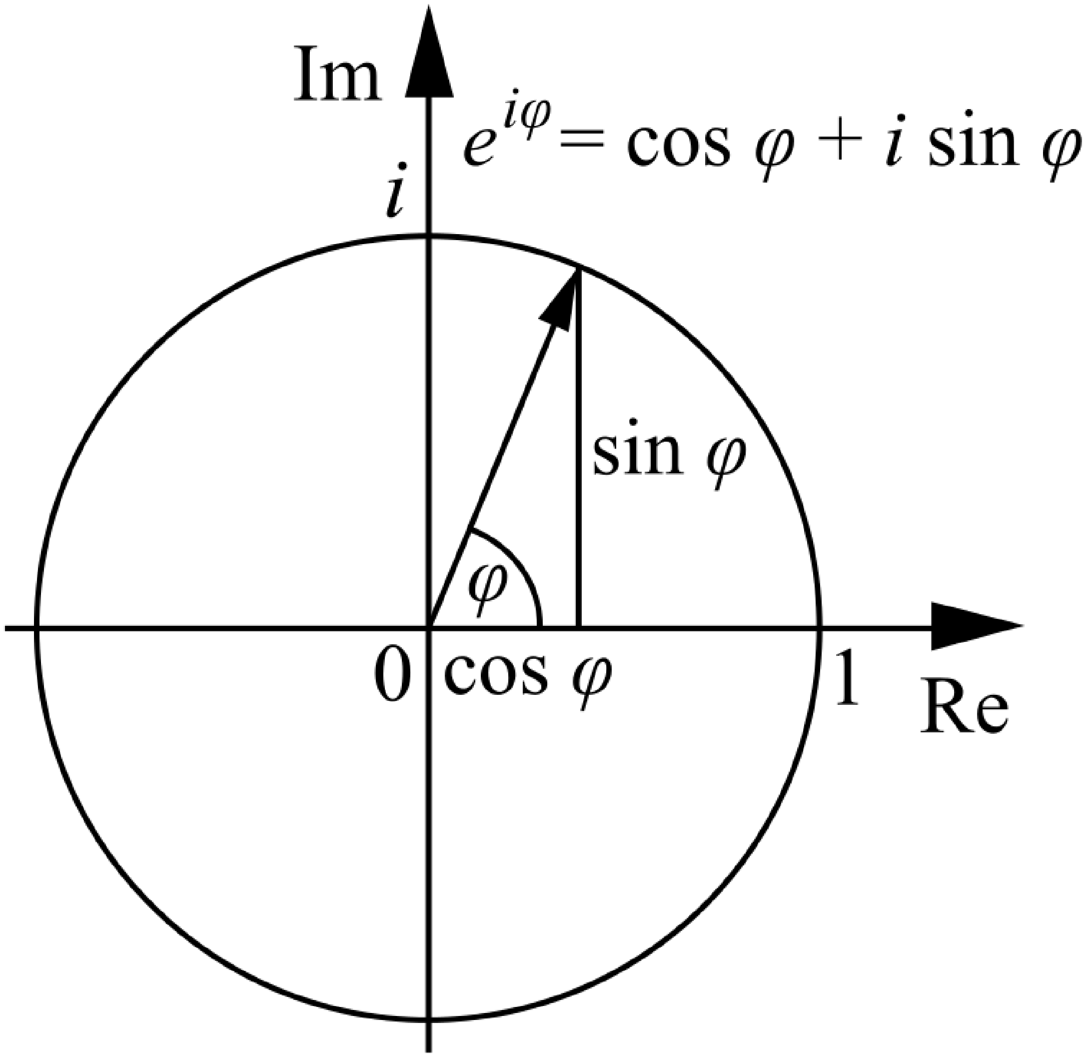

Wikipedia

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

Fourier Bases

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

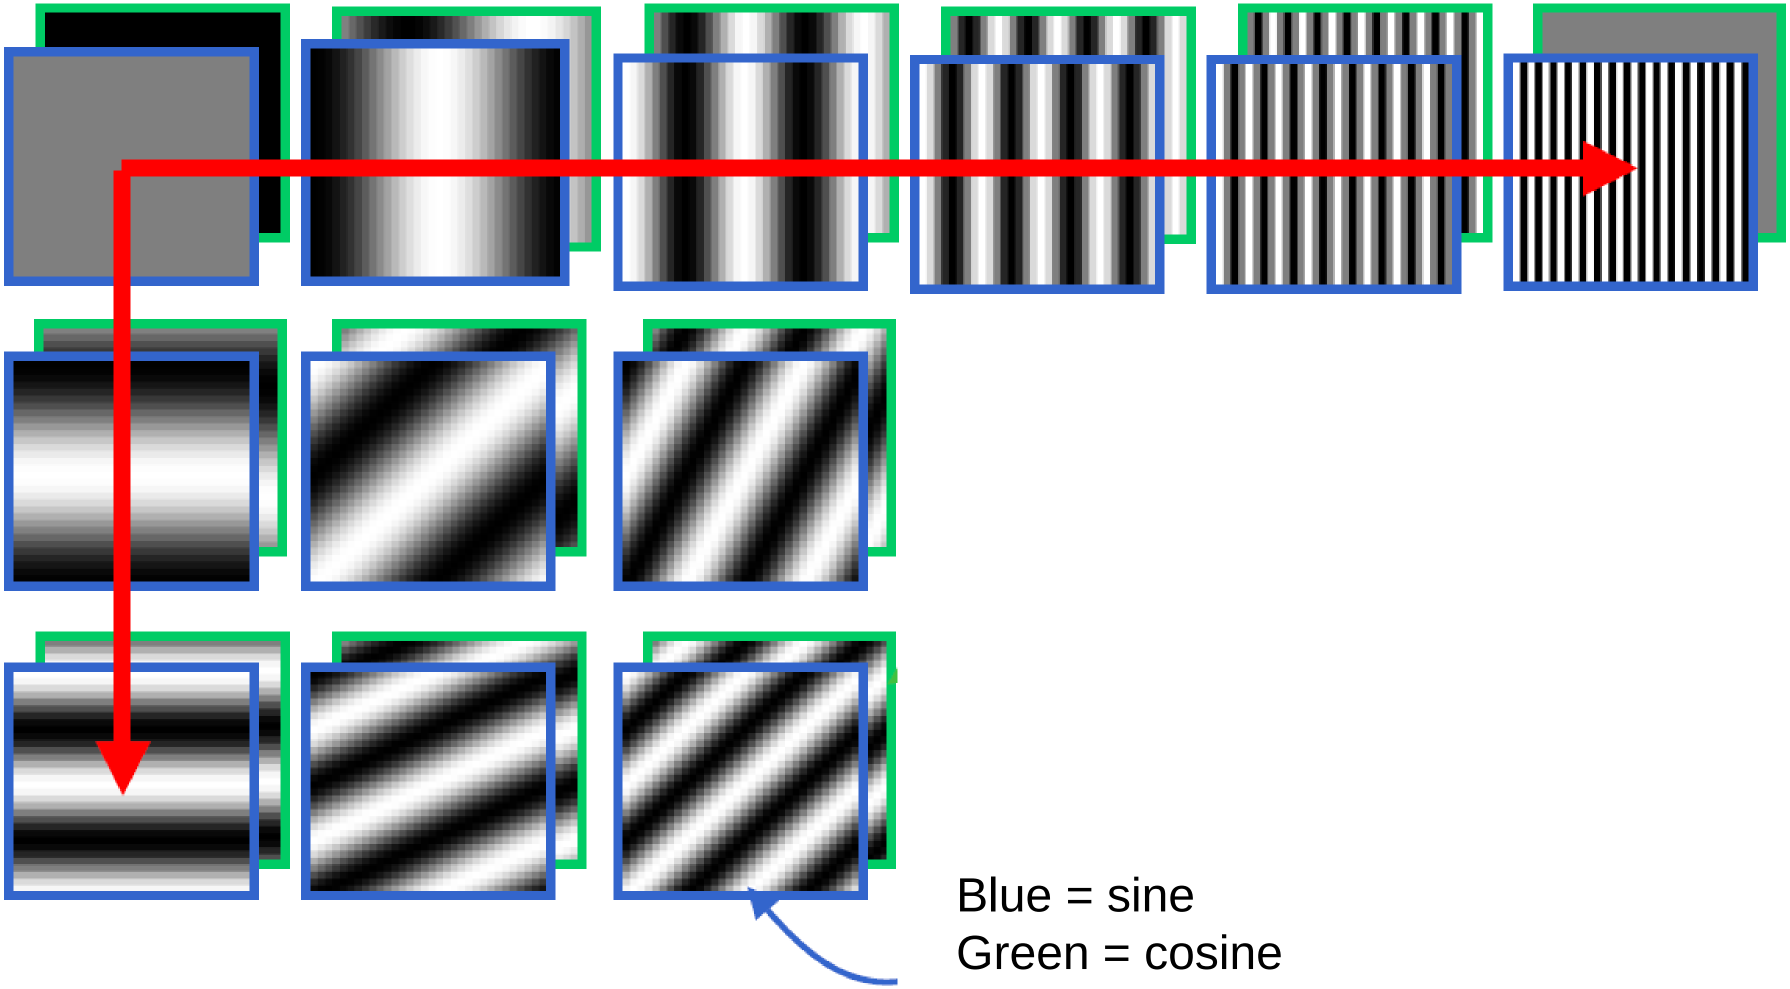

This change of basis is the Fourier Transform

Hays

Basis reconstruction


Full image


First 1 basis fn

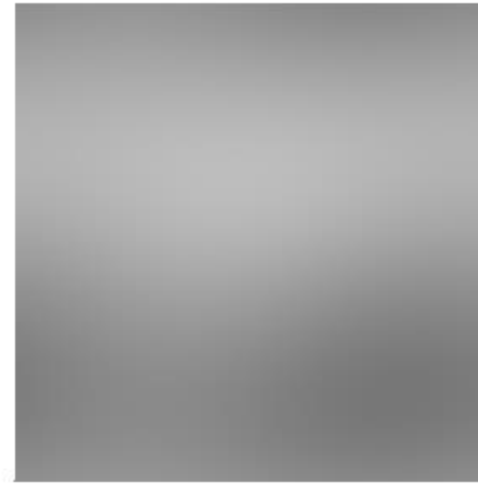
First 4 basis fns

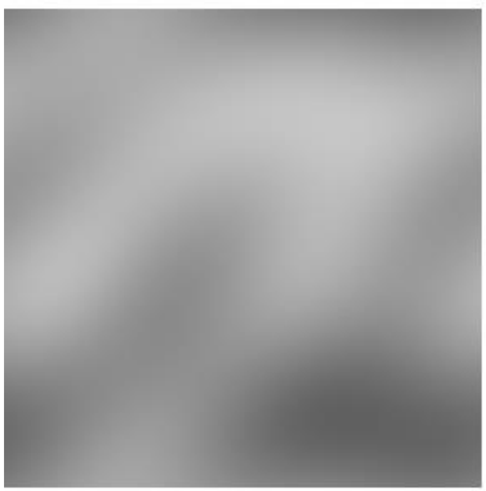
First 9 basis fns

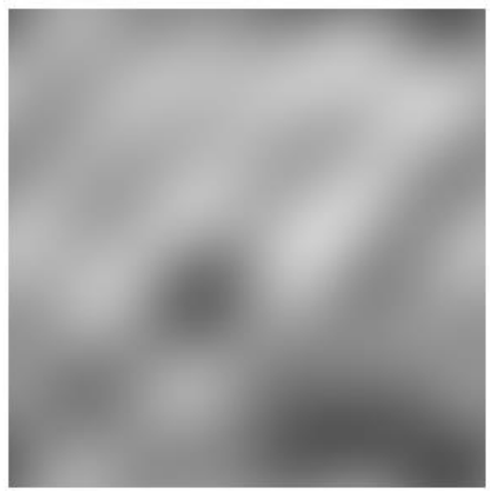
First 16 basis fns

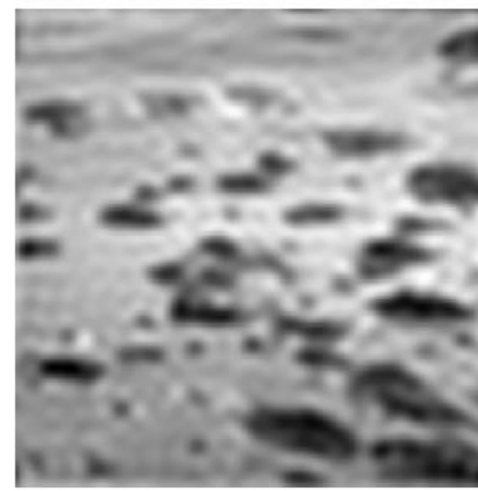
First 400 basis fns

Danny Alexander

Man-made Scene

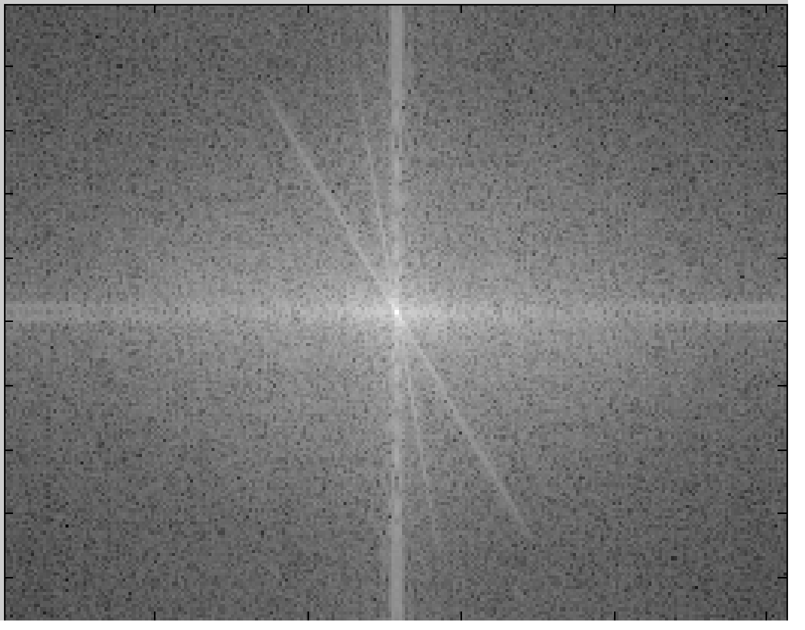

What does it mean to be at pixel x,y?
What does it mean to be more or less bright in the Fourier decomposition image?

Now we can edit frequencies!

Low and High Pass filtering

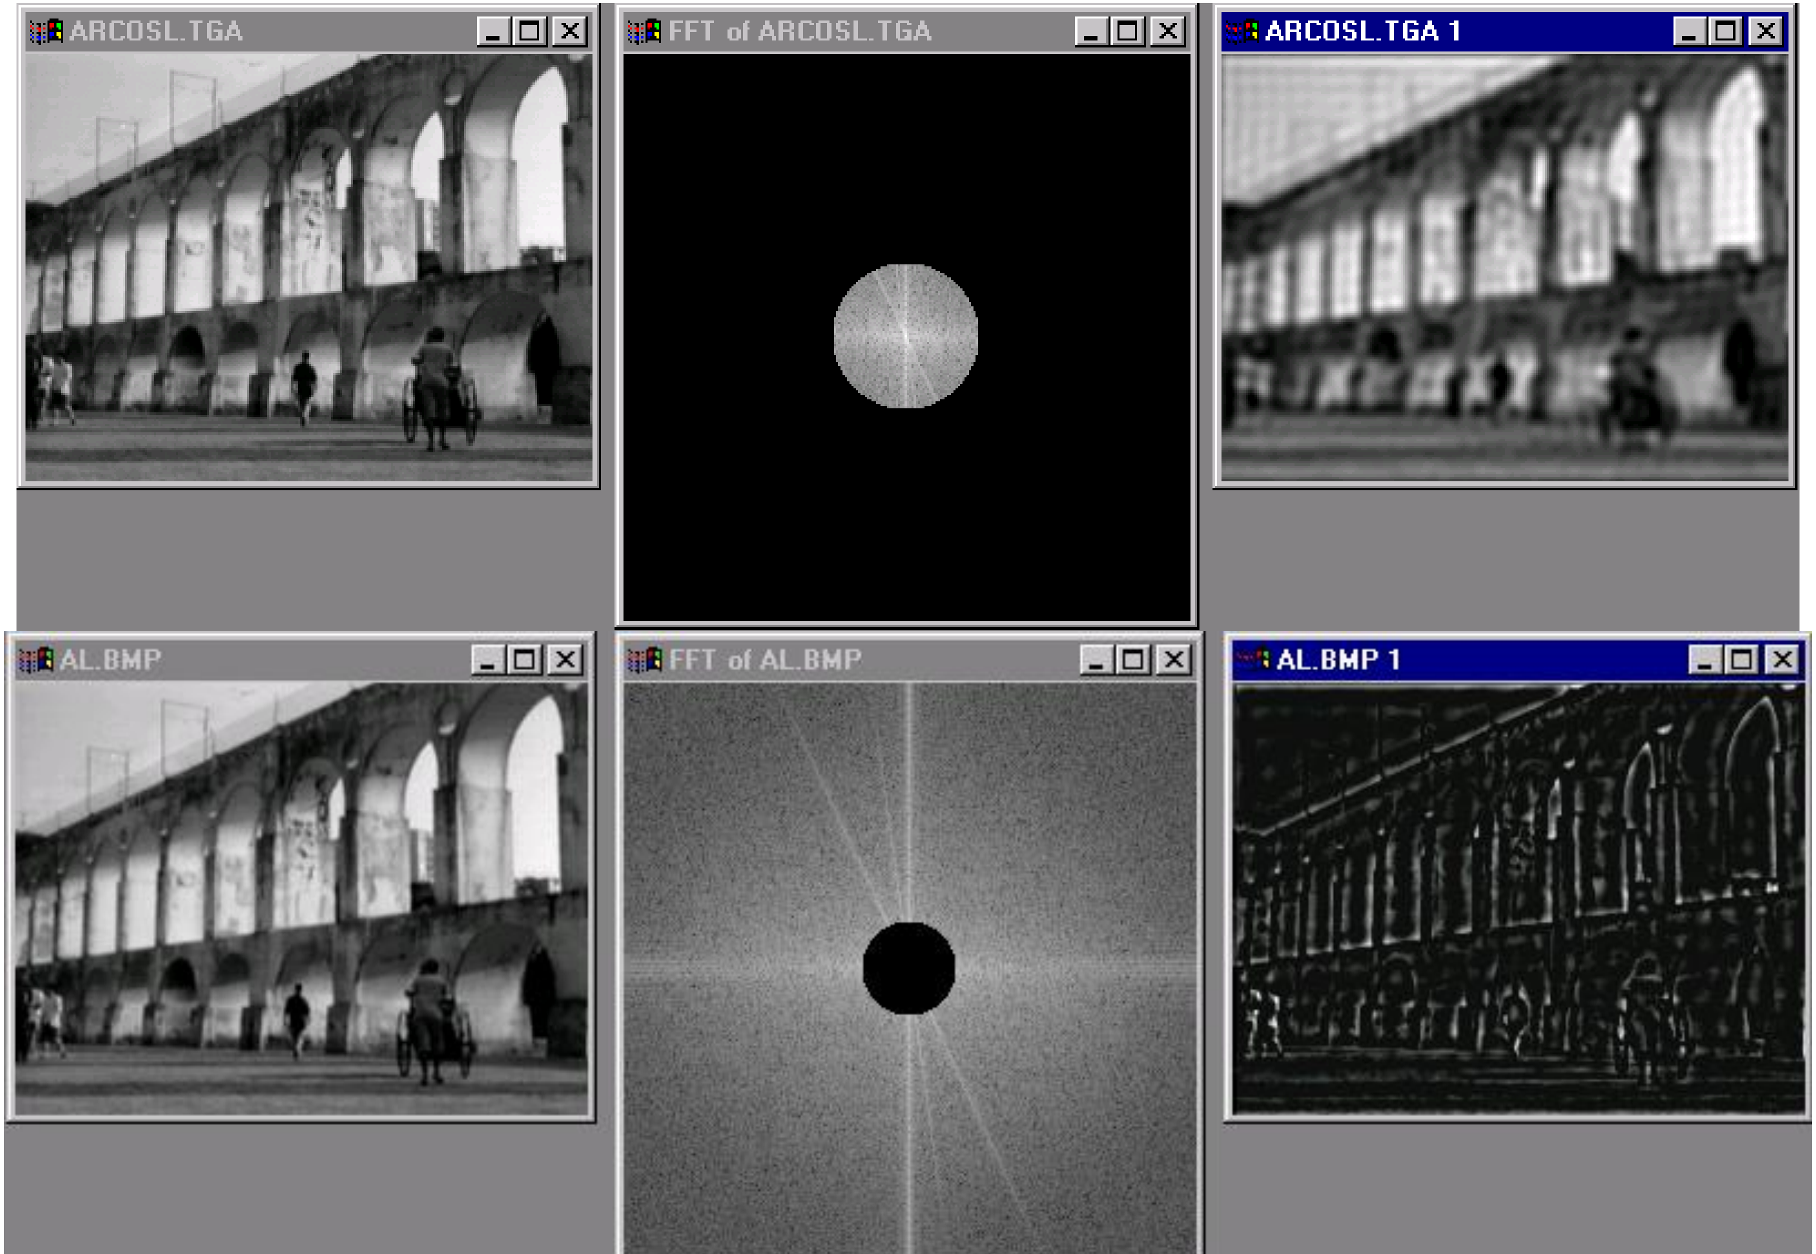

Removing frequency bands

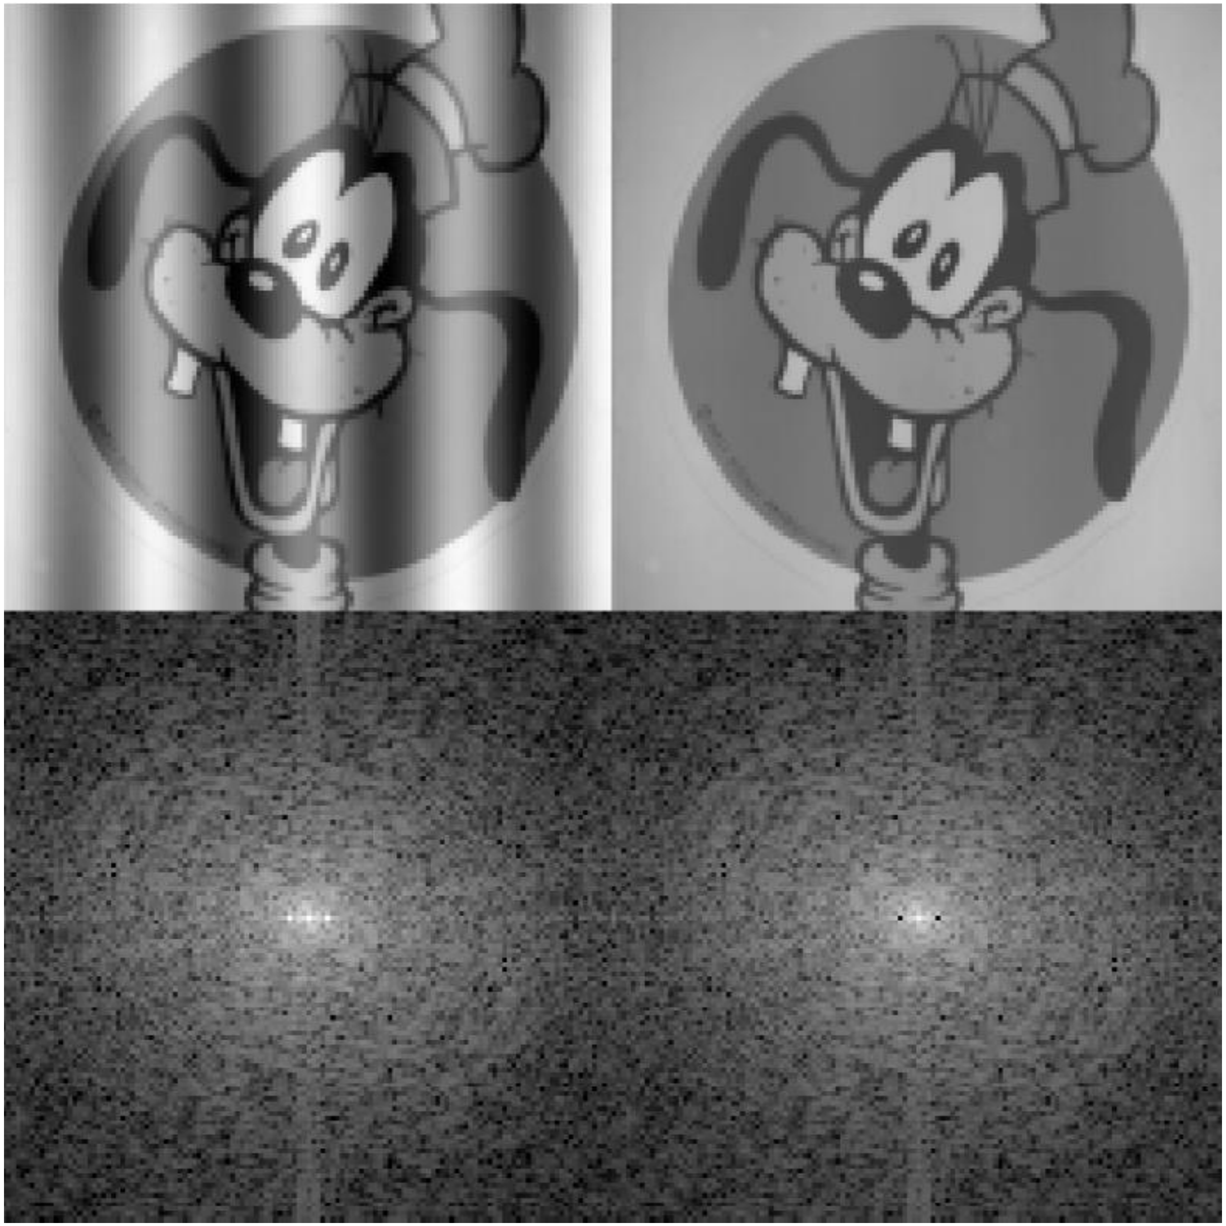

Brayer

High pass filtering + orientation

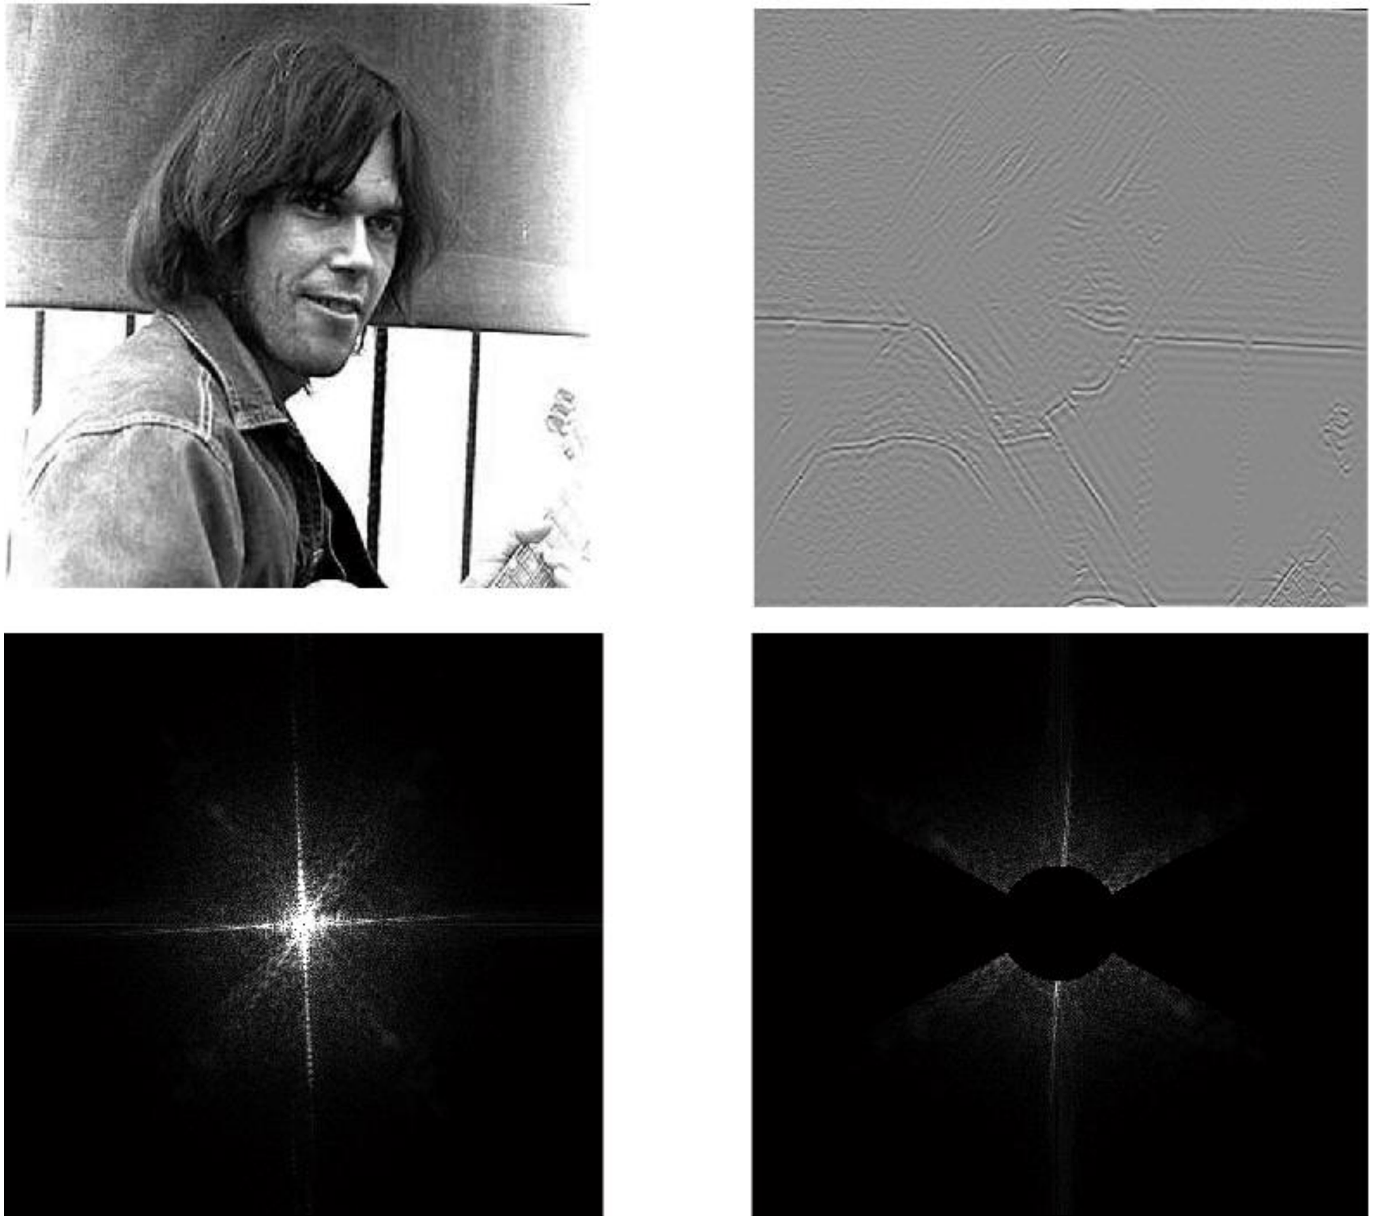

What about phase?

Efros

What about phase?

Amplitude

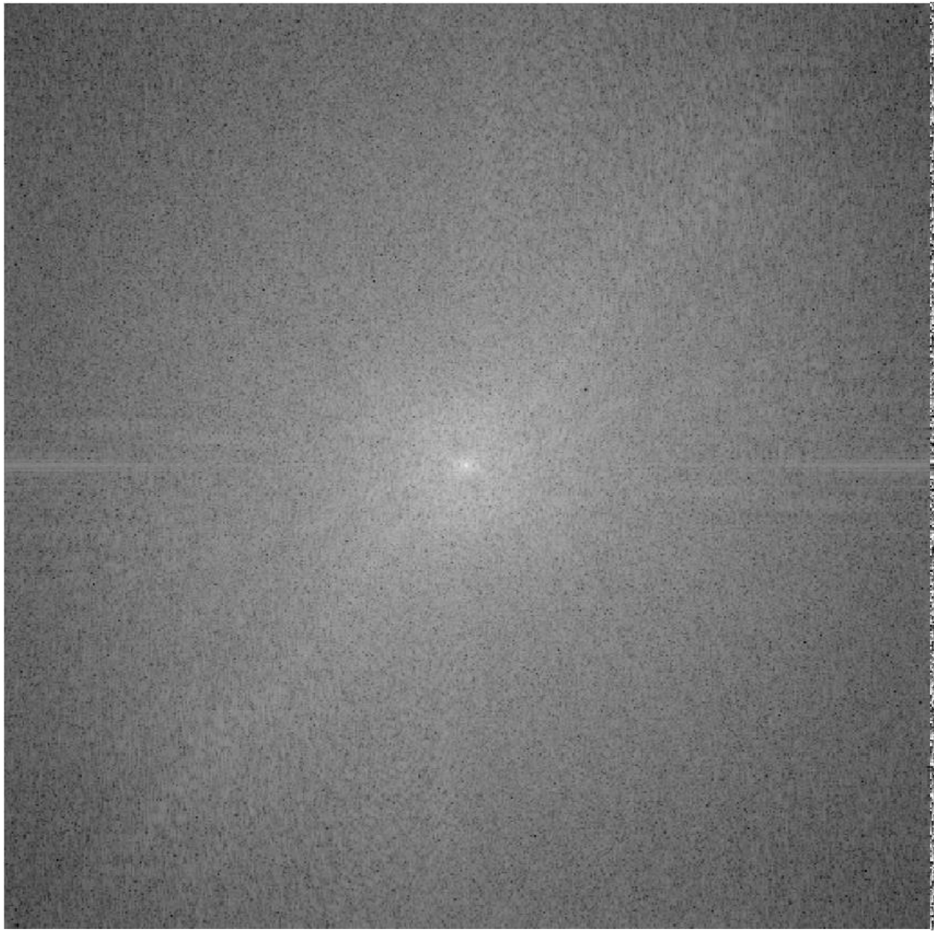

Phase

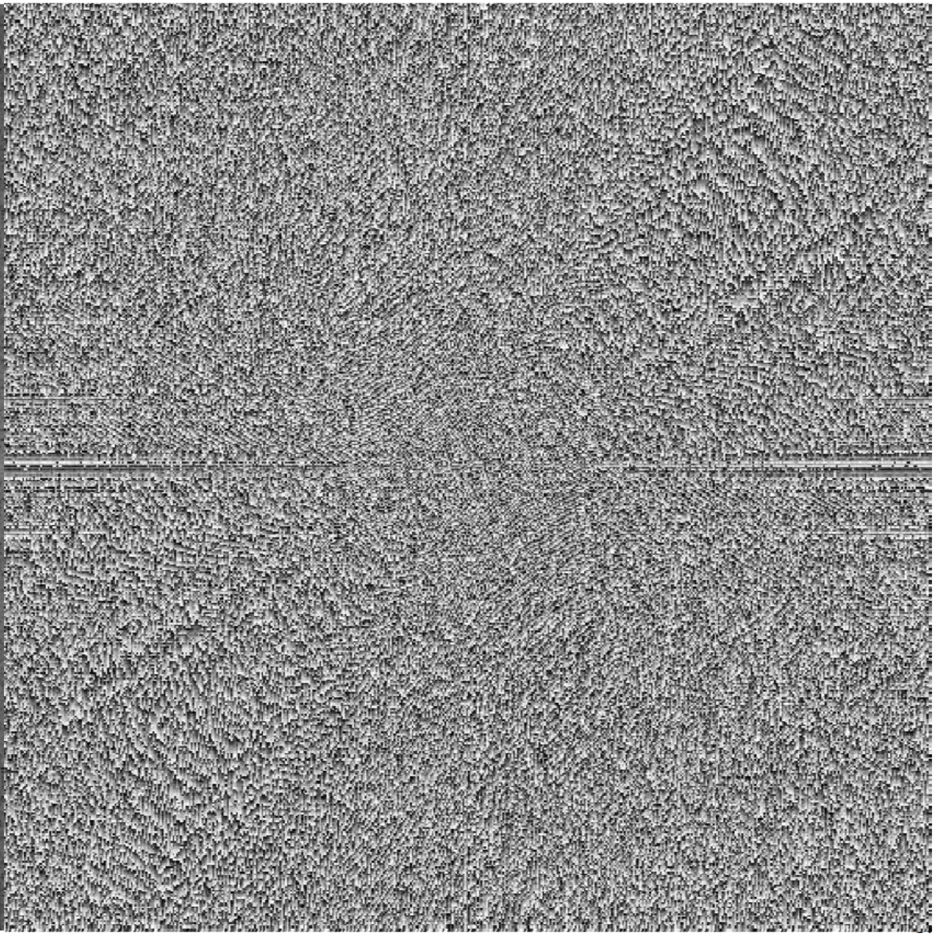

Efros

What about phase?

Efros

What about phase?

Amplitude

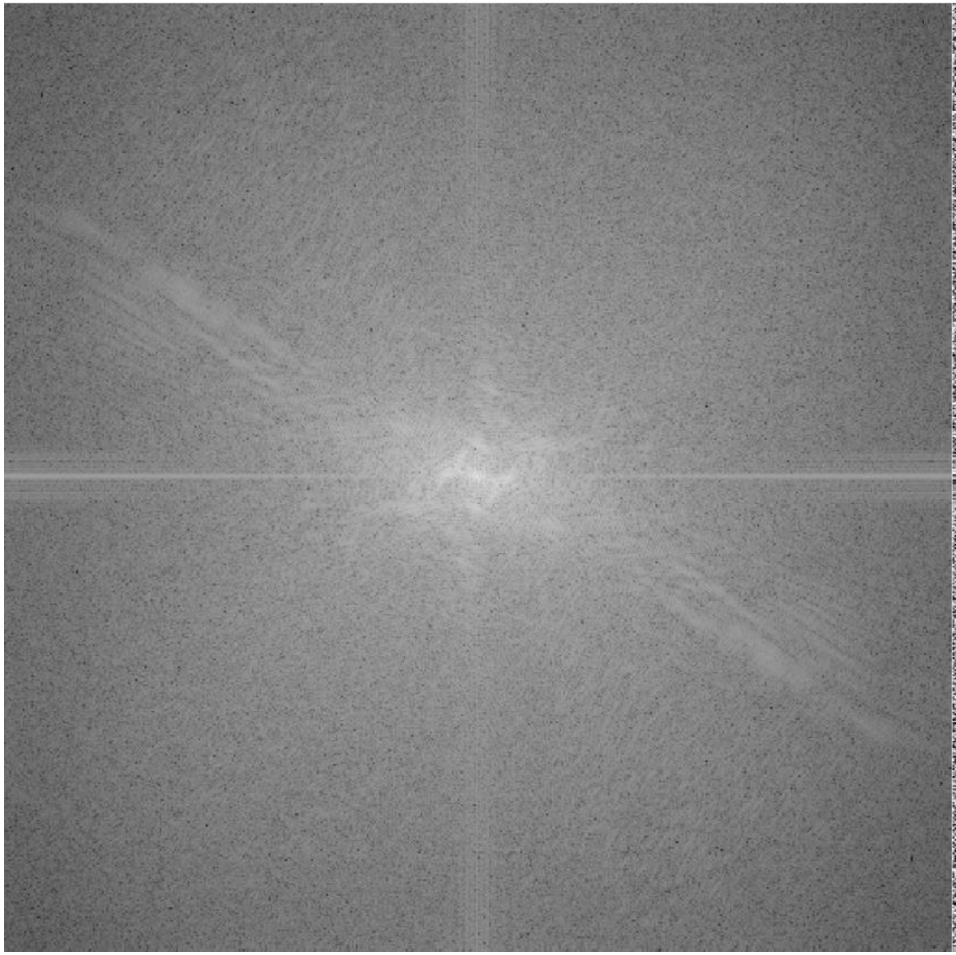

Phase

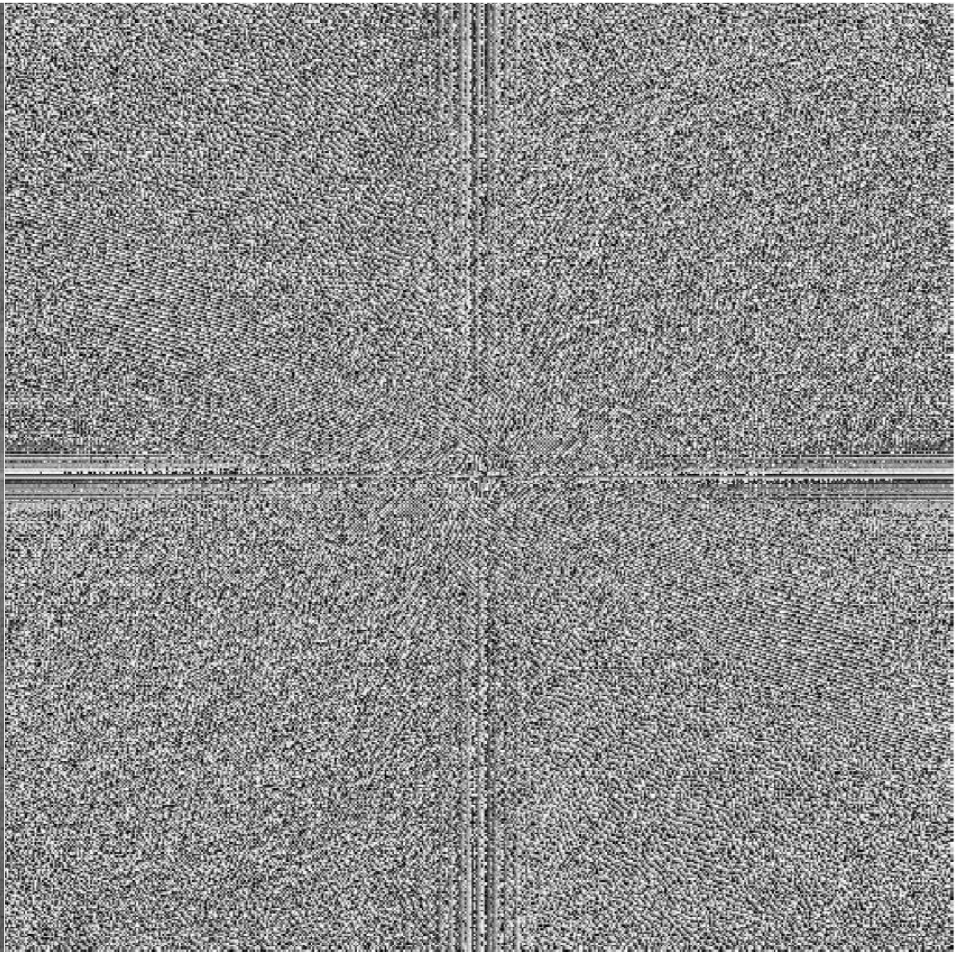

Efros

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

Cheebra

Zebra phase, cheetah amplitude

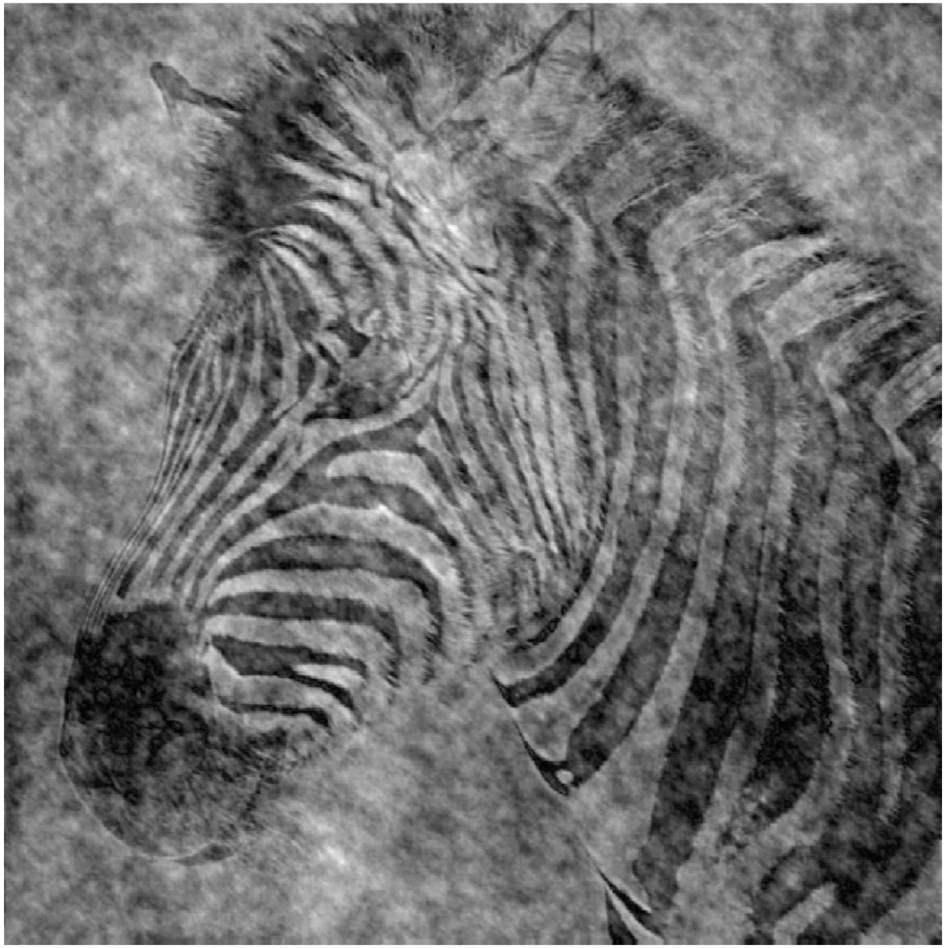

Cheetah phase, zebra amplitude

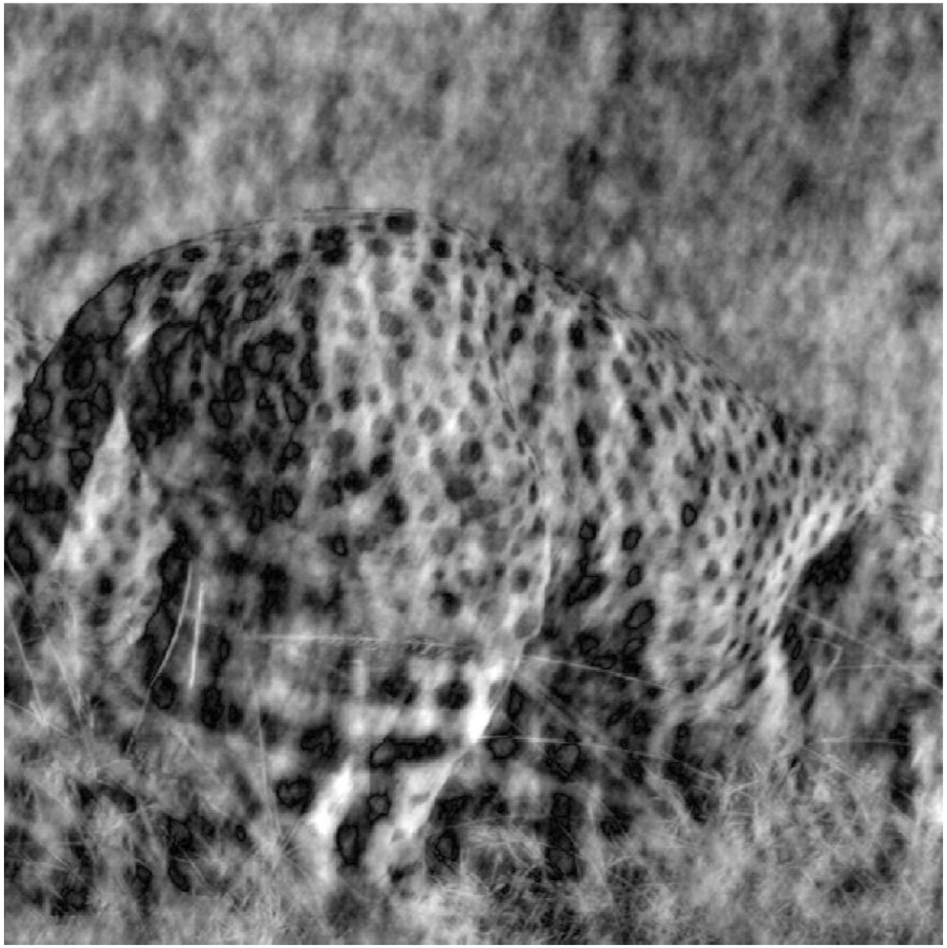

Efros

- 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
  - Not quite clear why

Efros

We stopped here in class.

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

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

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

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

Hays

Filtering in spatial domain

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

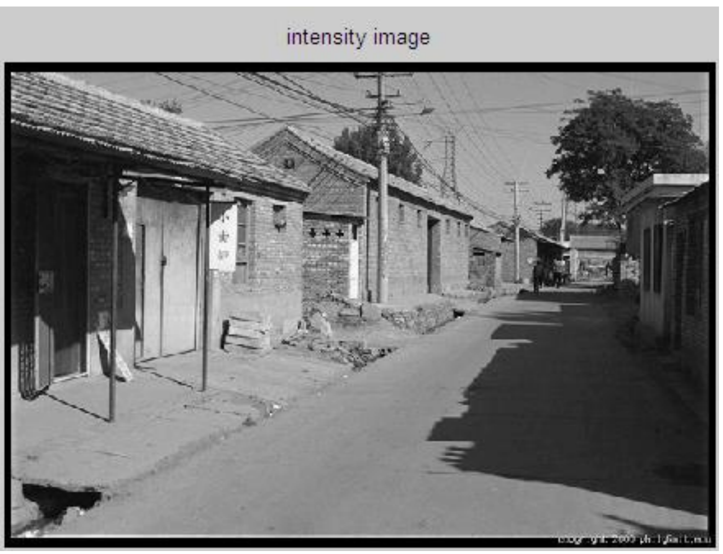

*

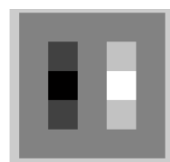

=

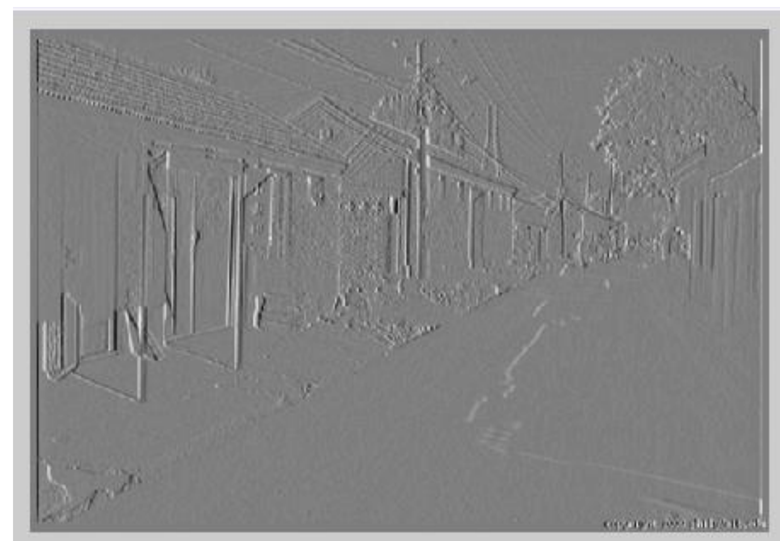

Hays

Filtering in frequency domain

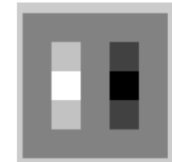

FFT

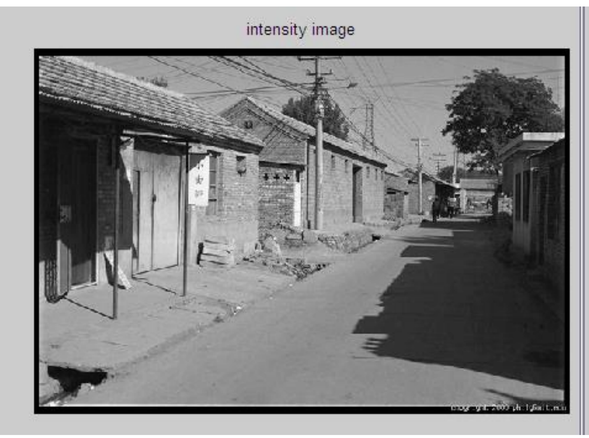

FFT

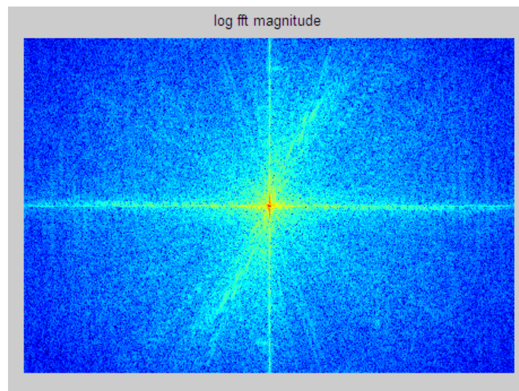

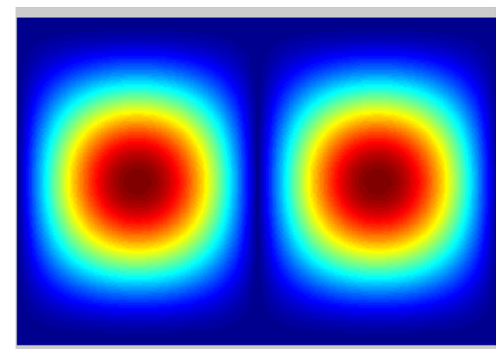

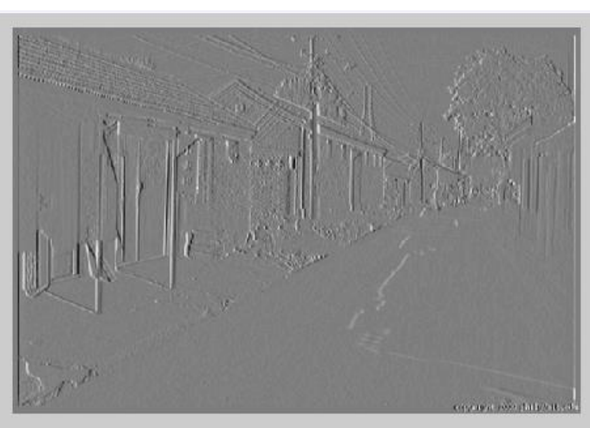

Inverse FFT

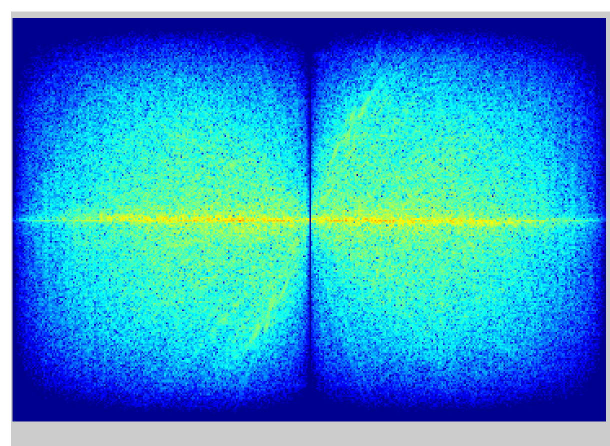

Slide: Hoiem

Fast Fourier Transform in Matlab

- Filtering with fft (fft2 -> 2D)

```
im = double(imread('…'))/255;
im = rgb2gray(im); % "im" should be a gray-scale floating point image
[imh, imw] = size(im);

hs = 50; % filter half-size
fil = fspecial('gaussian', hs*2+1, 10);

fftsize = 1024; % should be order of 2 (for speed) and include padding
im_fft  = fft2(im,  fftsize, fftsize);                        % 1) fft im with padding
fil_fft = fft2(fil, fftsize, fftsize);                        % 2) fft fil, pad to same size as image
im_fil_fft = im_fft .* fil_fft;                               % 3) multiply fft images
im_fil = ifft2(im_fil_fft);                                   % 4) inverse fft2
im_fil = im_fil(1+hs:size(im,1)+hs, 1+hs:size(im, 2)+hs); % 5) remove padding
```

- Displaying with fft

```
figure(1), imagesc(log(abs(fftshift(im_fft)))), axis image, colormap jet
```

Slide: Hoiem

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

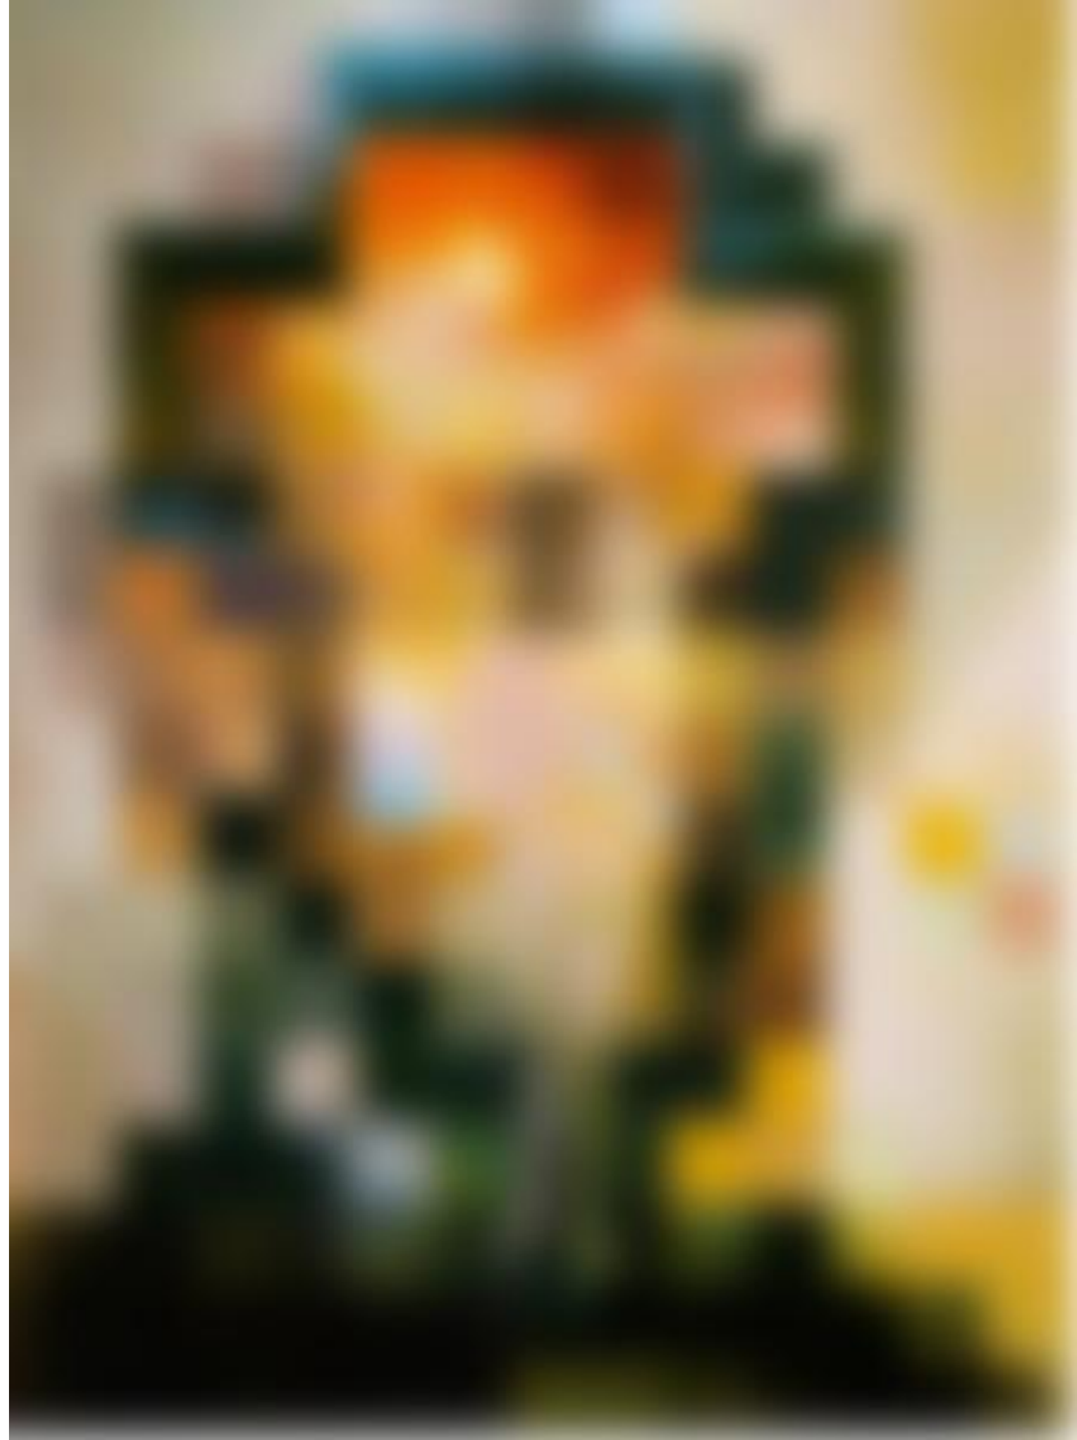

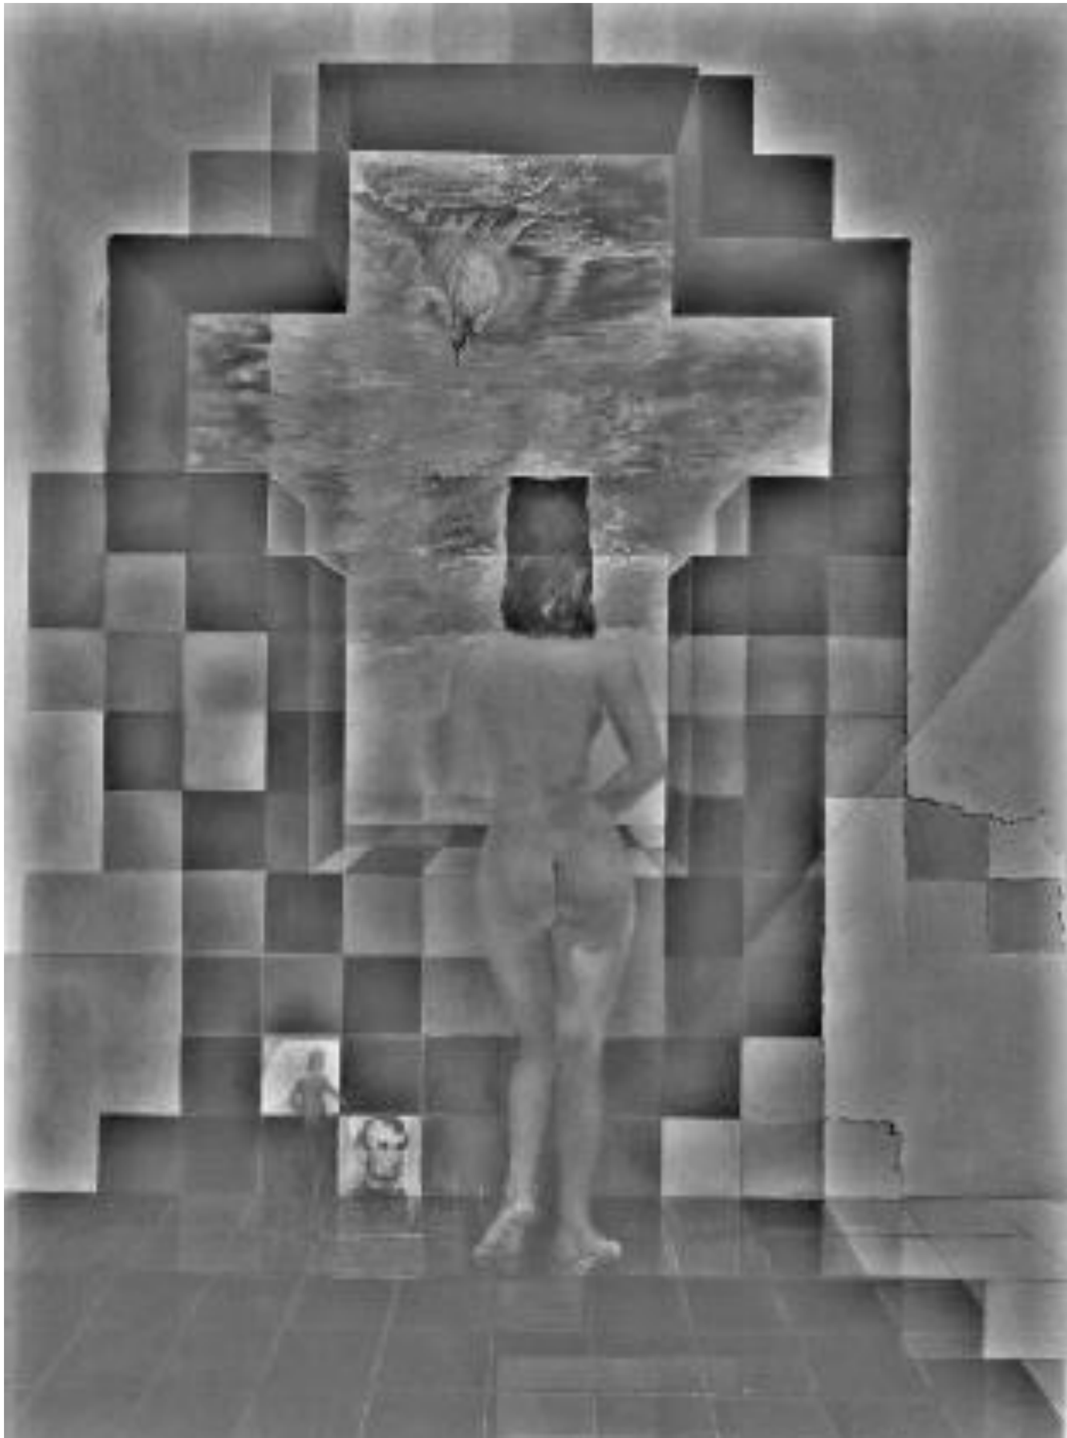

On Friday:

- More frequency analysis with Fourier.
- Resampling and image pyramids.