The Frequency Domain

Somewhere in Cinque Terre, May 2005

Many slides borrowed from Steve Seitz

15-463: Computational Photography
Alexei Efros, CMU, Fall 2011

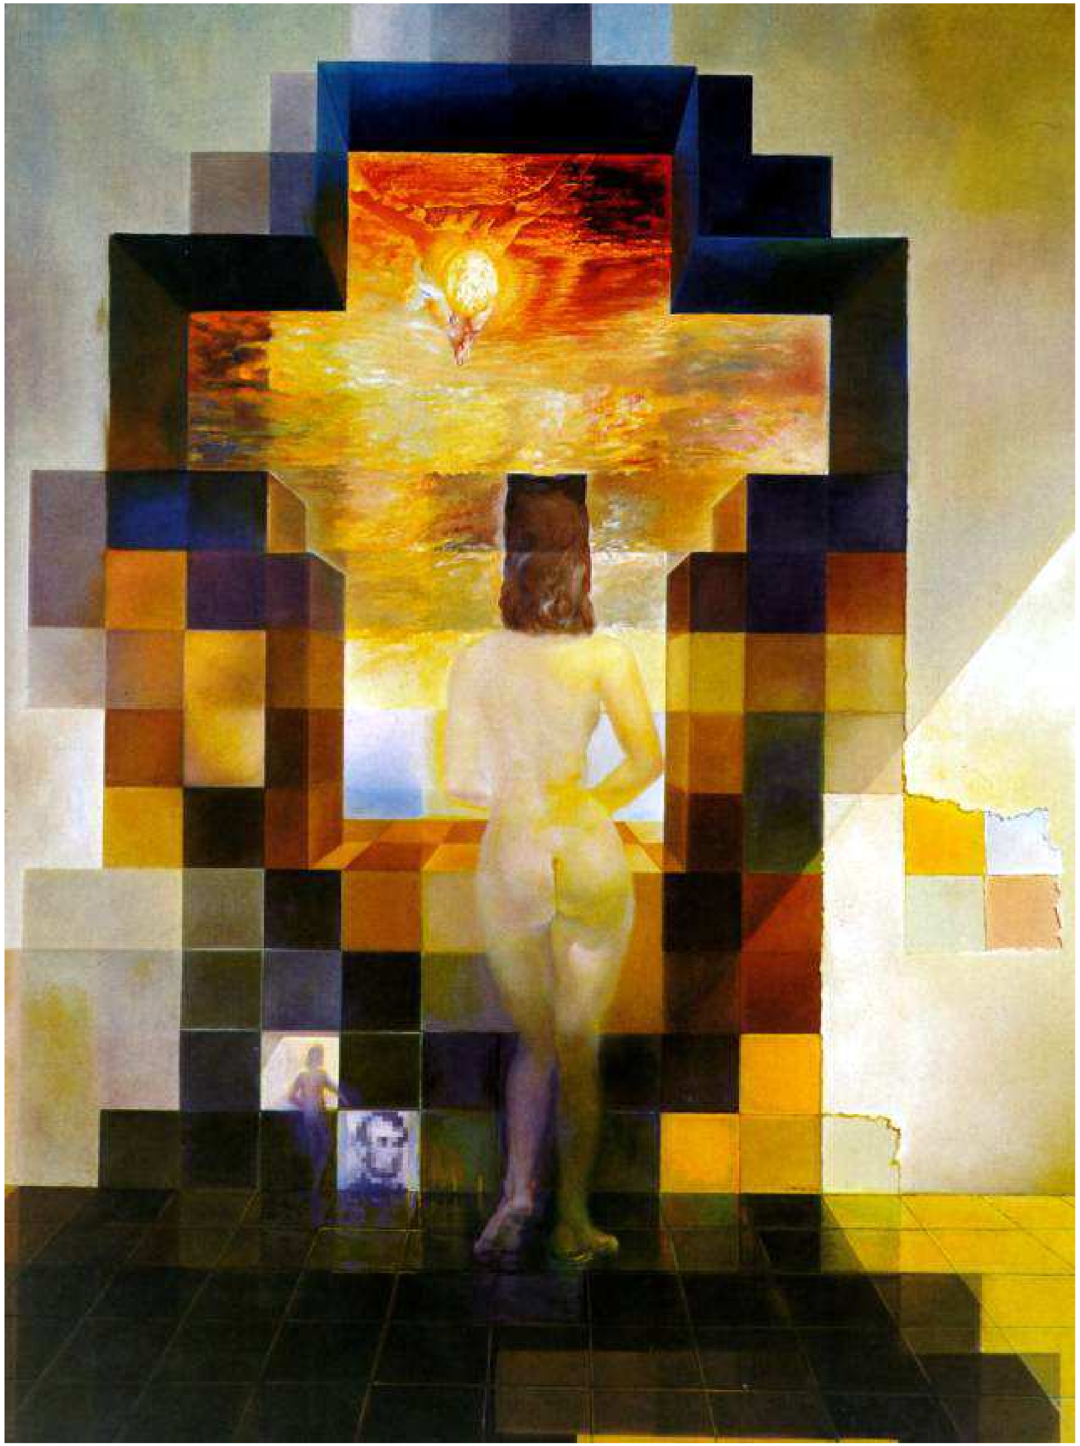

Salvador Dali
"Gala Contemplating the Mediterranean Sea, which at 30 meters becomes the portrait of Abraham Lincoln", 1976

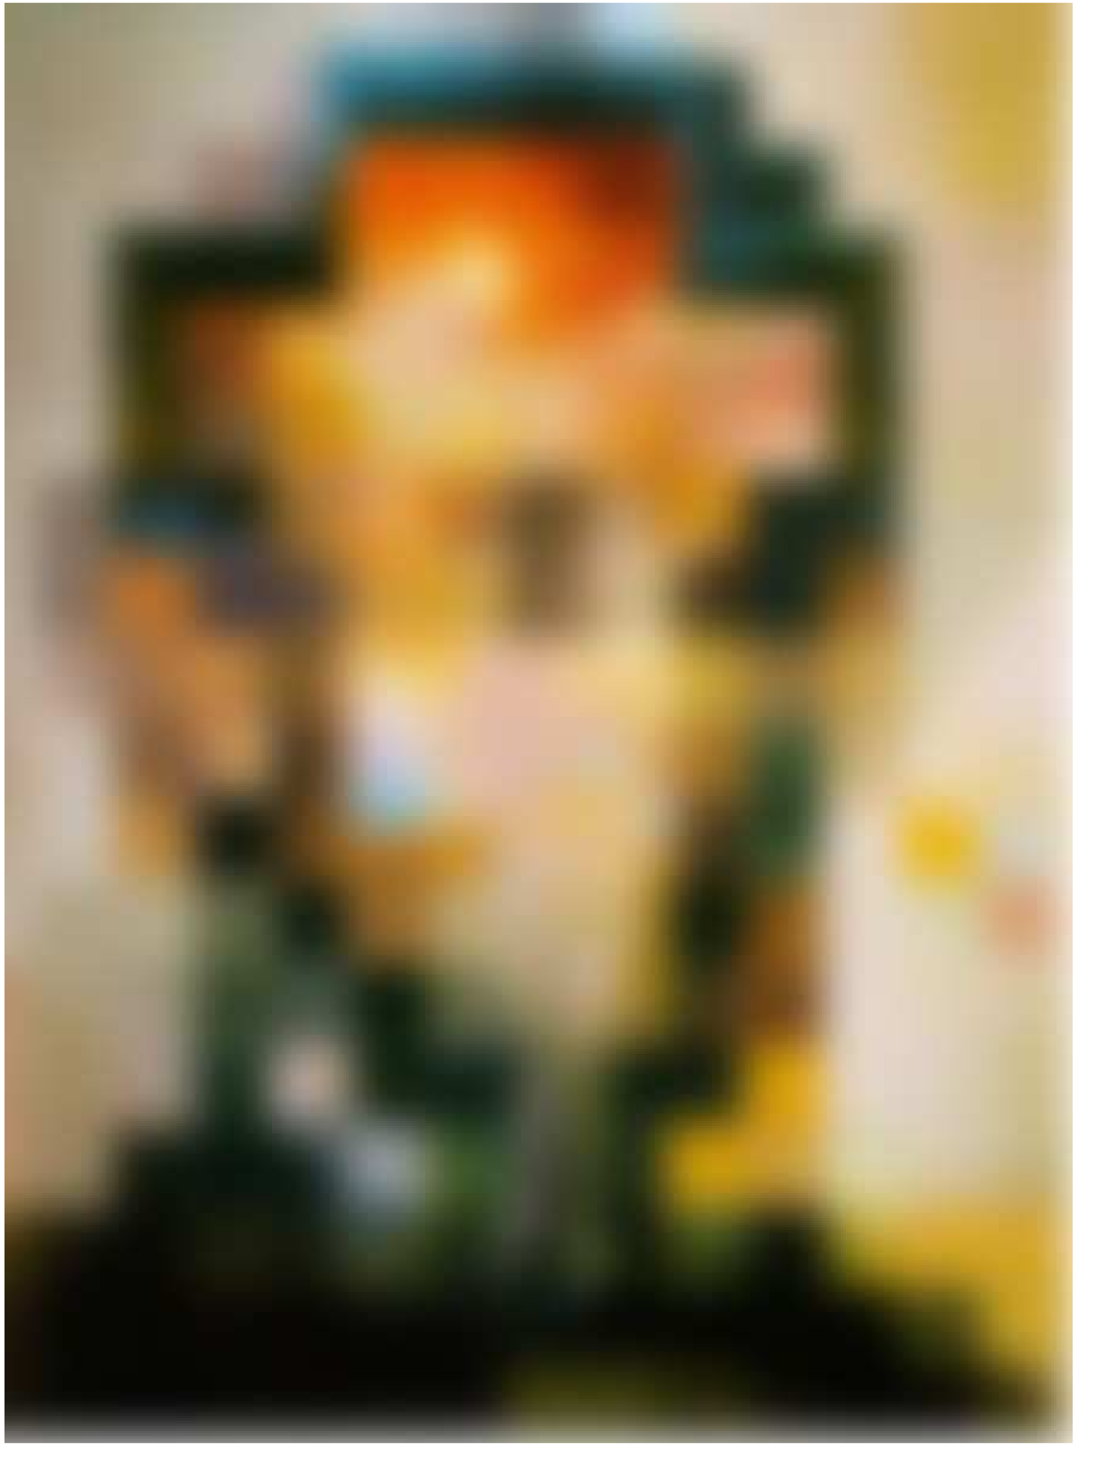

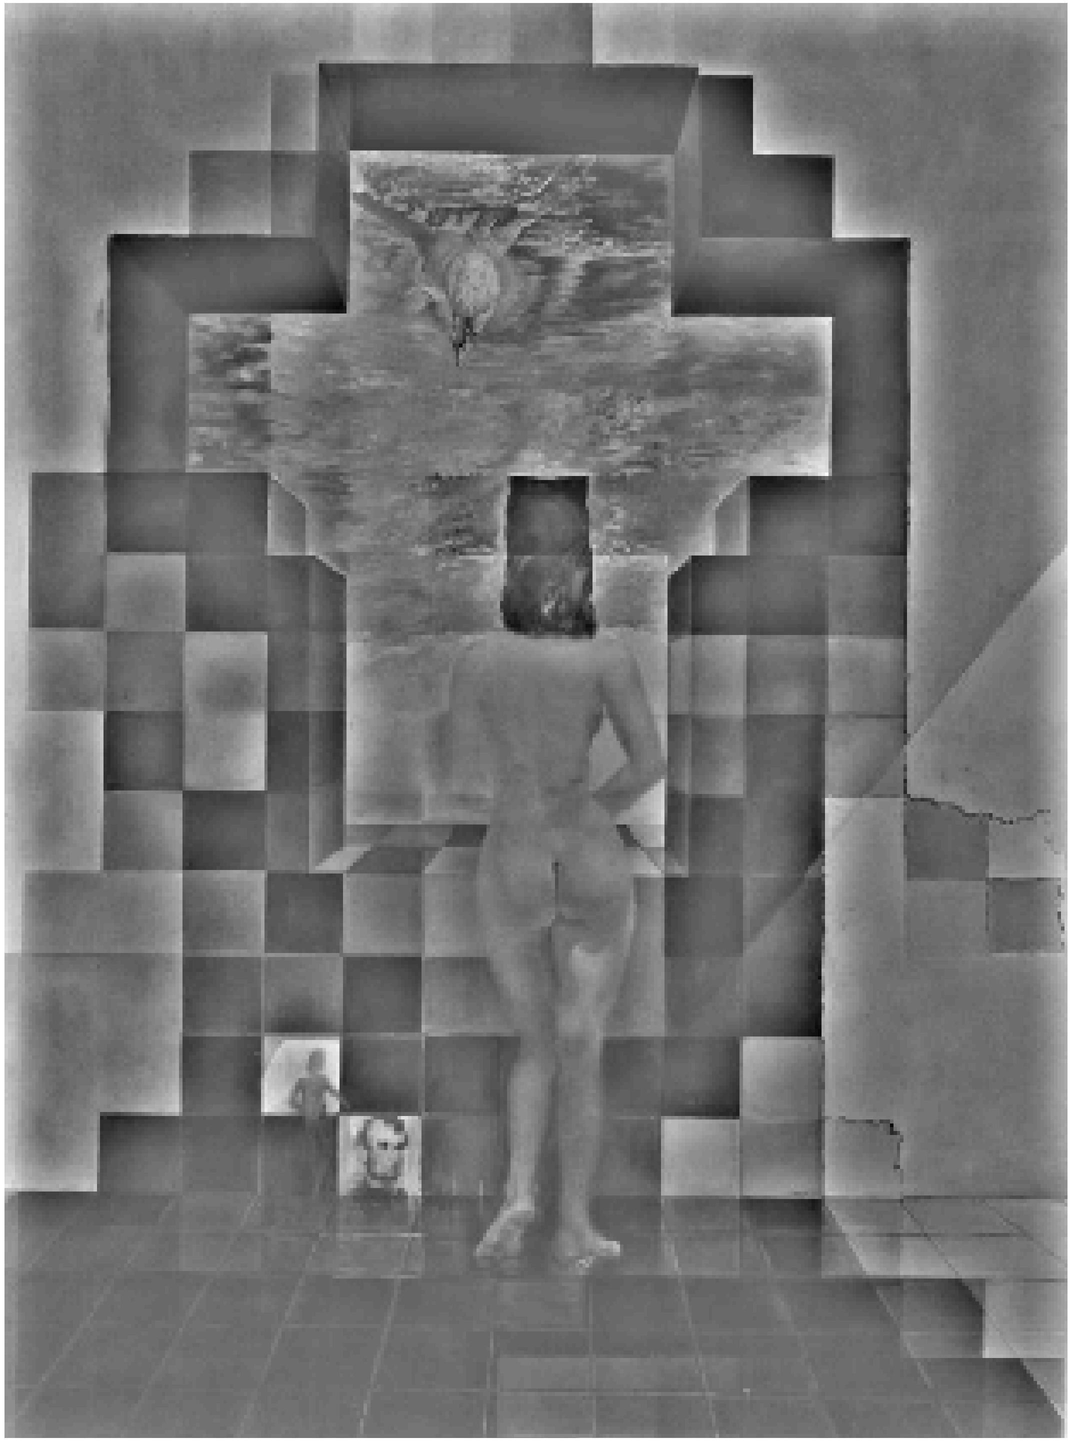

A nice set of basis

Teases away fast vs. slow changes in the image.

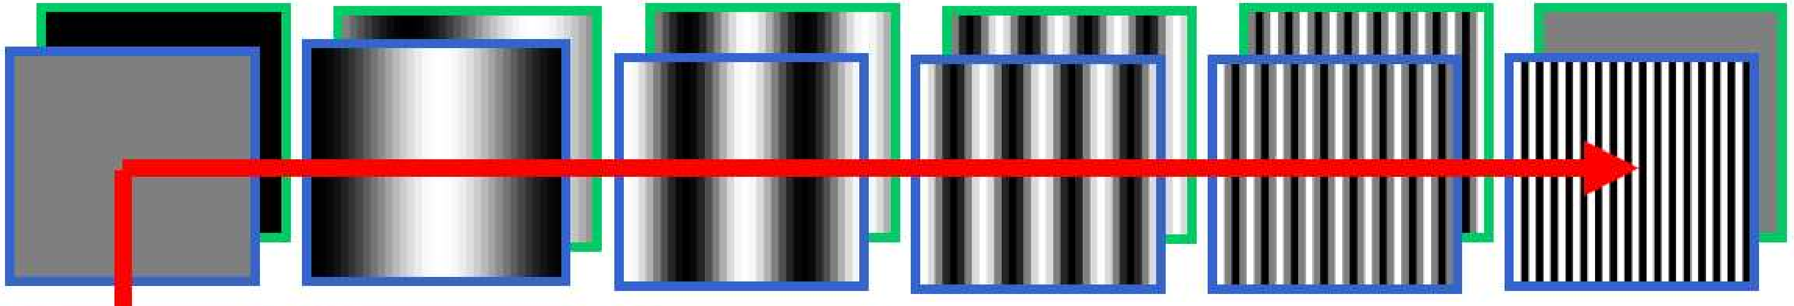

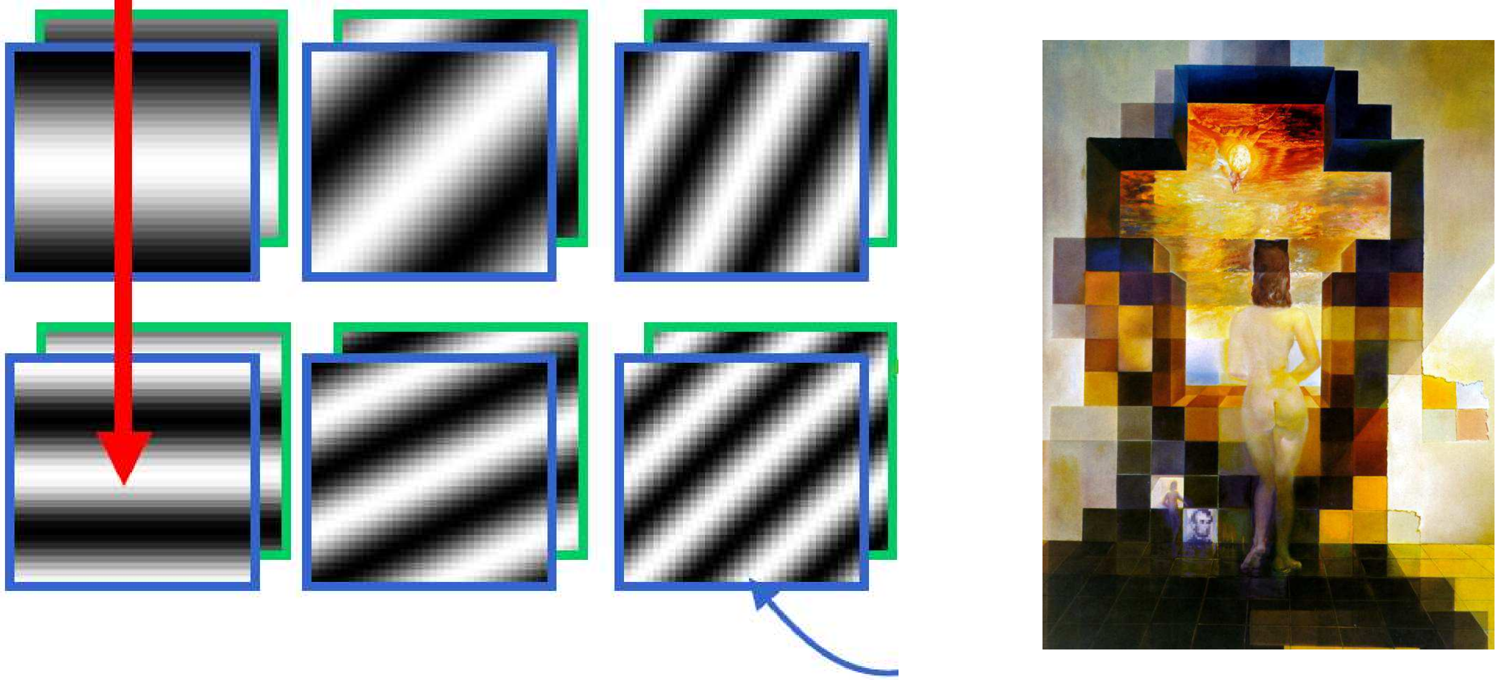

This change of basis has a special name…

Jean Baptiste Joseph Fourier (1768-1830)

...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.

had crazy idea (1807

Any *univariate function ca be rewritten as a weighted sum of sines and cosines 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
- there are some subtle restrictions

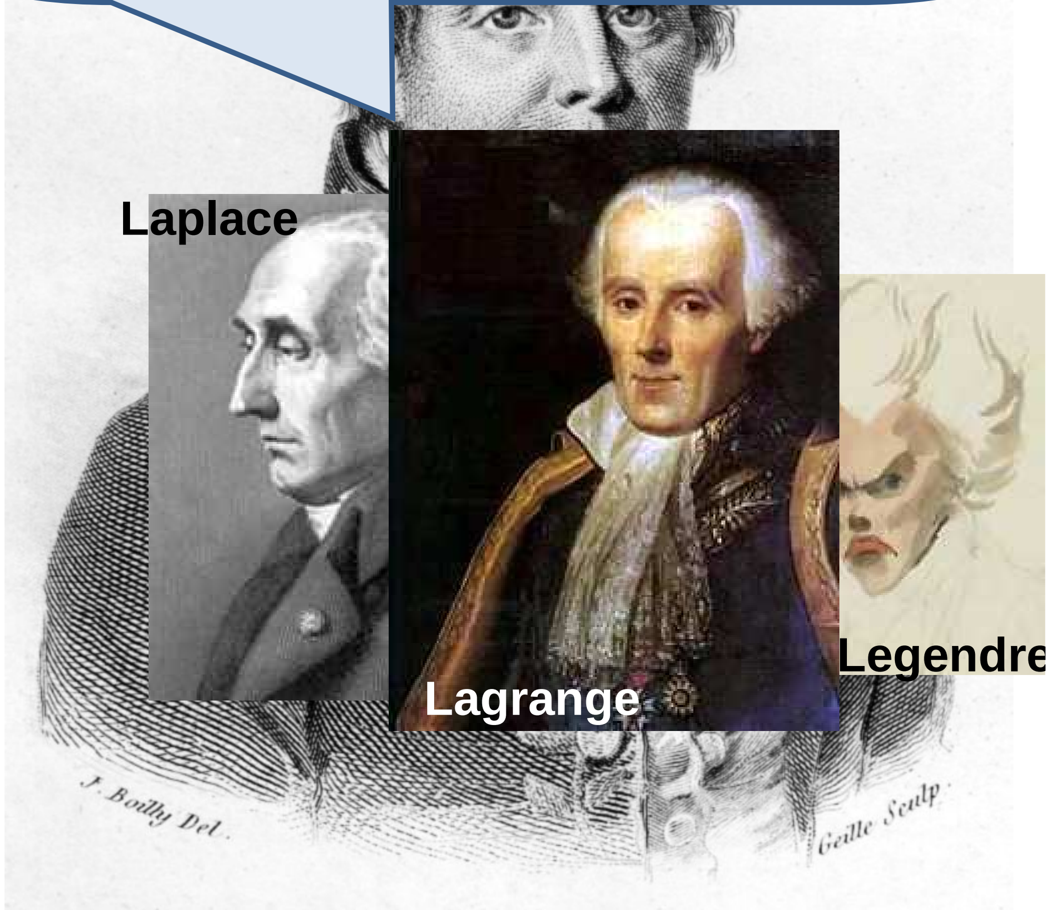

A sum of sines

Our building block:

$$A\sin(\omega x + \phi)$$

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

How many degrees of freedom?

What does each control?

Which one encodes the coarse vs. fine structure of the signal?

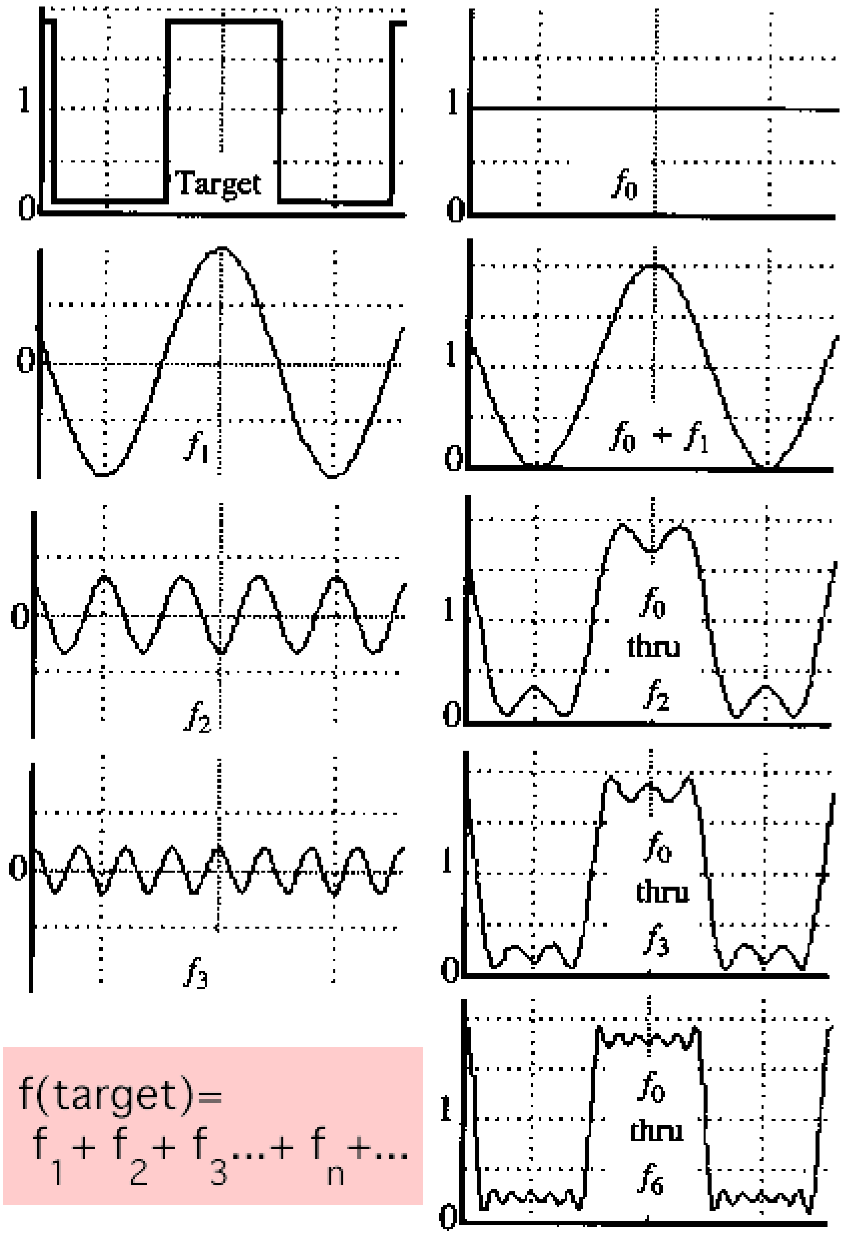

f(target)=
$f_1 + f_2 + f_3 ... + f_n + ...$

Fourier Transform

We want to understand the frequency ω of our signal. So, let's reparametrize the signal by ω instead of *x*:

f(x) → Fourier Transform → ***F(ω)***

For every ω from 0 to inf, ***F(ω)*** holds the amplitude *A* and phase ϕ of the corresponding sine $A\sin(\omega x + \phi)$

- How can *F* hold both? Complex number trick!

$$F(\omega) = R(\omega) + iI(\omega)$$

$$A = \pm\sqrt{R(\omega)^2 + I(\omega)^2} \qquad \phi = \tan^{-1}\frac{I(\omega)}{R(\omega)}$$

We can always go back:

F(ω) → Inverse Fourier Transform → ***f(x)***

Time and Frequency

example : $g(t) = \sin(2pf\,t) + (1/3)\sin(2p(3f)\,t)$

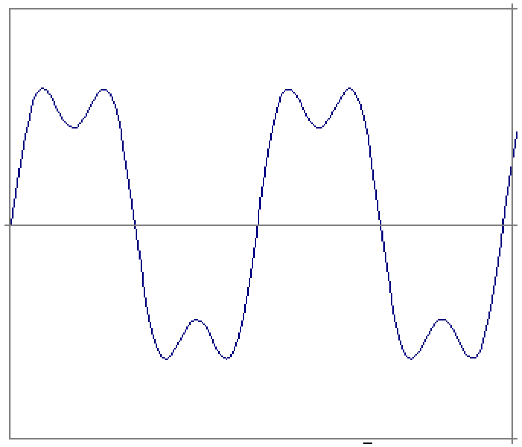

Time and Frequency

example : $g(t) = \sin(2pf\ t) + (1/3)\sin(2p(3f)\ t)$

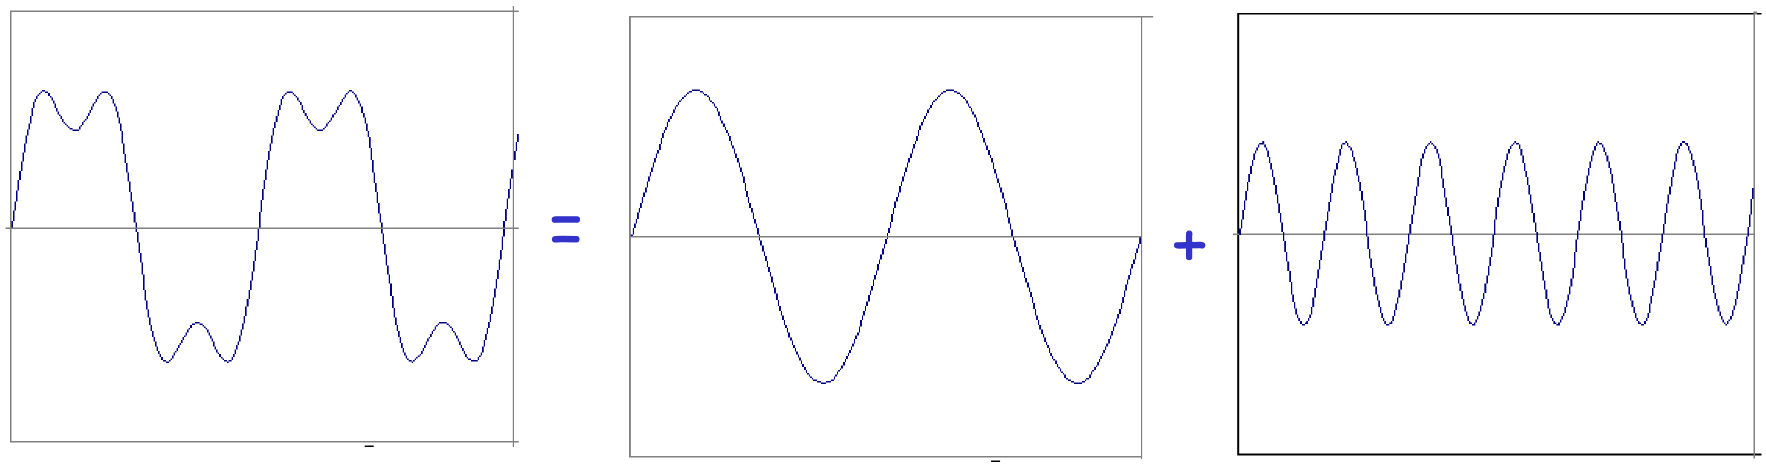

Frequency Spectra

example : $g(t) = \sin(2pf\ t) + (1/3)\sin(2p(3f)\ t)$

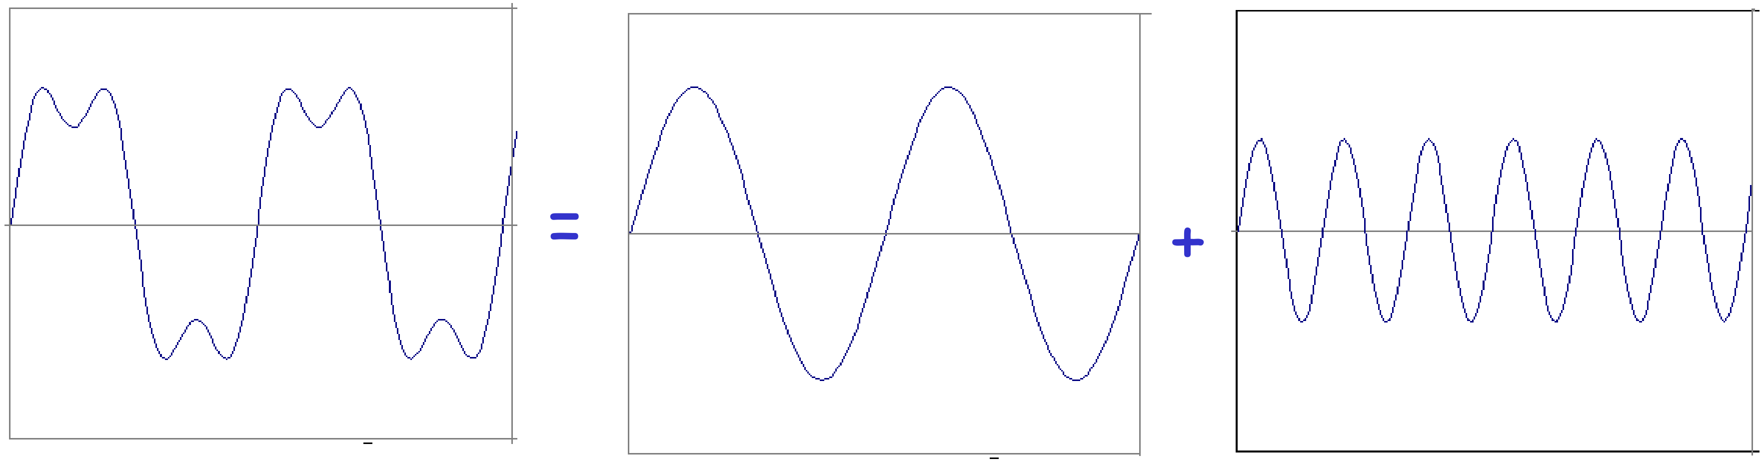

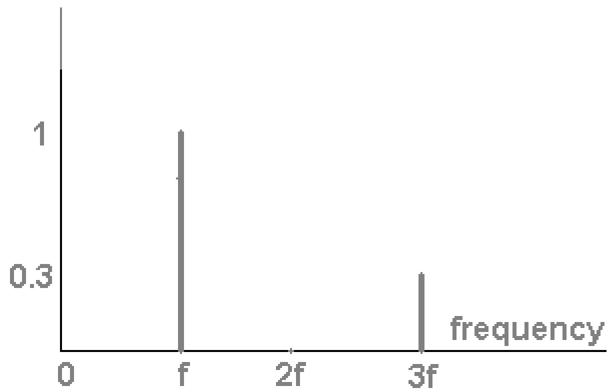

Frequency Spectra

Usually, frequency is more interesting than the phase

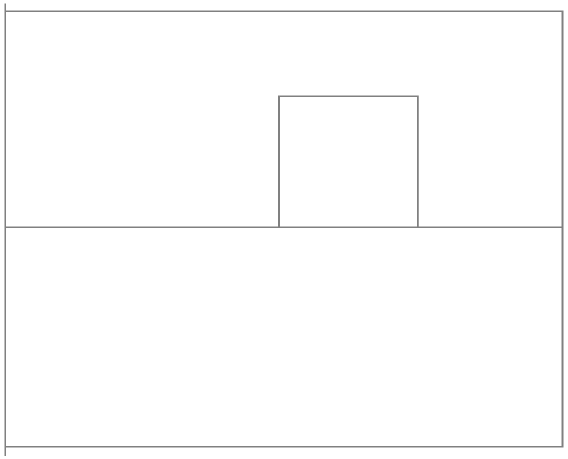

Frequency Spectra

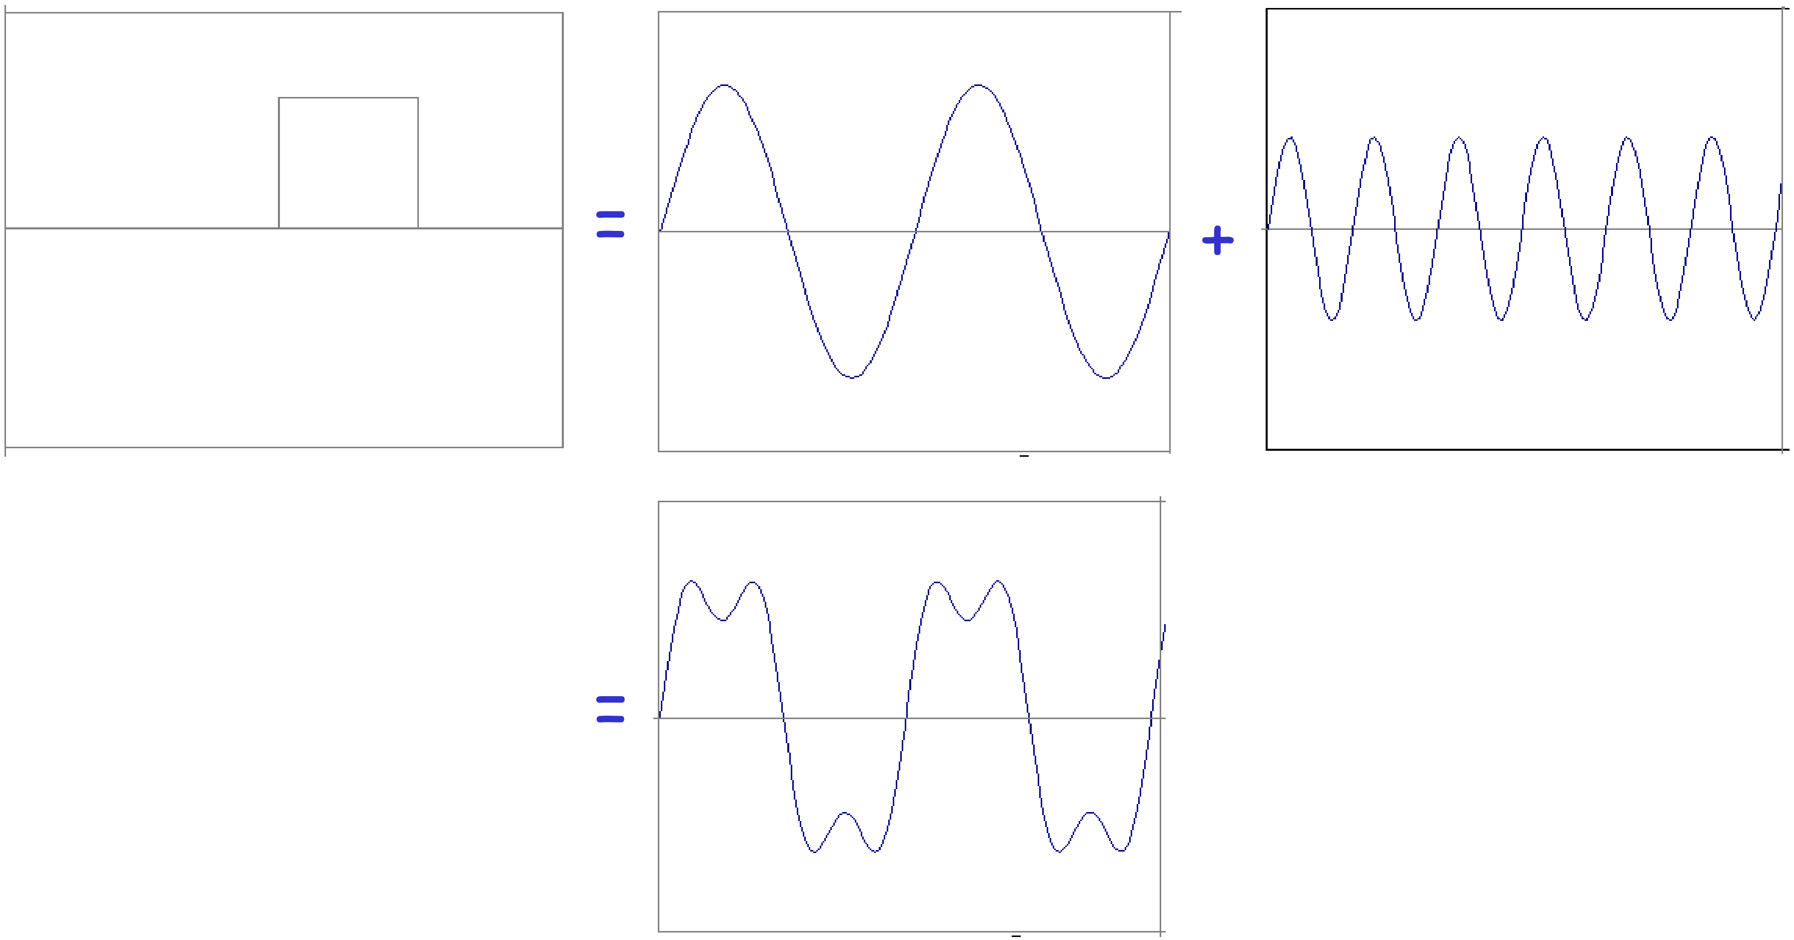

Frequency Spectra

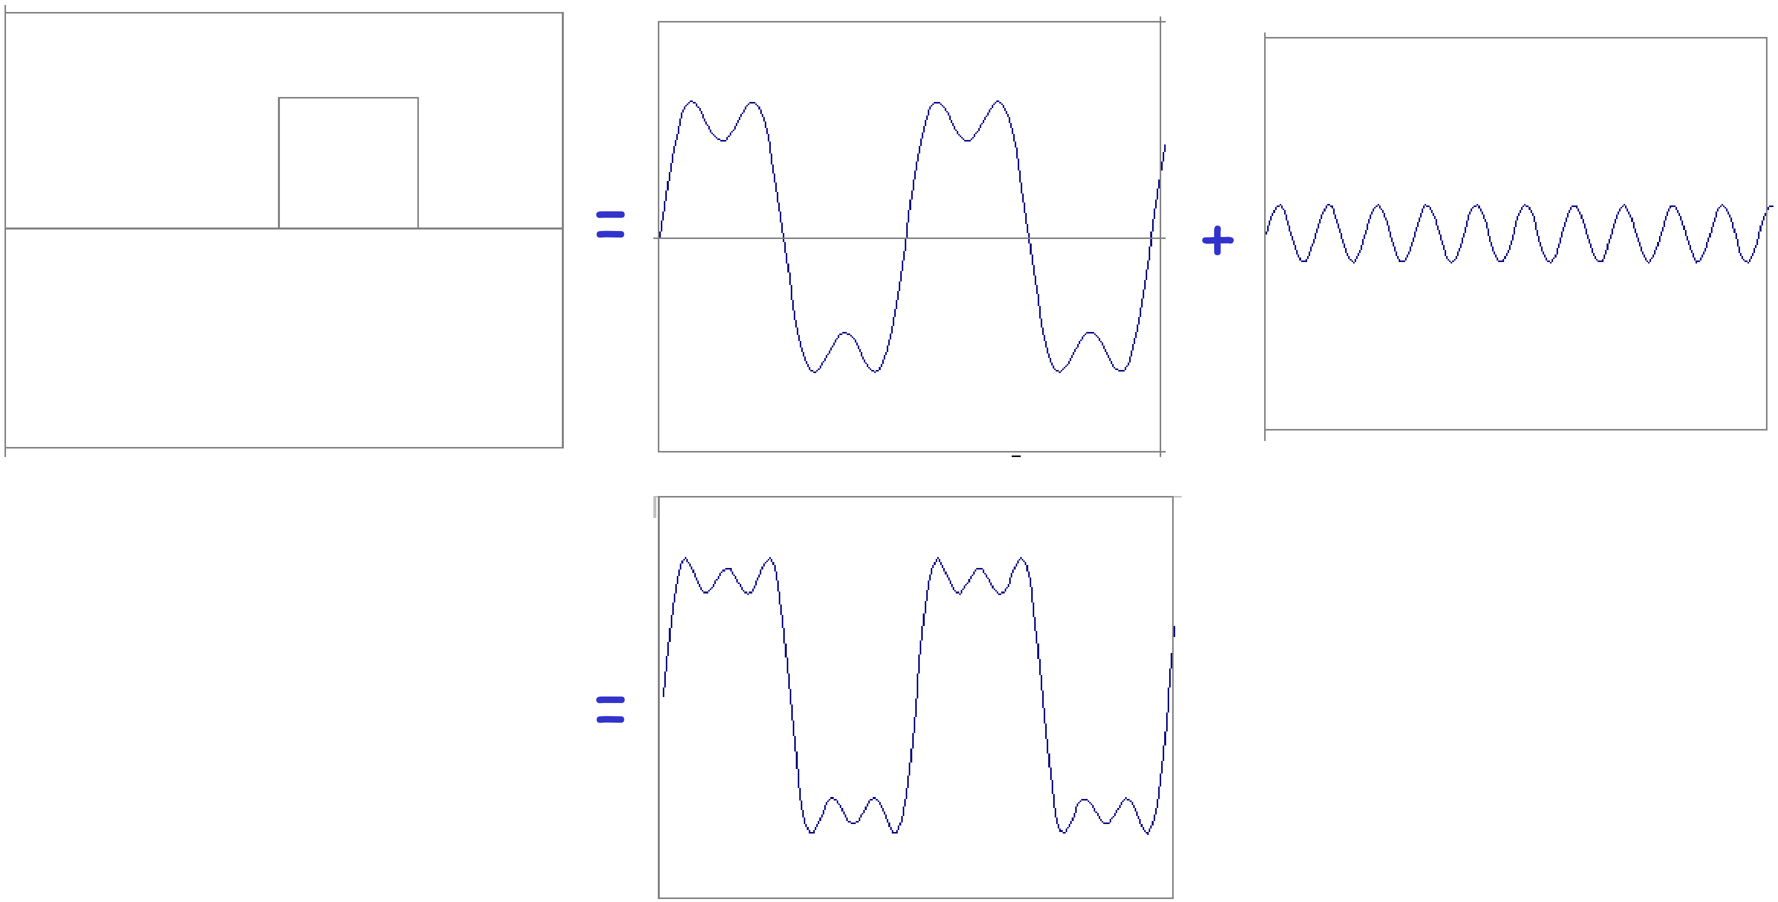

Frequency Spectra

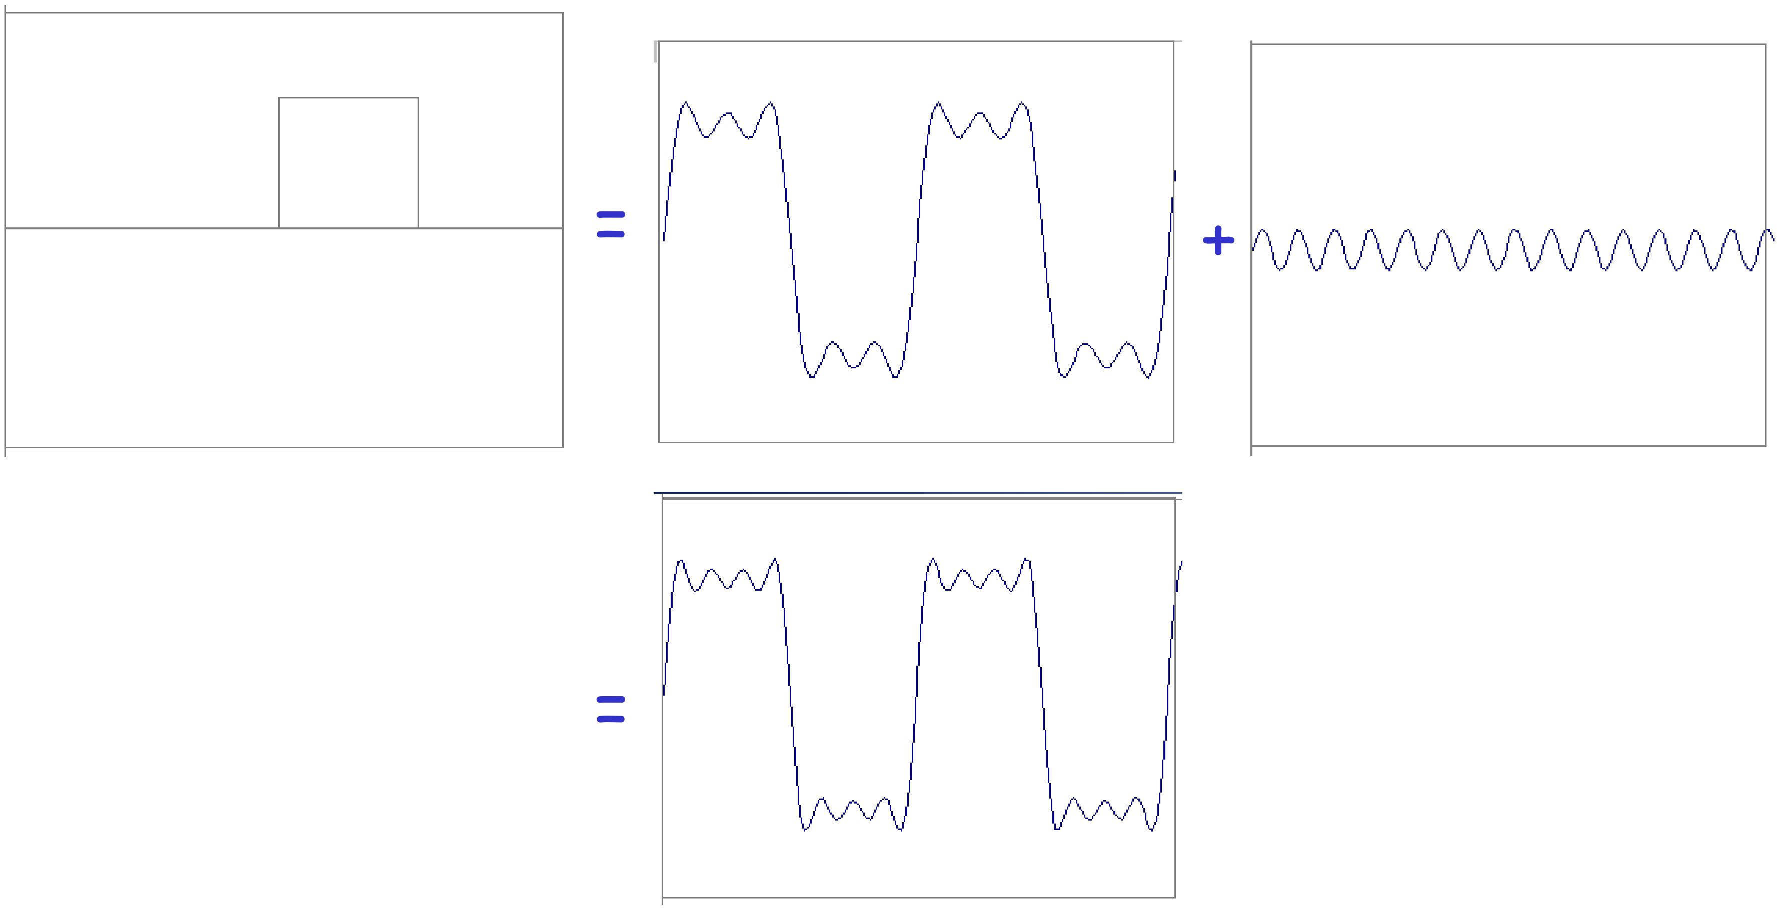

Frequency Spectra

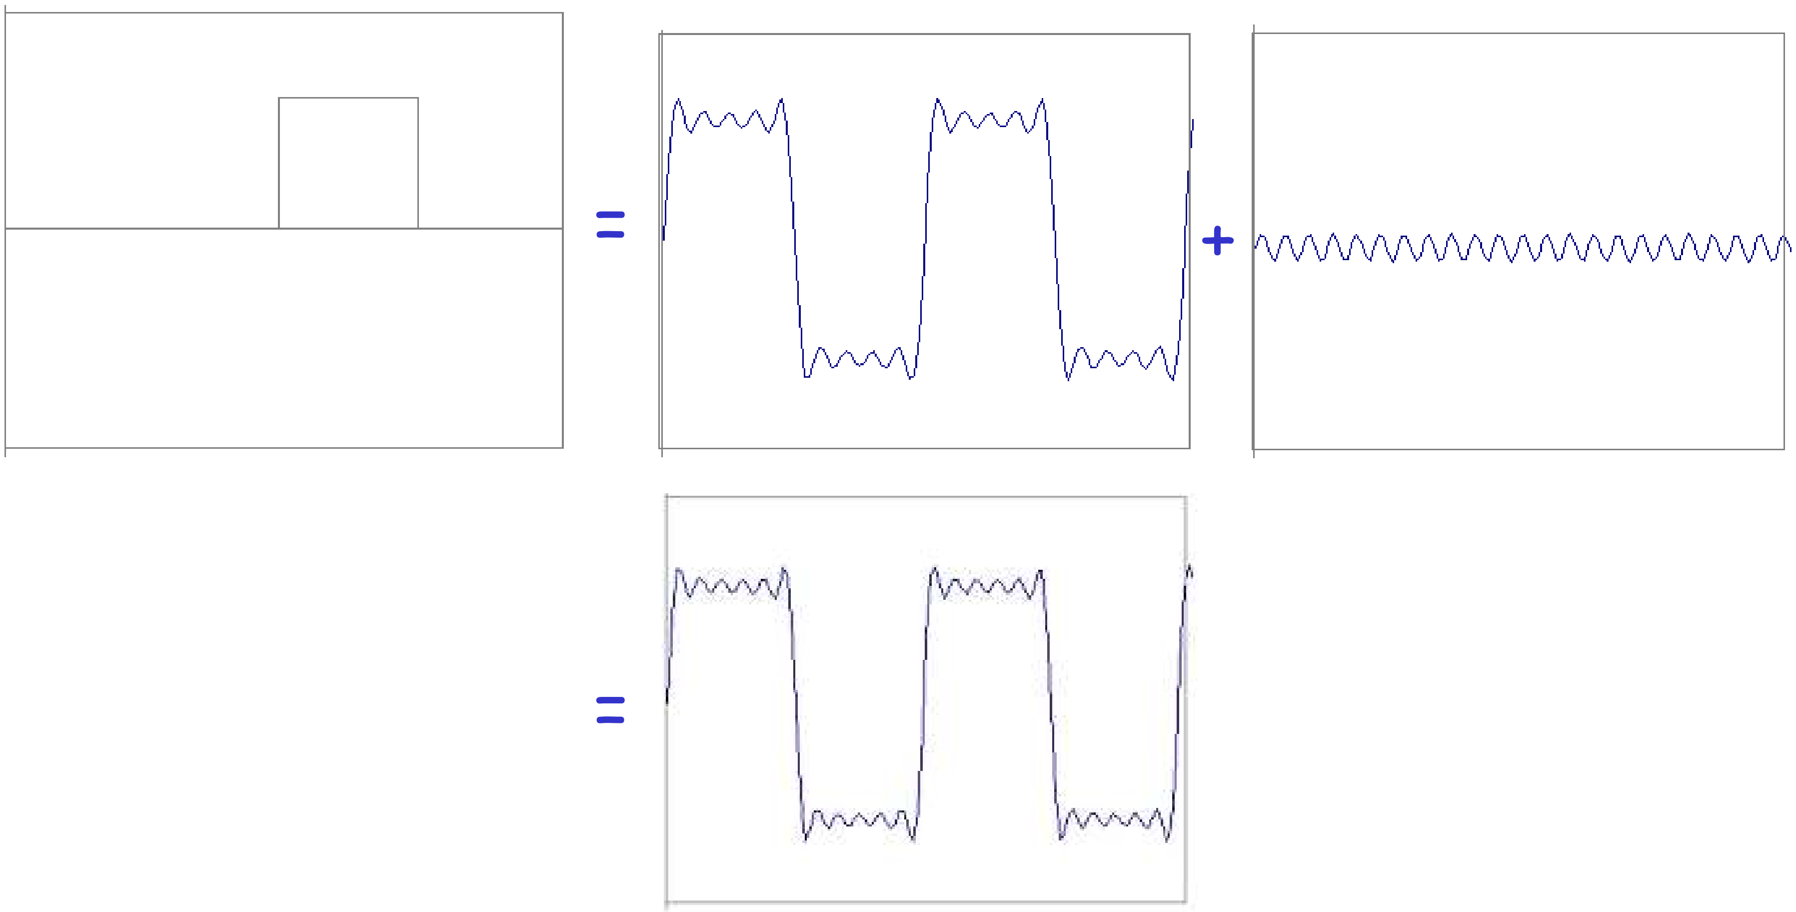

Frequency Spectra

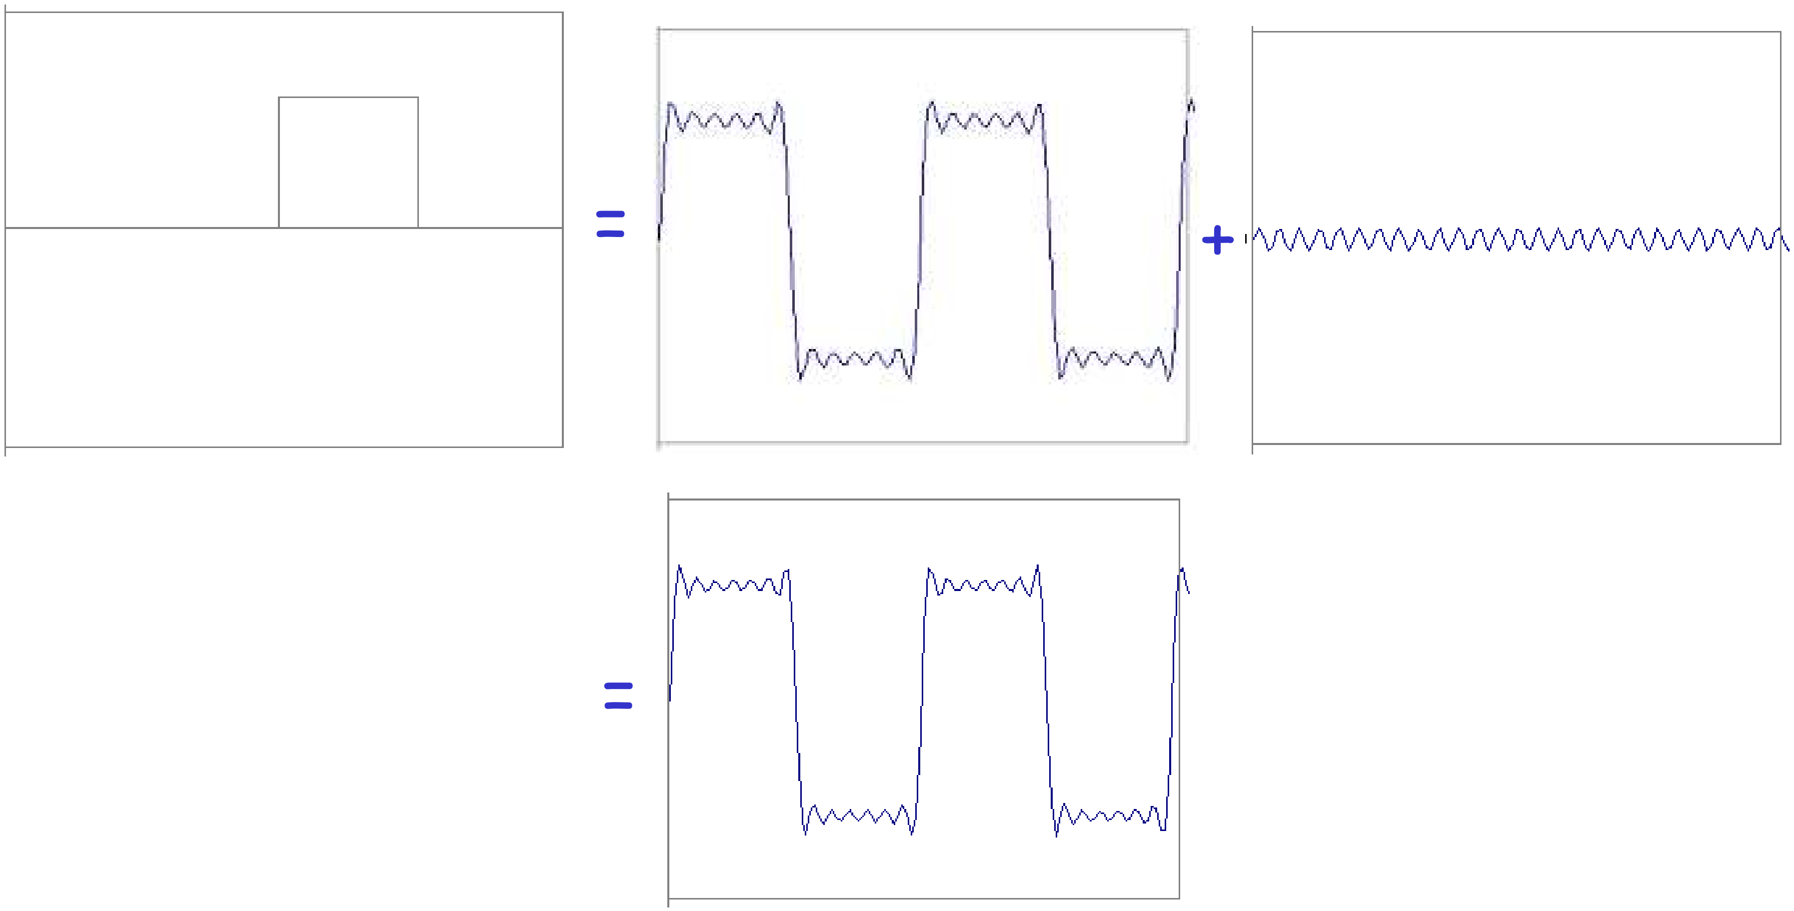

Frequency Spectra

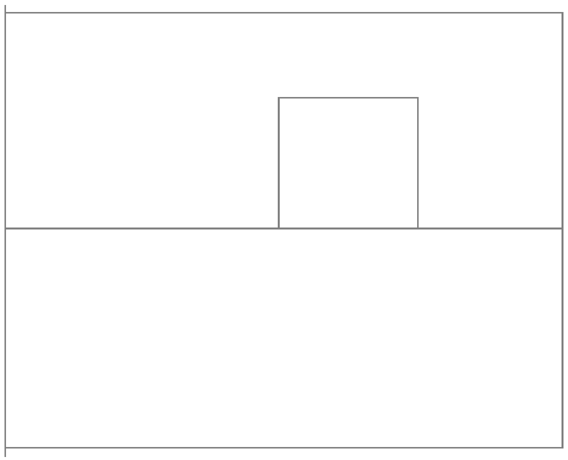

$$= A\sum_{k=1}^{\infty}\frac{1}{k}\sin(2\pi kt)$$

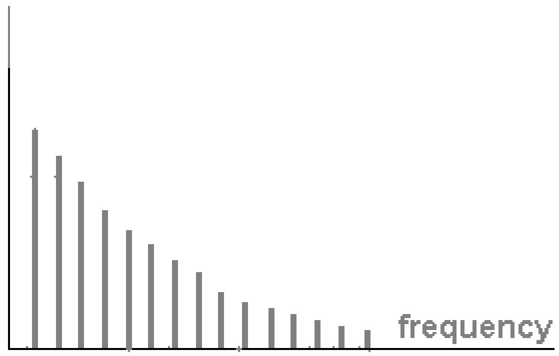

Frequency Spectra

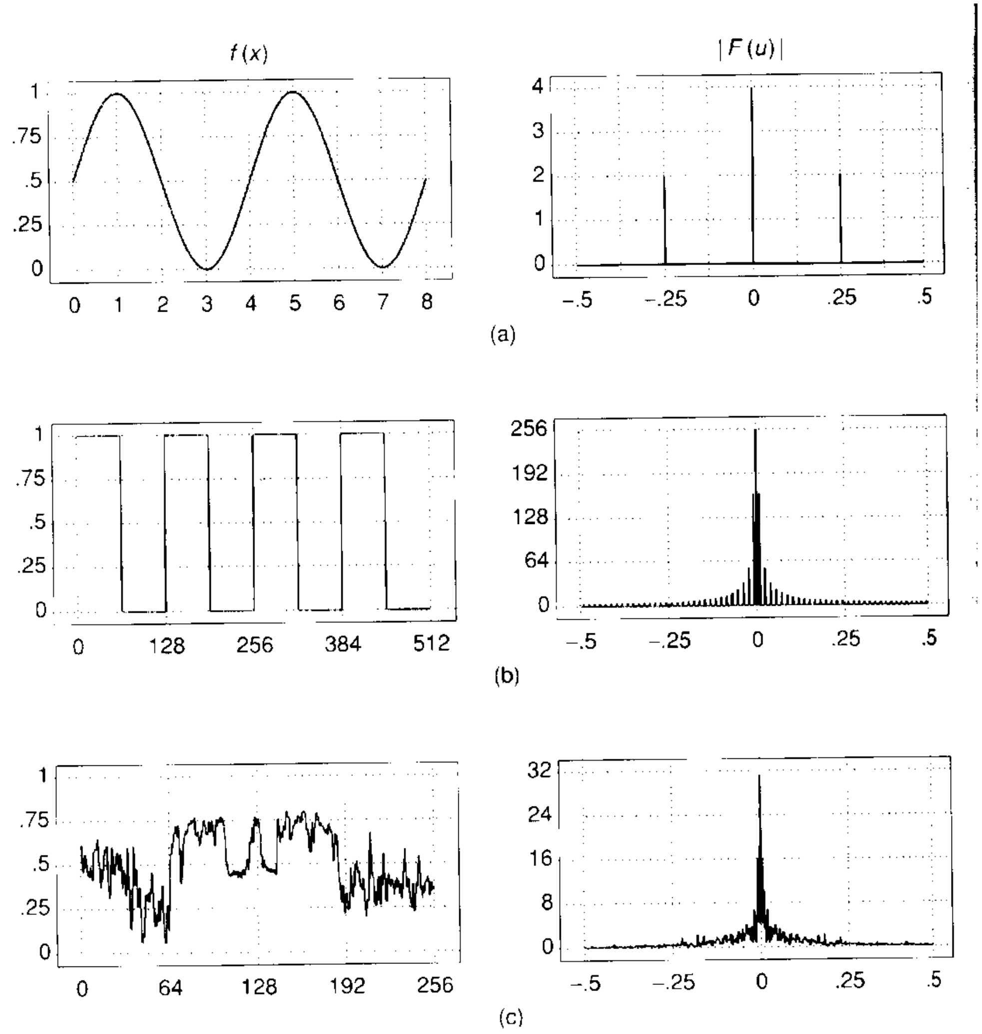

Extension to 2D

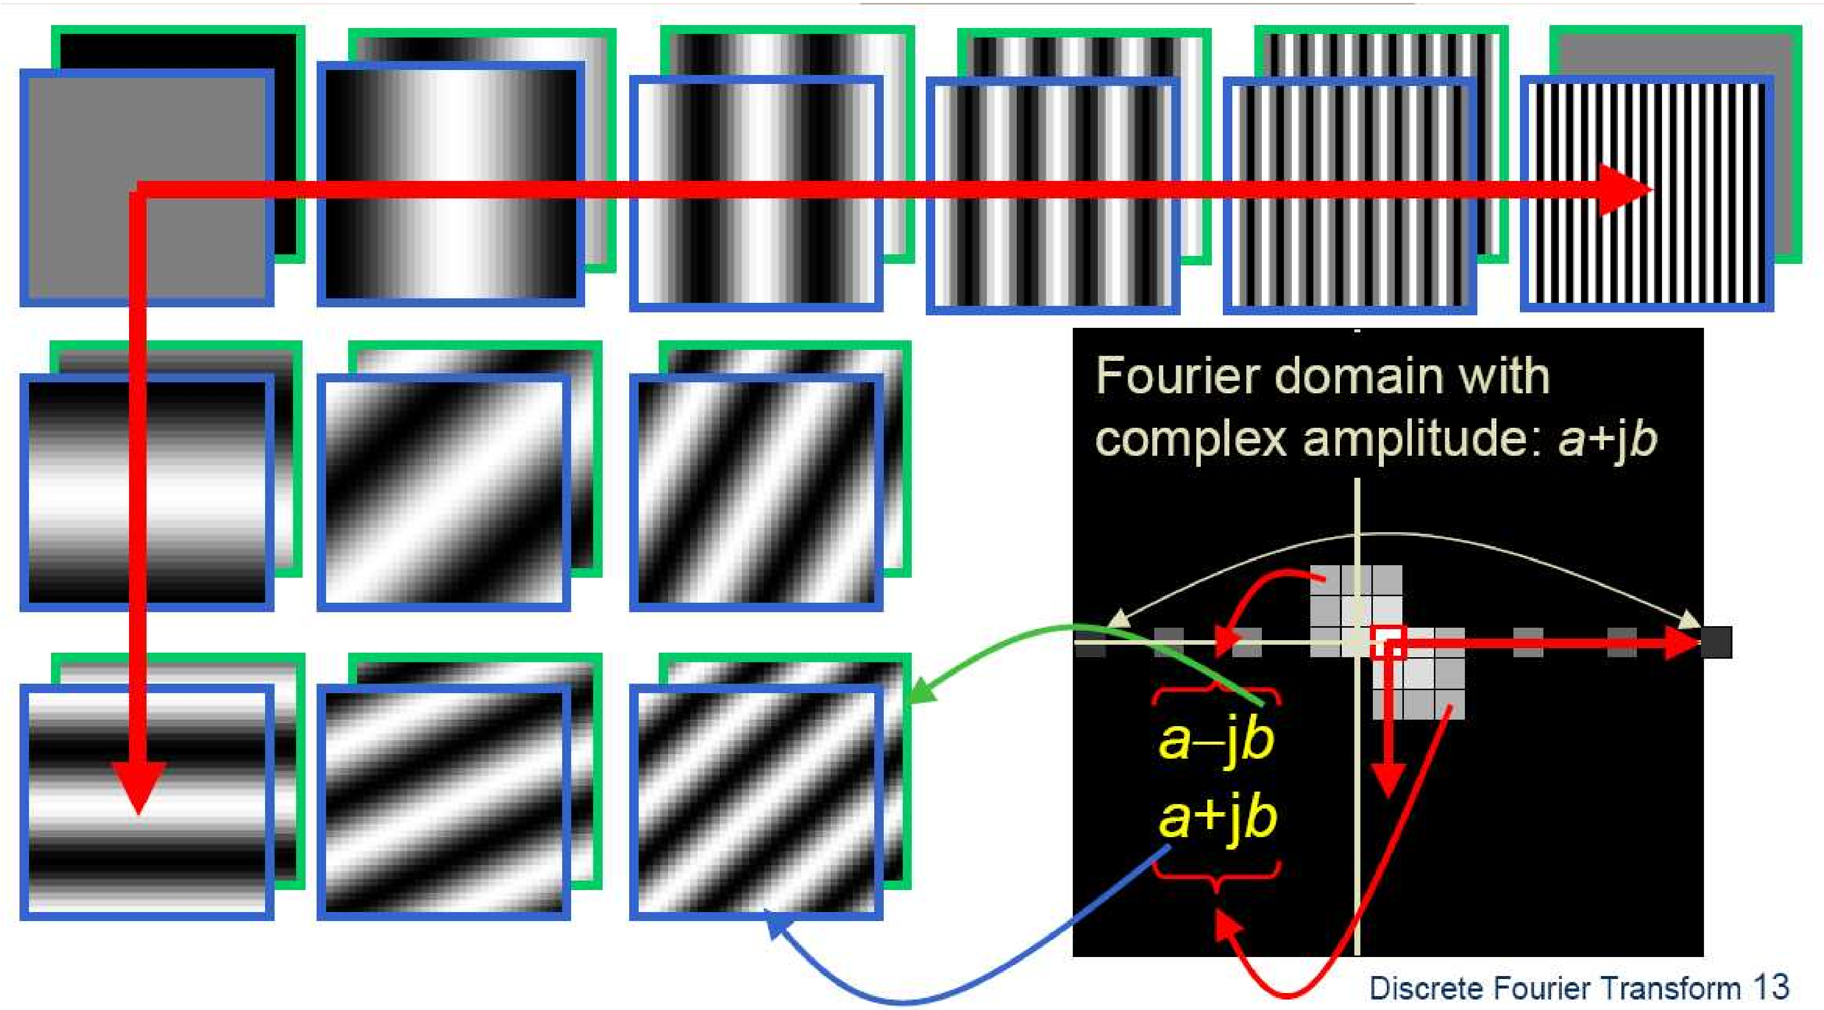

in Matlab, check out: imagesc(log(abs(fftshift(fft2(im)))));

Fourier analysis in images

Intensity Image

Fourier Image

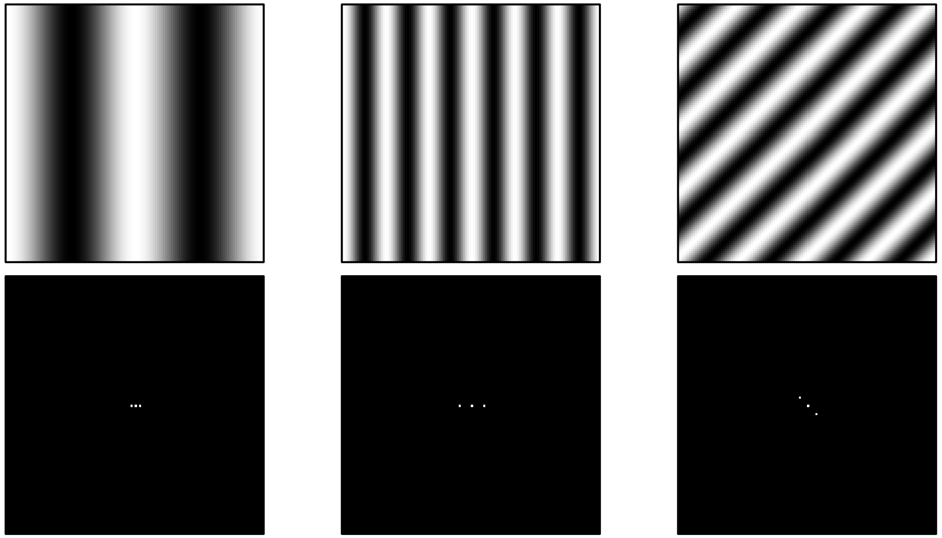

http://sharp.bu.edu/~slehar/fourier/fourier.html#filtering

Signals can be composed

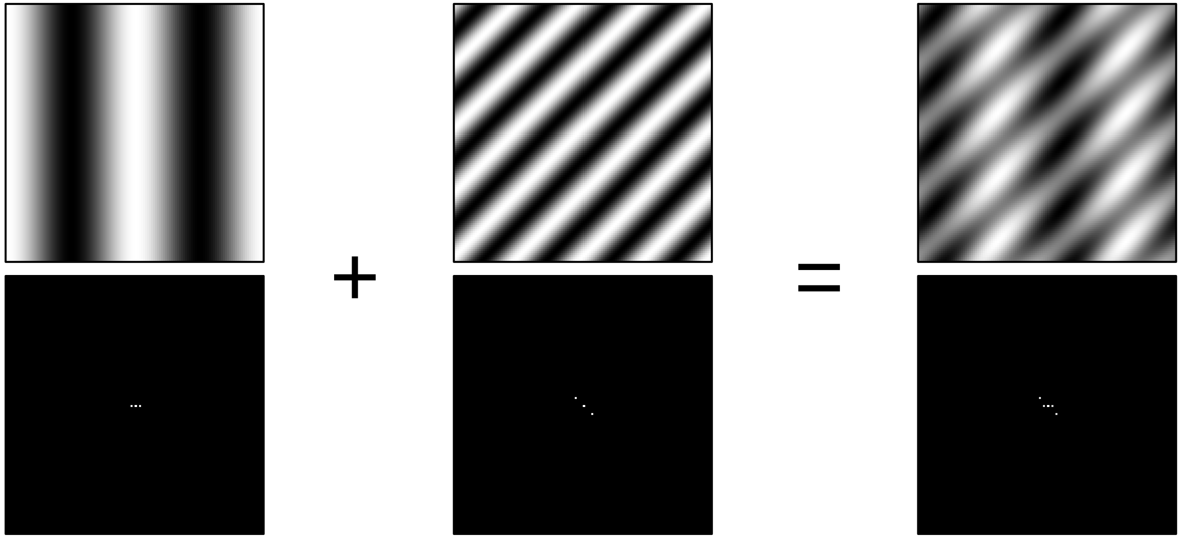

http://sharp.bu.edu/~slehar/fourier/fourier.html#filtering
More: http://www.cs.unm.edu/~brayer/vision/fourier.html

Man-made Scene

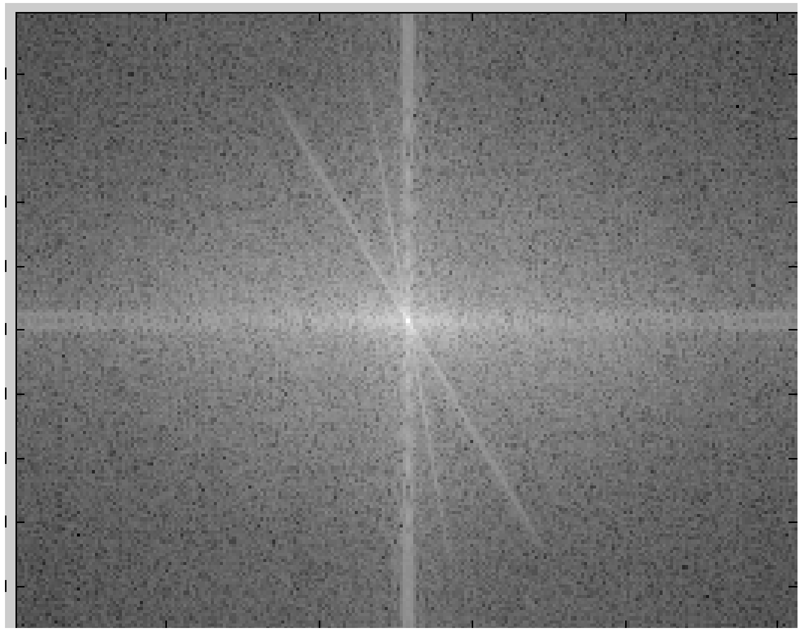

Can change spectrum, then reconstruct

Low and High Pass filtering

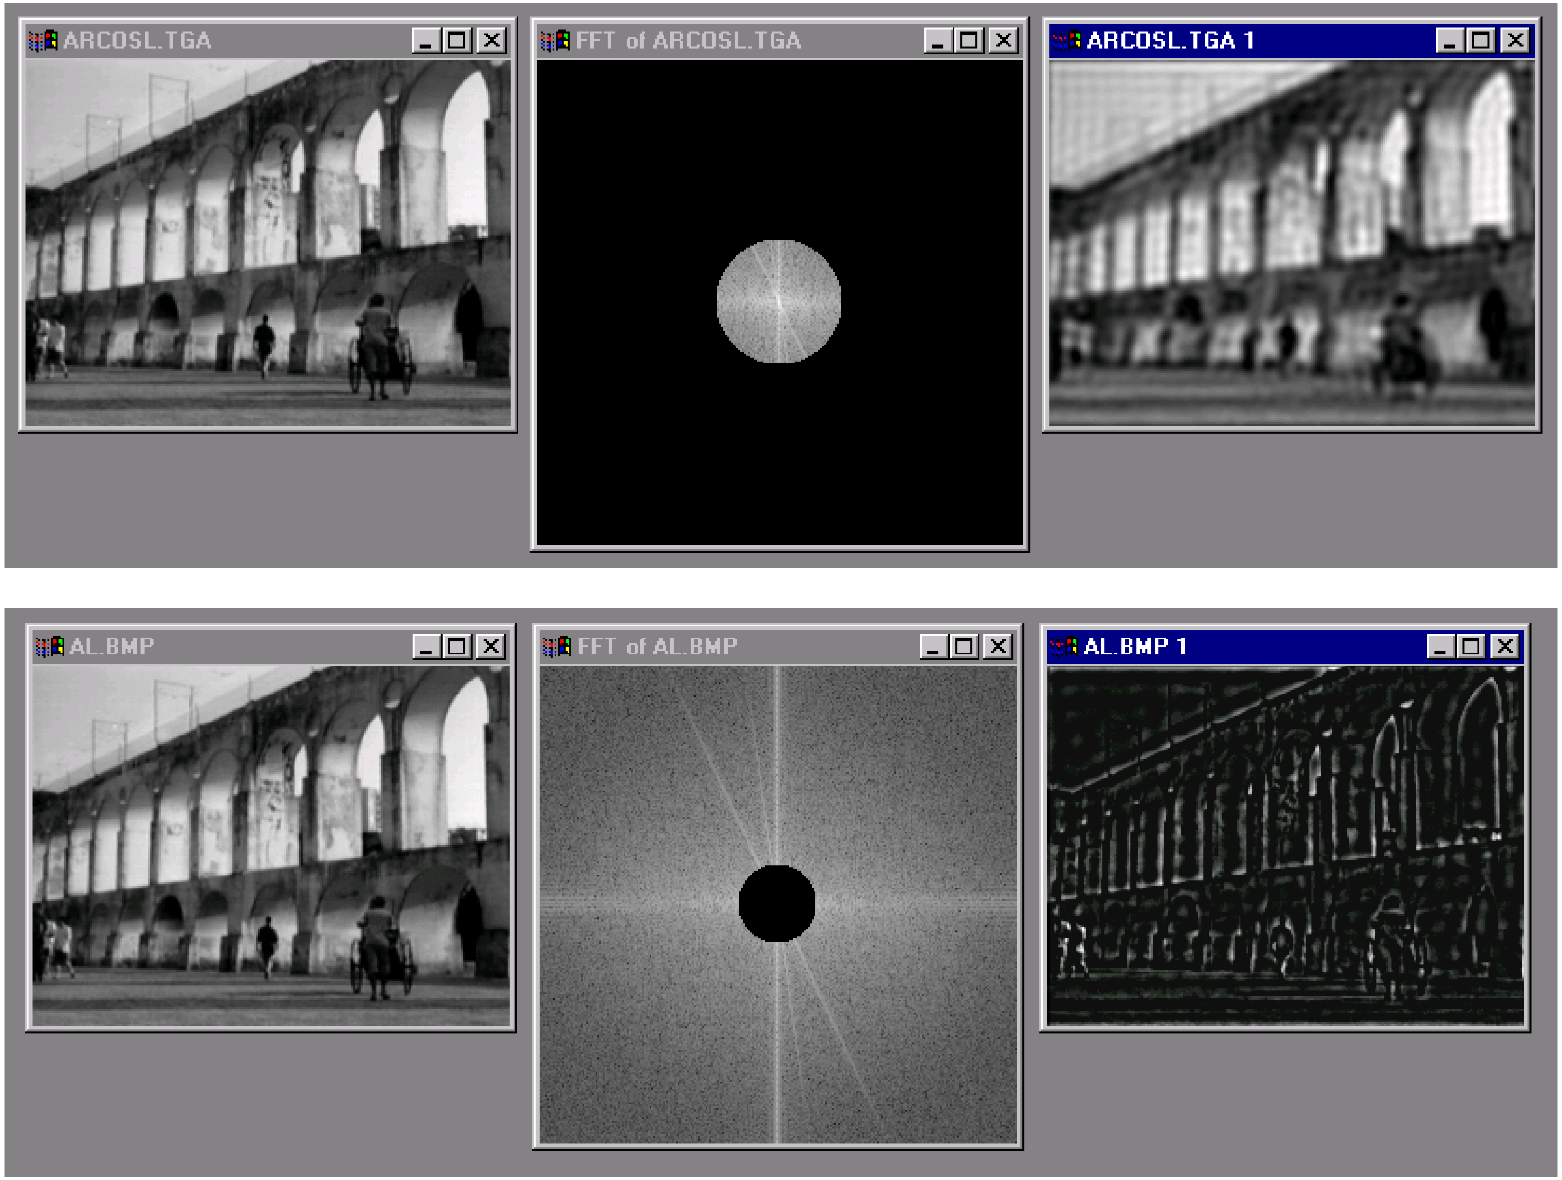

The Convolution Theorem

The greatest thing since sliced (banana) bread!

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

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

- The inverse Fourier transform of the product of two Fourier transforms is the convolution of the two inverse Fourier transforms

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

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

2D convolution theorem example

$f(x,y)$

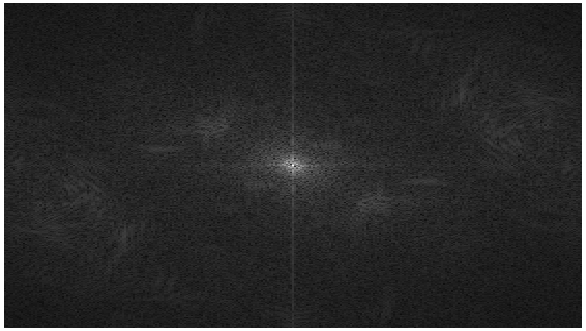

$|F(s_x,s_y)|$

$*$ $\times$

$h(x,y)$

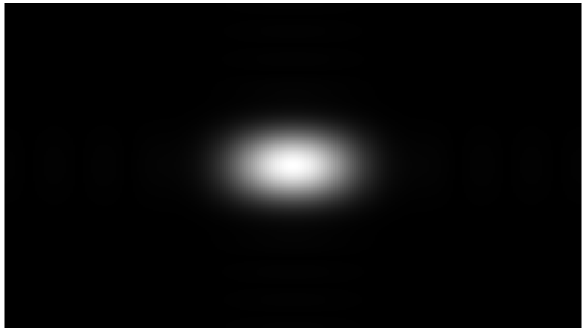

$|H(s_x,s_y)|$

$\Downarrow$ $\Downarrow$

$g(x,y)$

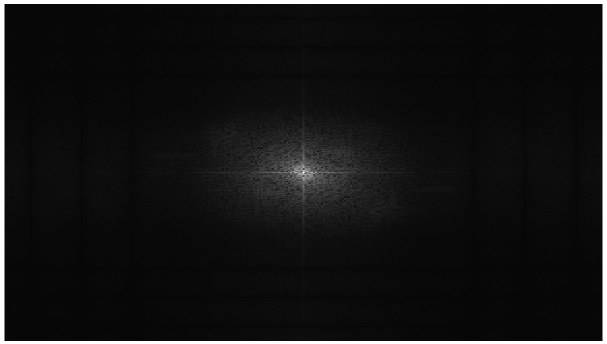

$|G(s_x,s_y)|$

Filtering

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

Gaussian

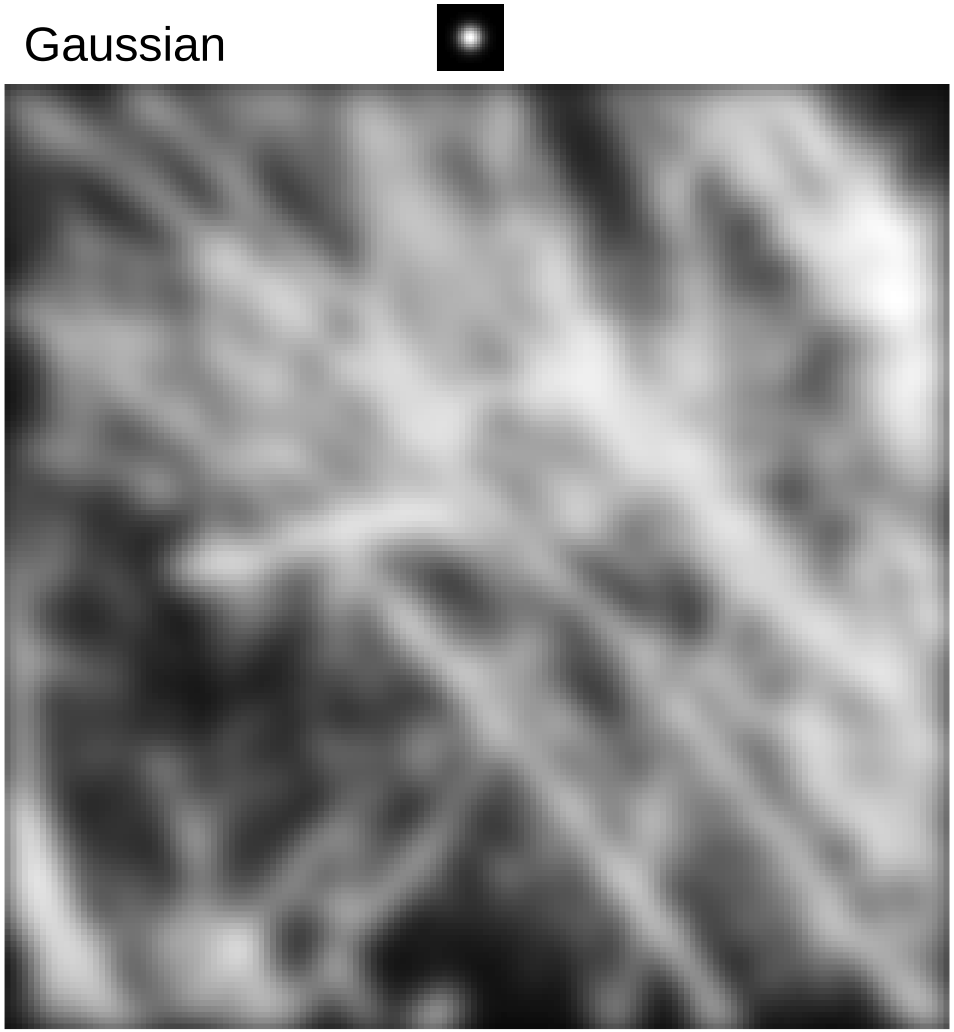

Box filter

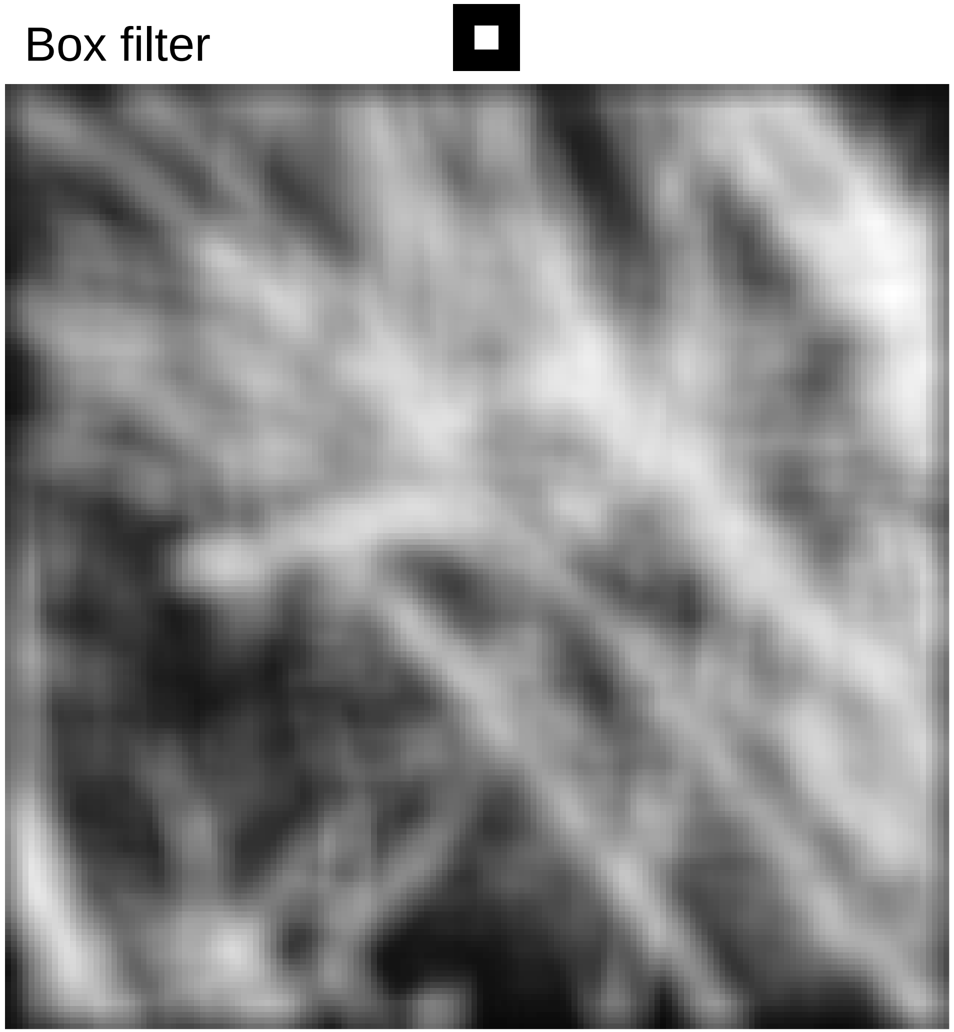

Gaussian

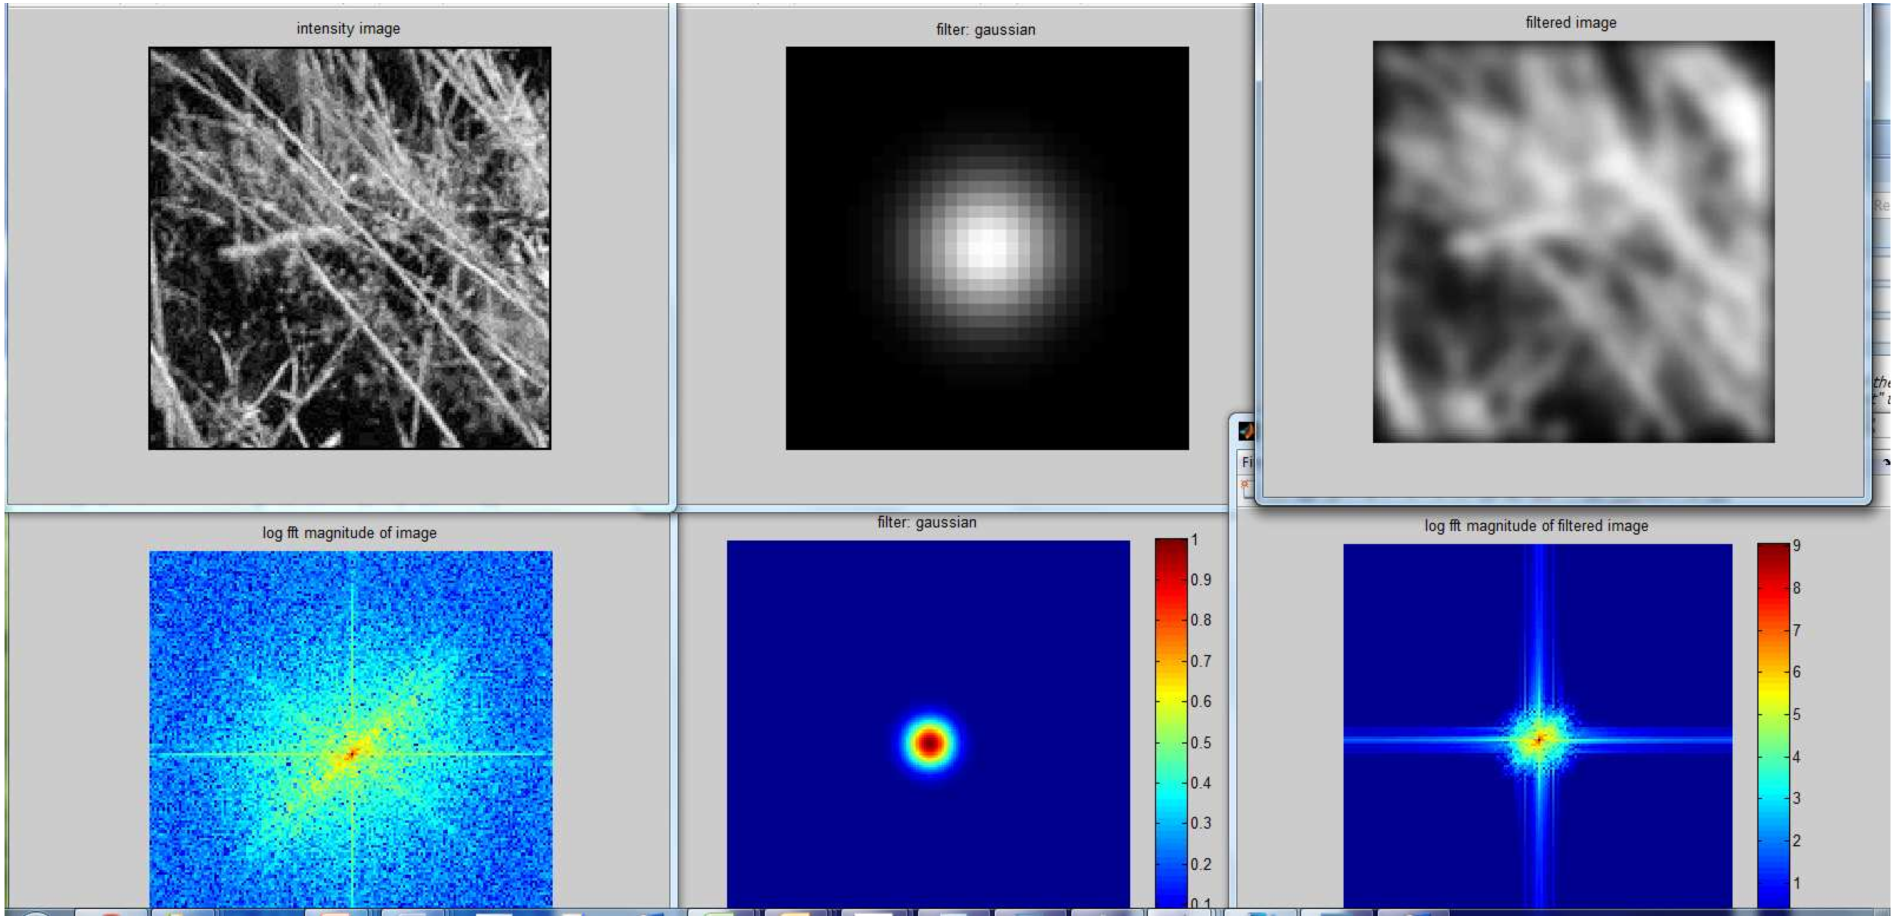

intensity image
filter: gaussian
filtered image
log fft magnitude of image
filter: gaussian
log fft magnitude of filtered image

Box Filter

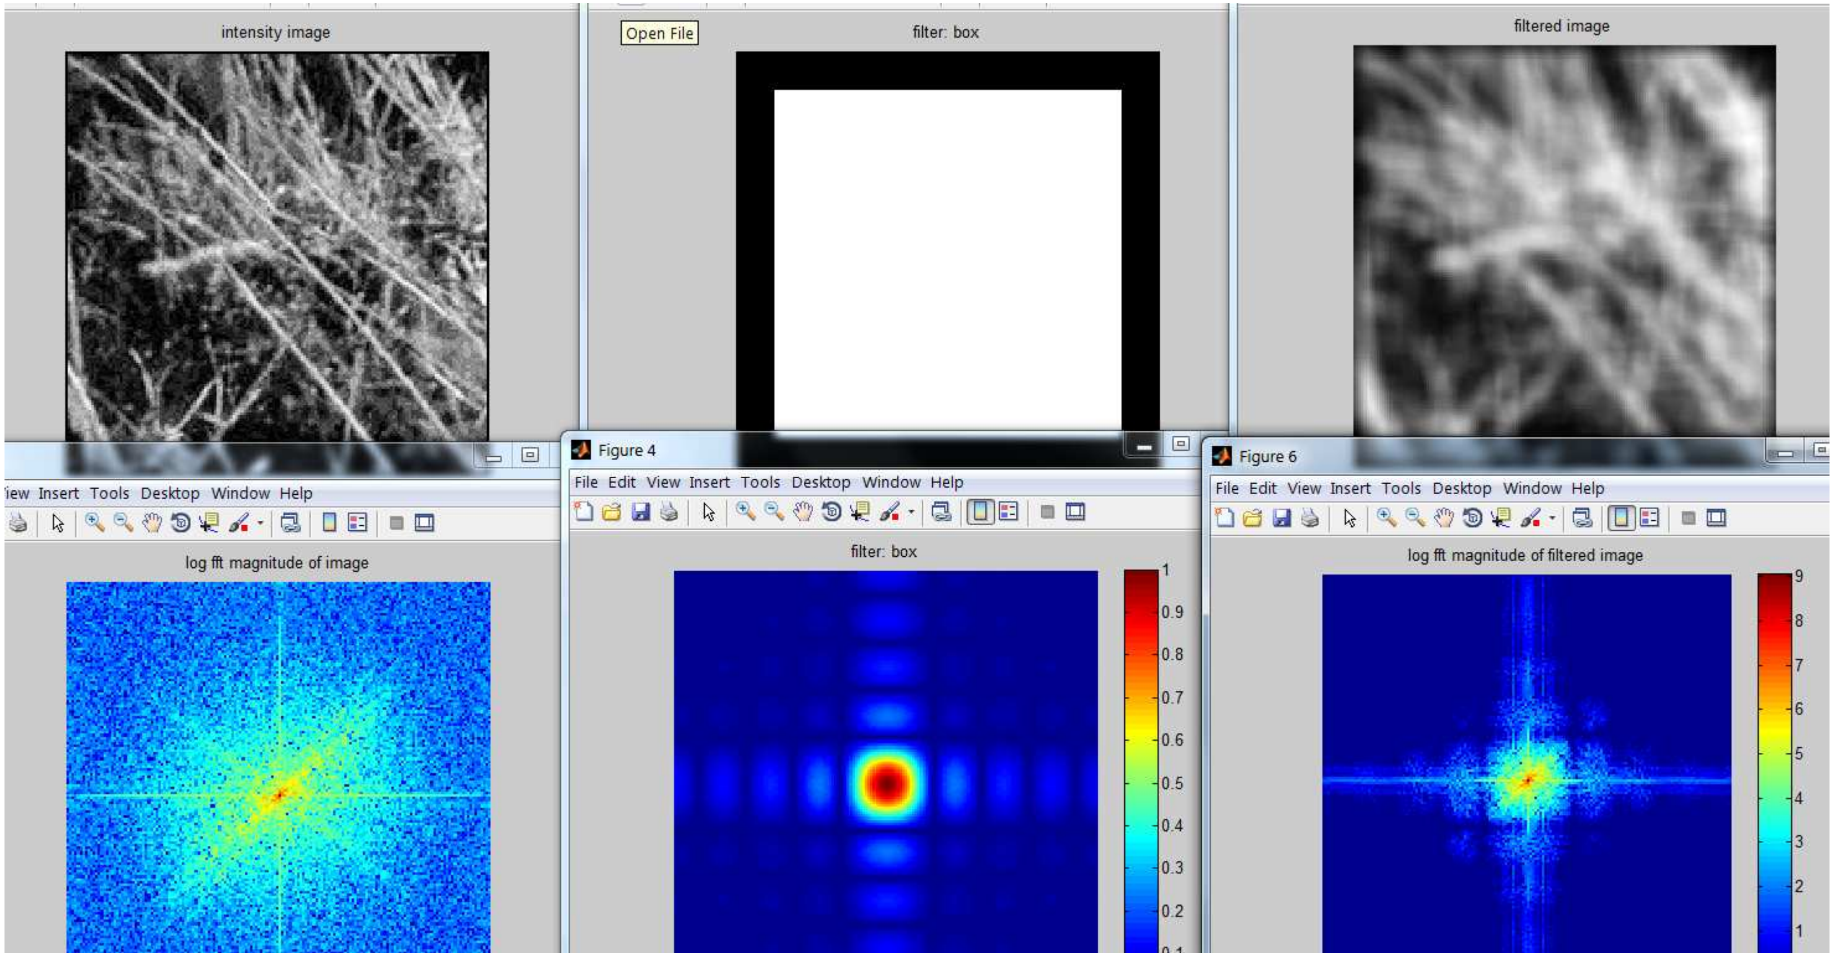
intensity image
Open File
filter: box
filtered image
Figure 4
Figure 6
File Edit View Insert Tools Desktop Window Help
log fft magnitude of image
filter: box
log fft magnitude of filtered image

Fourier Transform pairs

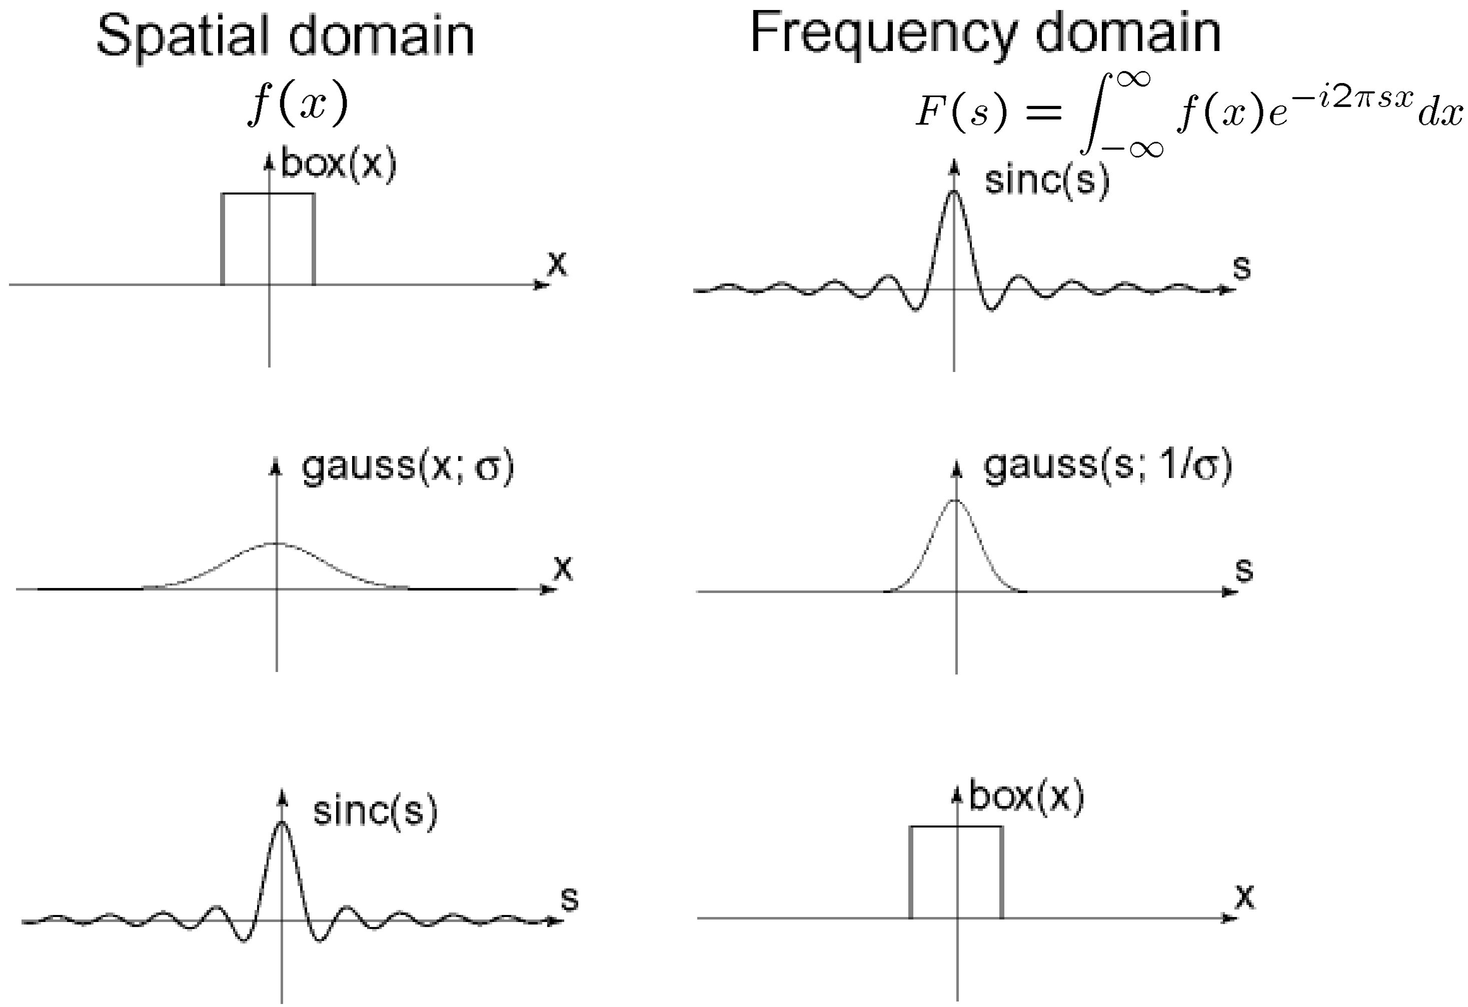

Low-pass, Band-pass, High-pass filters

low-pass:

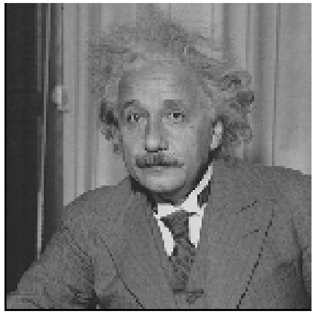 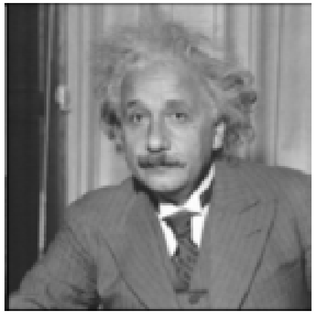 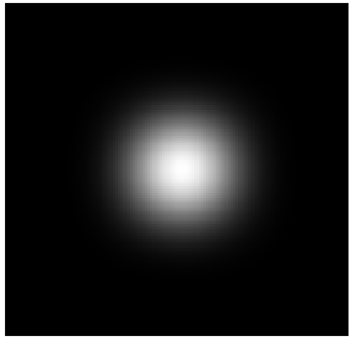

High-pass / band-pass:

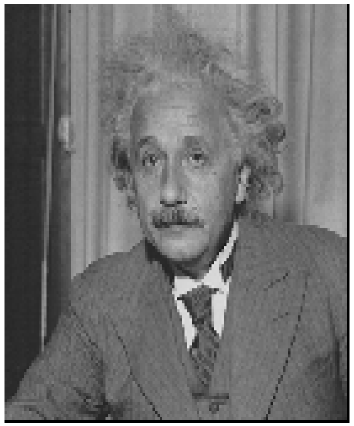  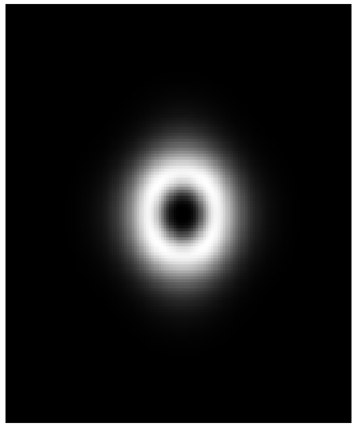

Edges in images

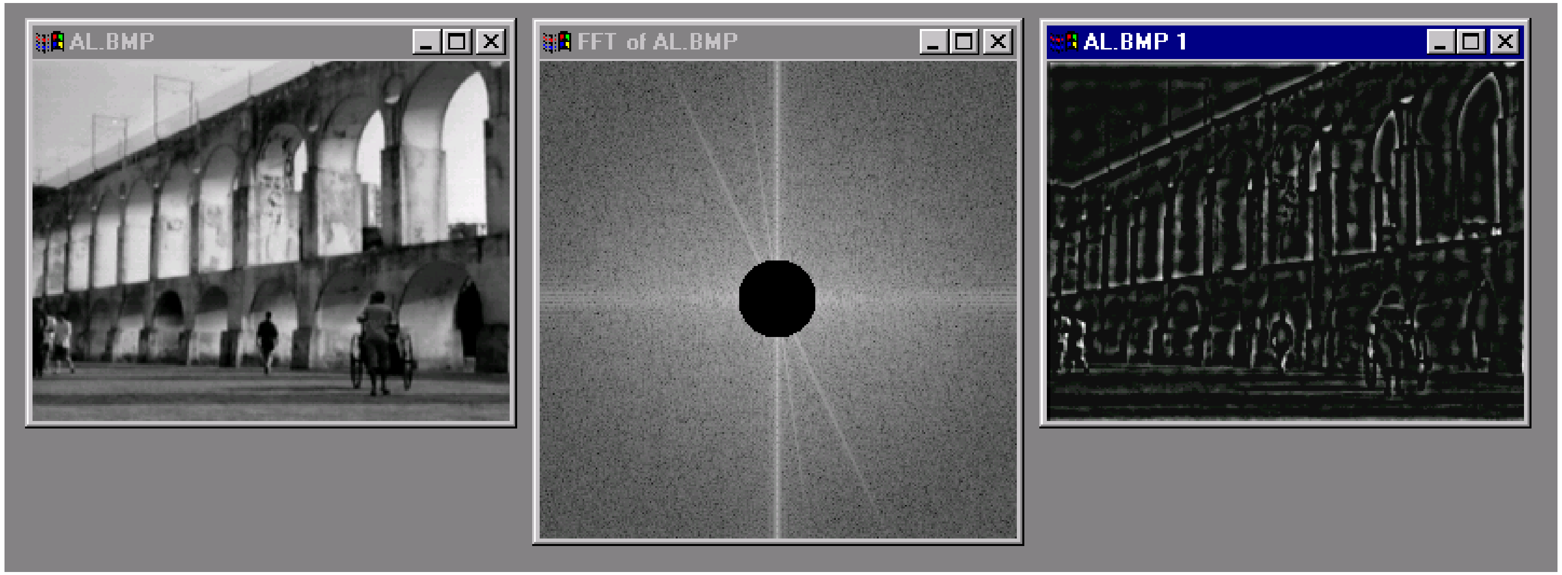

What does blurring take away?

original

What does blurring take away?

smoothed (5x5 Gaussian)

High-Pass filter

smoothed – original

Band-pass filtering

Gaussian Pyramid (low-pass images)

Laplacian Pyramid

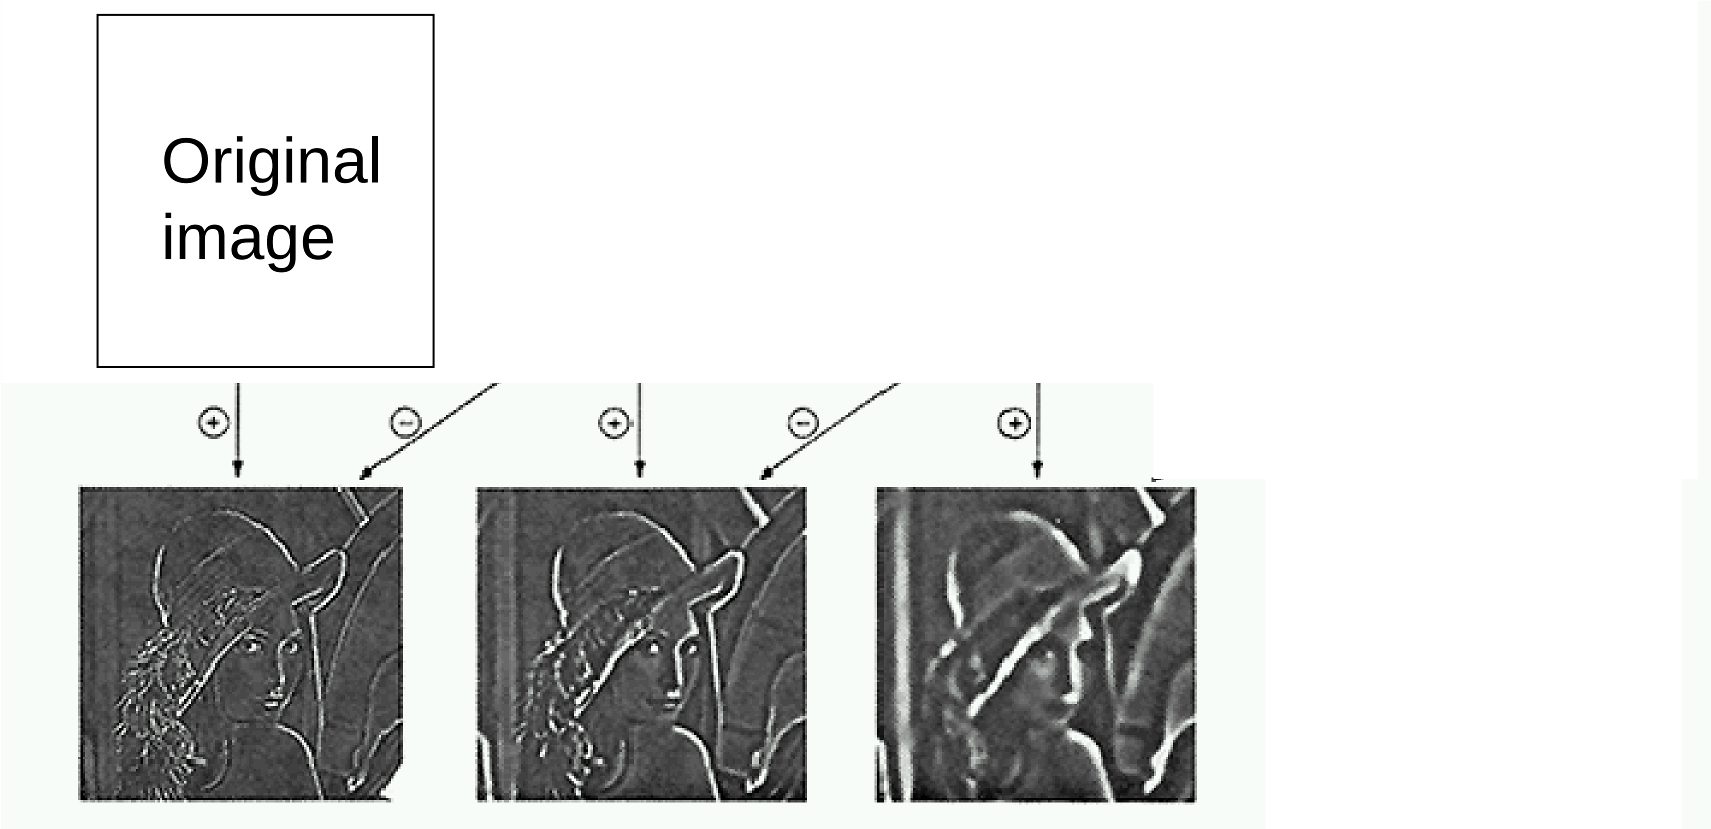

How can we reconstruct (collapse) this pyramid into the original image?

Why Laplacian?

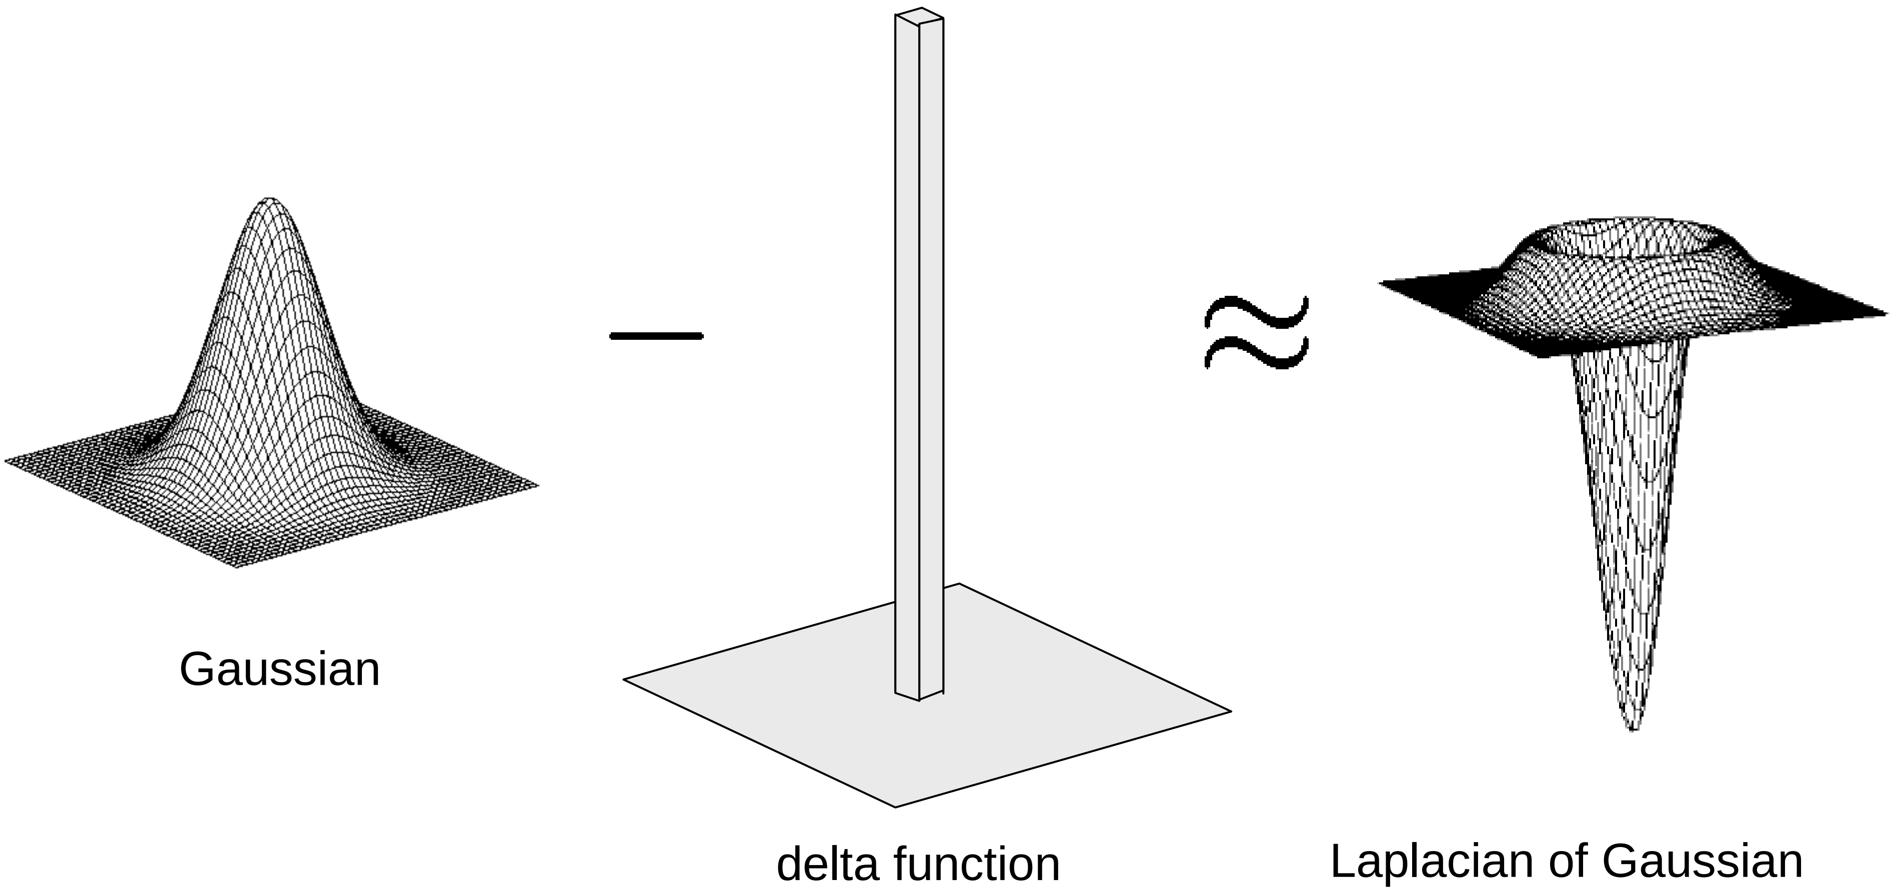

Project 1g: Hybrid Images

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

Gaussian Filter!

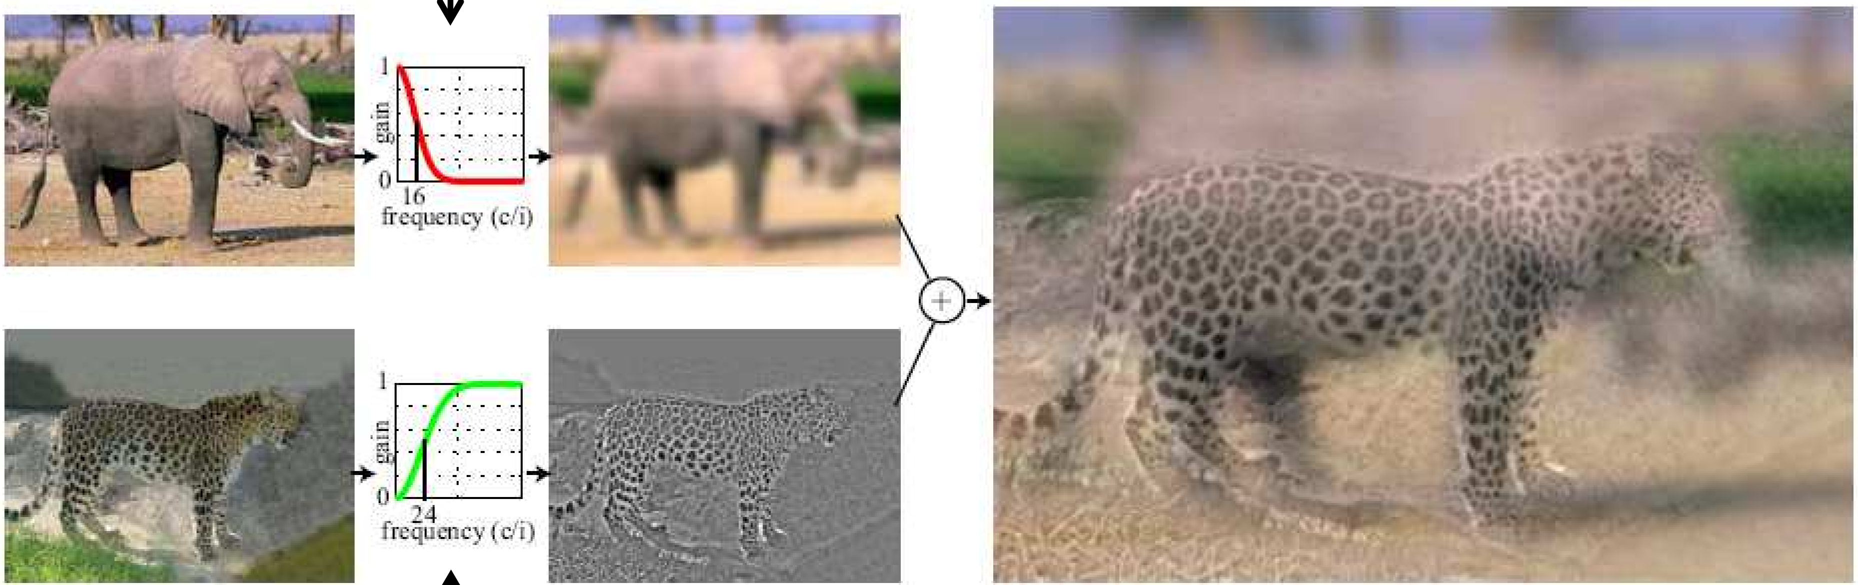

Laplacian Filter!

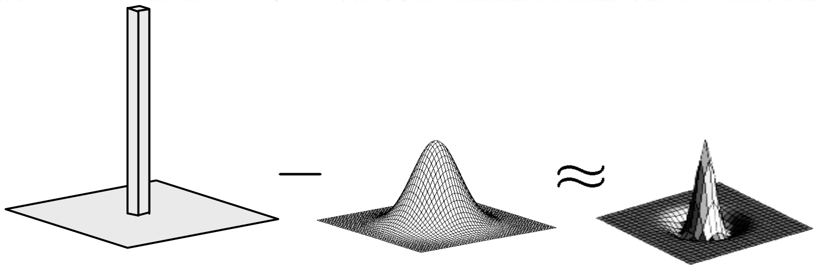

unit impulse — Gaussian ≈ Laplacian of Gaussian

Project Instructions:

http://www.cs.illinois.edu/class/fa10/cs498dwh/projects/hybrid/ComputationalPhotography_ProjectHybrid.html

Clues from Human Perception

Early processing in humans filters for various orientations and scales of frequency

Perceptual cues in the mid frequencies dominate perception

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

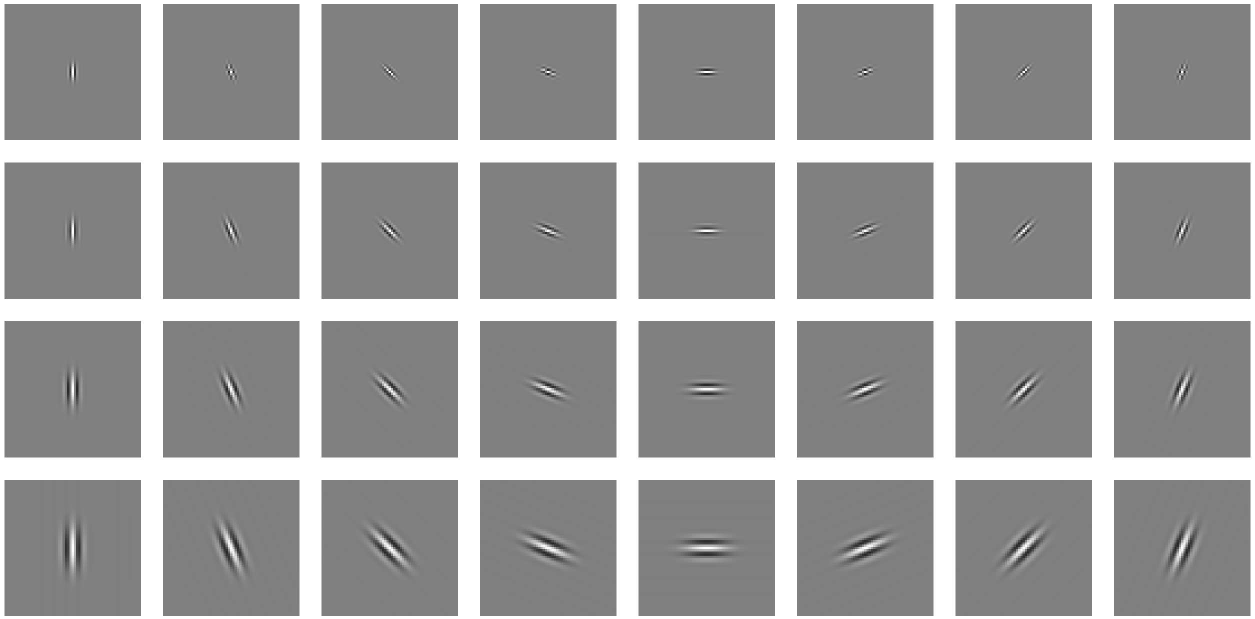

Early Visual Processing: Multi-scale edge and blob filters

Frequency Domain and Perception

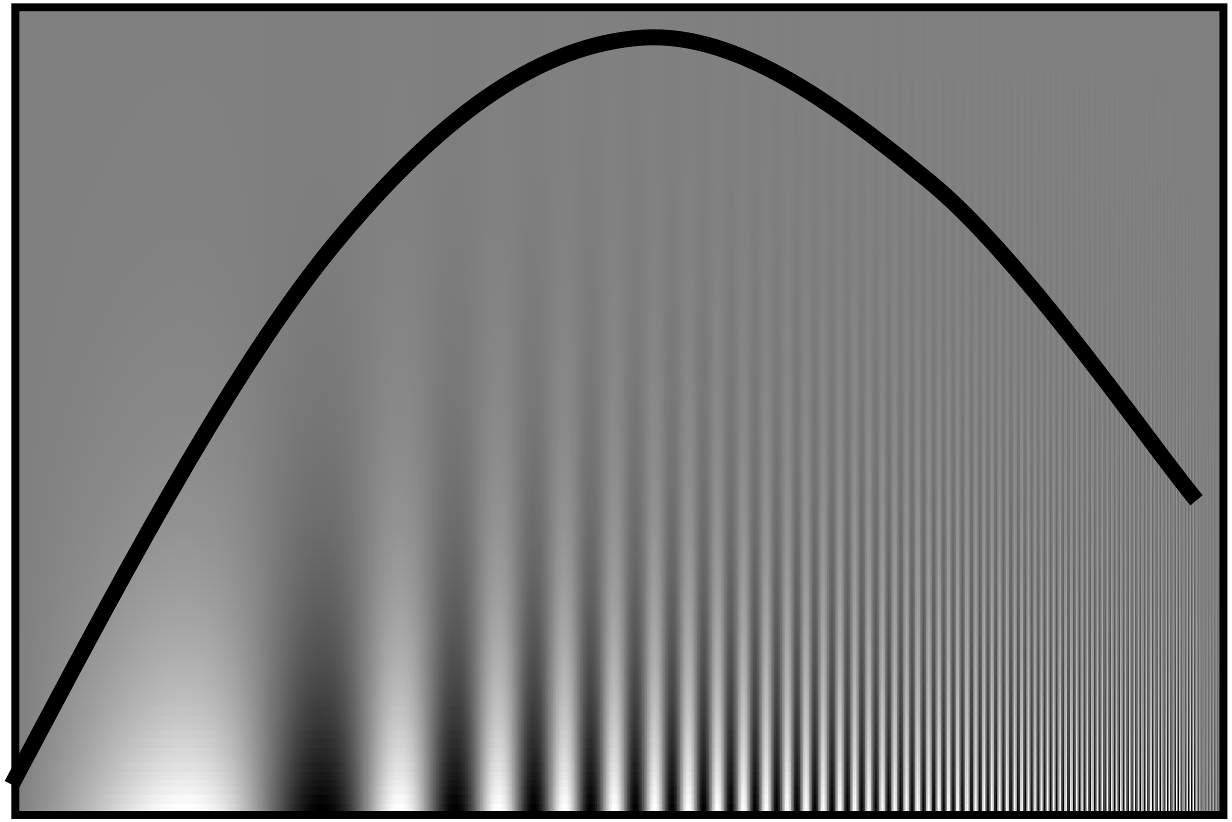

Campbell-Robson contrast sensitivity curve

Unsharp Masking

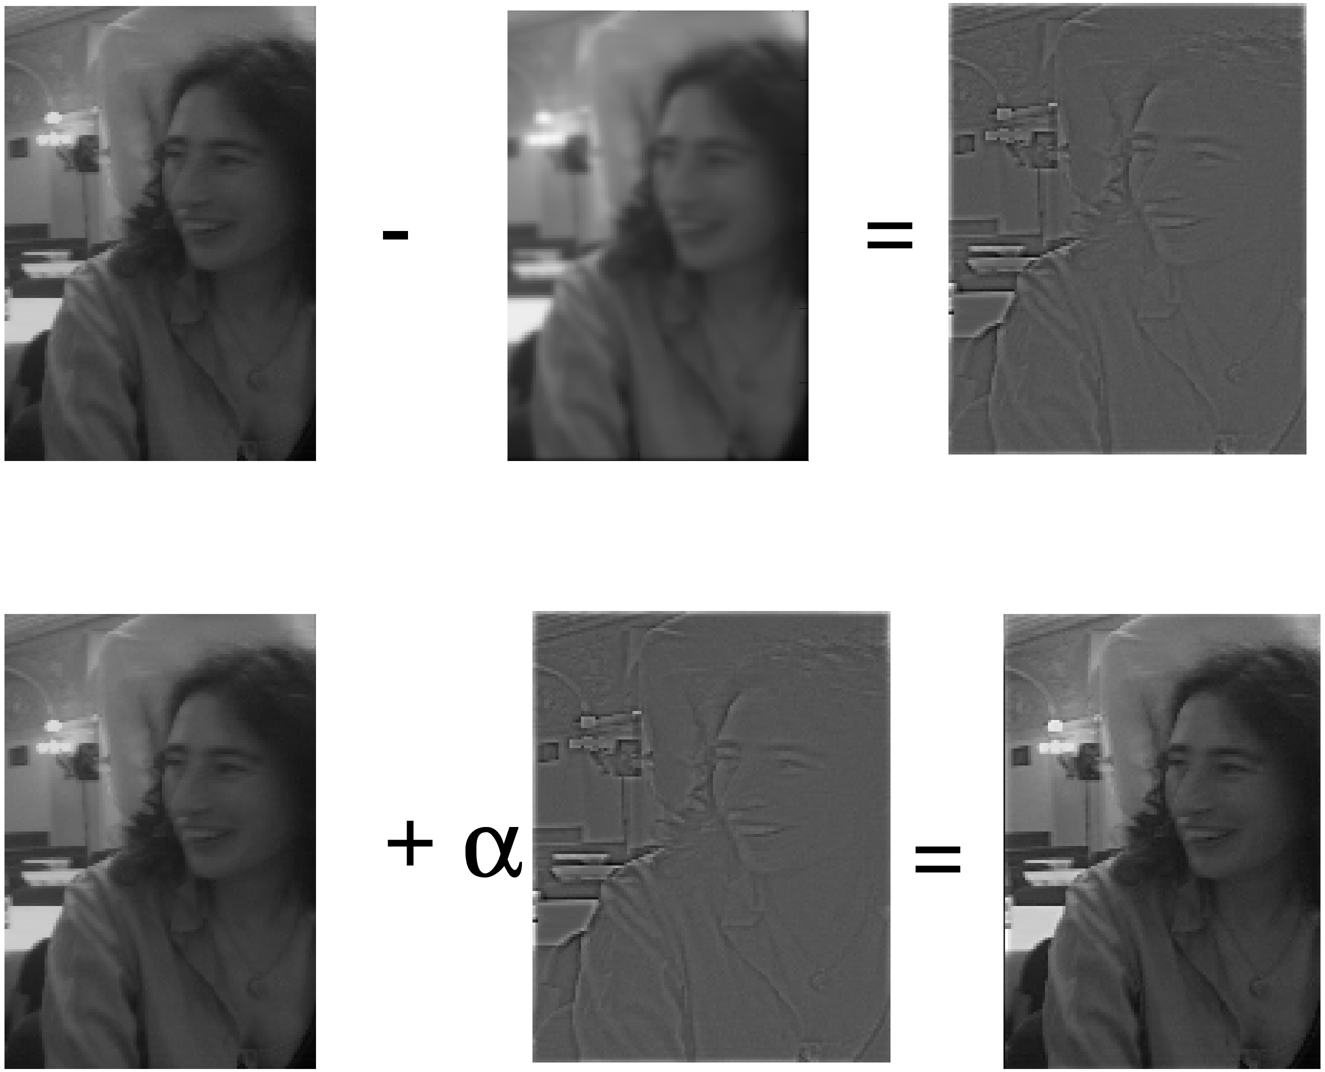

Freq. Perception Depends on Color

R G B

Lossy Image Compression (JPEG)

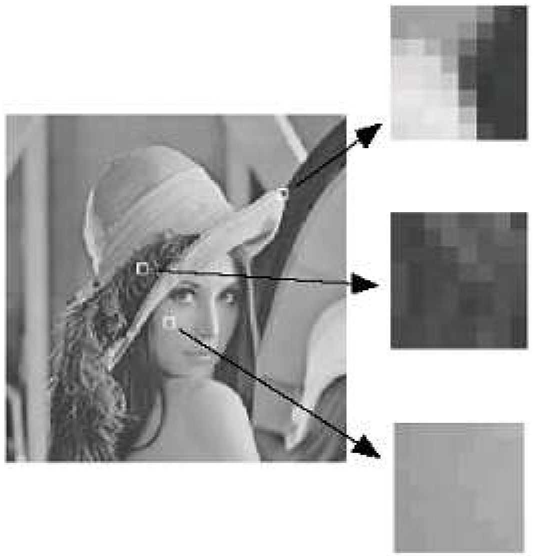

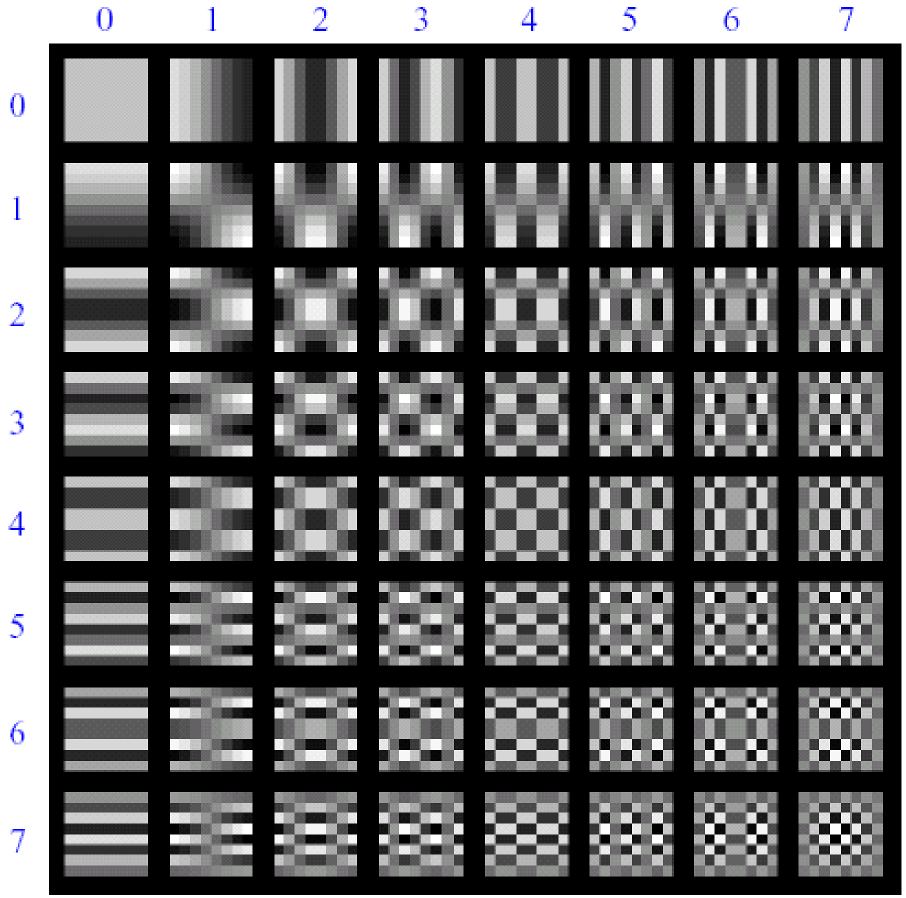

Block-based Discrete Cosine Transform (DCT)

Using DCT in JPEG

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

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

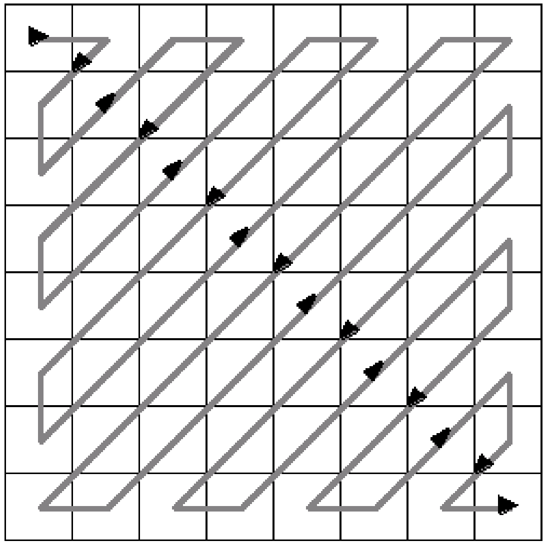

Image compression using DCT

Quantize

- More coarsely for high frequencies (which also tend to have smaller values)
- Many quantized high frequency values will be zero

Encode

- Can decode with inverse dct

Filter responses

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

Quantized values

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

Quantization table

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

JPEG Compression Summary

Subsample color by factor of 2

- People have bad resolution for color

Split into blocks (8x8, typically), subtract 128

For each block

a. Compute DCT coefficients for
b. Coarsely quantize
   - Many high frequency components will become zero
c. Encode (e.g., with Huffman coding)

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

Block size in JPEG

Block size

- small block
  - faster
  - correlation exists between neighboring pixels
- large block
  - better compression in smooth regions
- It’s 8x8 in standard JPEG

JPEG compression comparison

89k

12k

Image gradient

The gradient of an image:

$$\nabla f = \left[\frac{\partial f}{\partial x}, \frac{\partial f}{\partial y}\right]$$

The gradient points in the direction of most rapid change in intensity

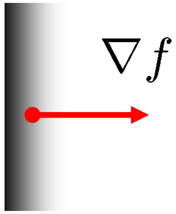

$$\nabla f = \left[\frac{\partial f}{\partial x}, 0\right]$$

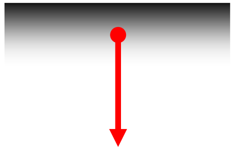

$$\nabla f = \left[0, \frac{\partial f}{\partial y}\right]$$

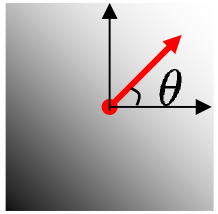

$$\nabla f = \left[\frac{\partial f}{\partial x}, \frac{\partial f}{\partial y}\right]$$

The gradient direction is given by:

$$\theta = \tan^{-1}\left(\frac{\partial f}{\partial y} / \frac{\partial f}{\partial x}\right)$$

- how does this relate to the direction of the edge?

The *edge strength* is given by the gradient magnitude

$$||\nabla f|| = \sqrt{\left(\frac{\partial f}{\partial x}\right)^2 + \left(\frac{\partial f}{\partial y}\right)^2}$$

Effects of noise

Consider a single row or column of the image

- Plotting intensity as a function of position gives a signal

$f(x)$

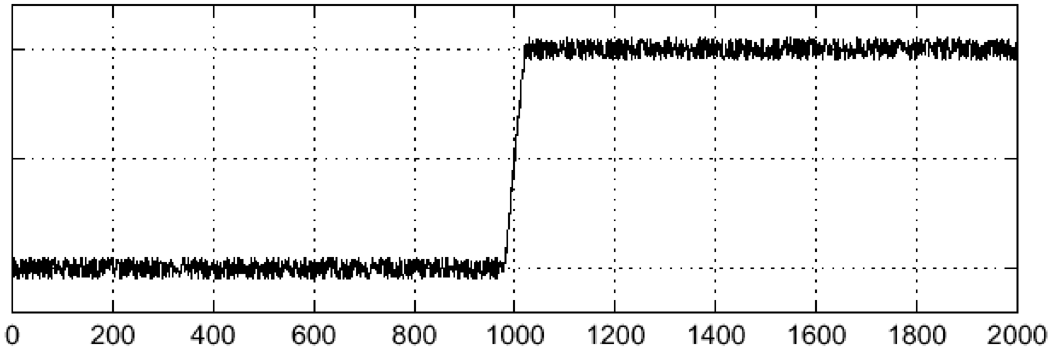

How to compute a derivative?

$\frac{d}{dx}f(x)$

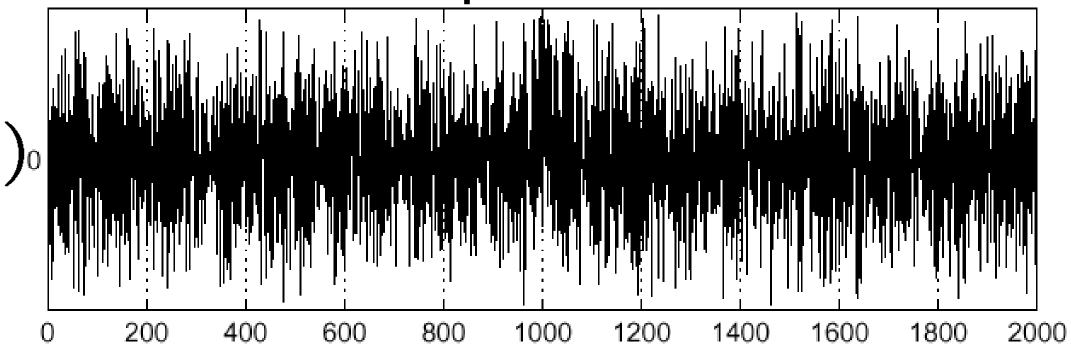

Where is the edge?

Solution: smooth first

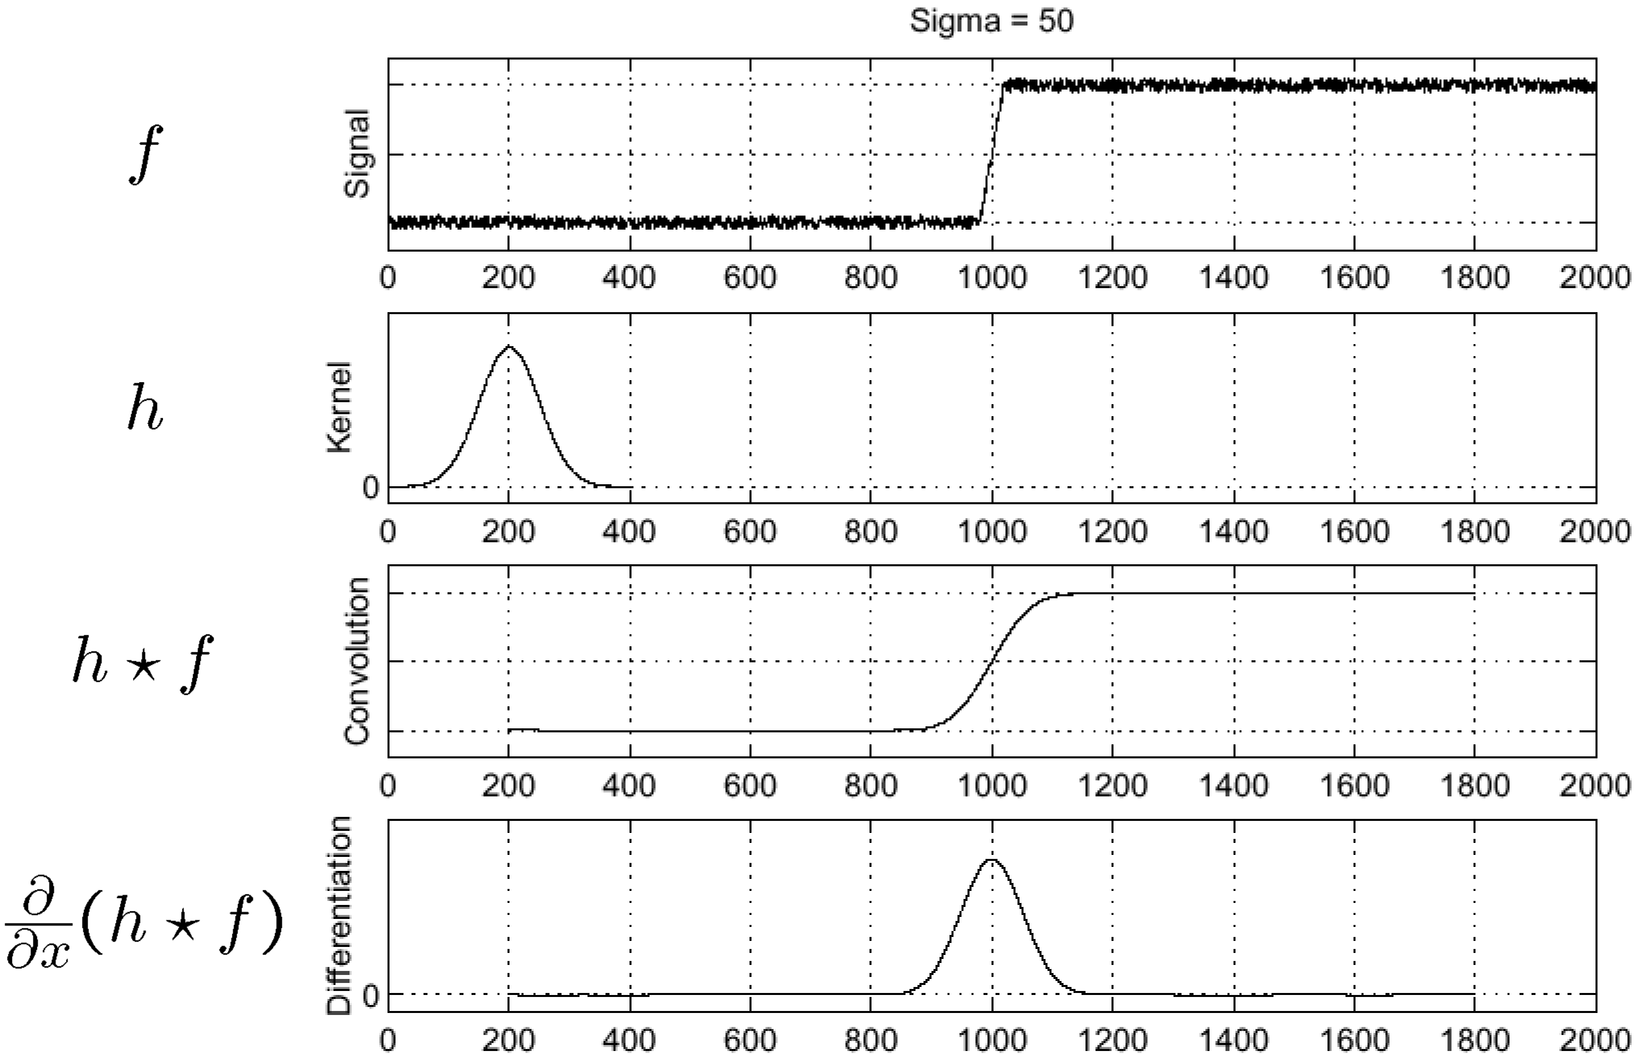

Where is the edge? Look for peaks in $\frac{\partial}{\partial x}(h \star f)$

Derivative theorem of convolution

$$\frac{\partial}{\partial x}(h \star f) = (\frac{\partial}{\partial x}h) \star f$$

This saves us one operation:

f

$\frac{\partial}{\partial x}h$

$(\frac{\partial}{\partial x}h) \star f$

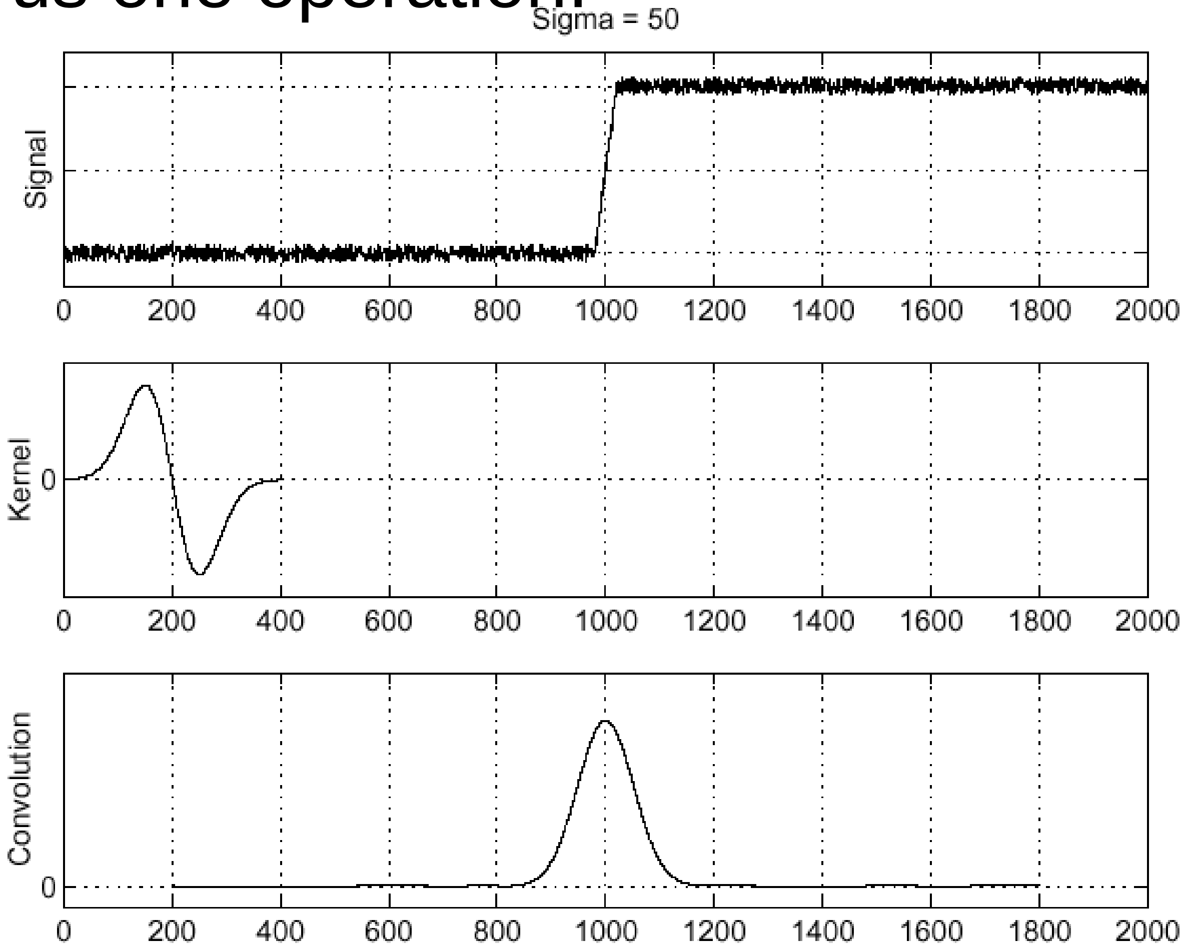

Laplacian of Gaussian

Consider $\frac{\partial^2}{\partial x^2}(h \star f)$

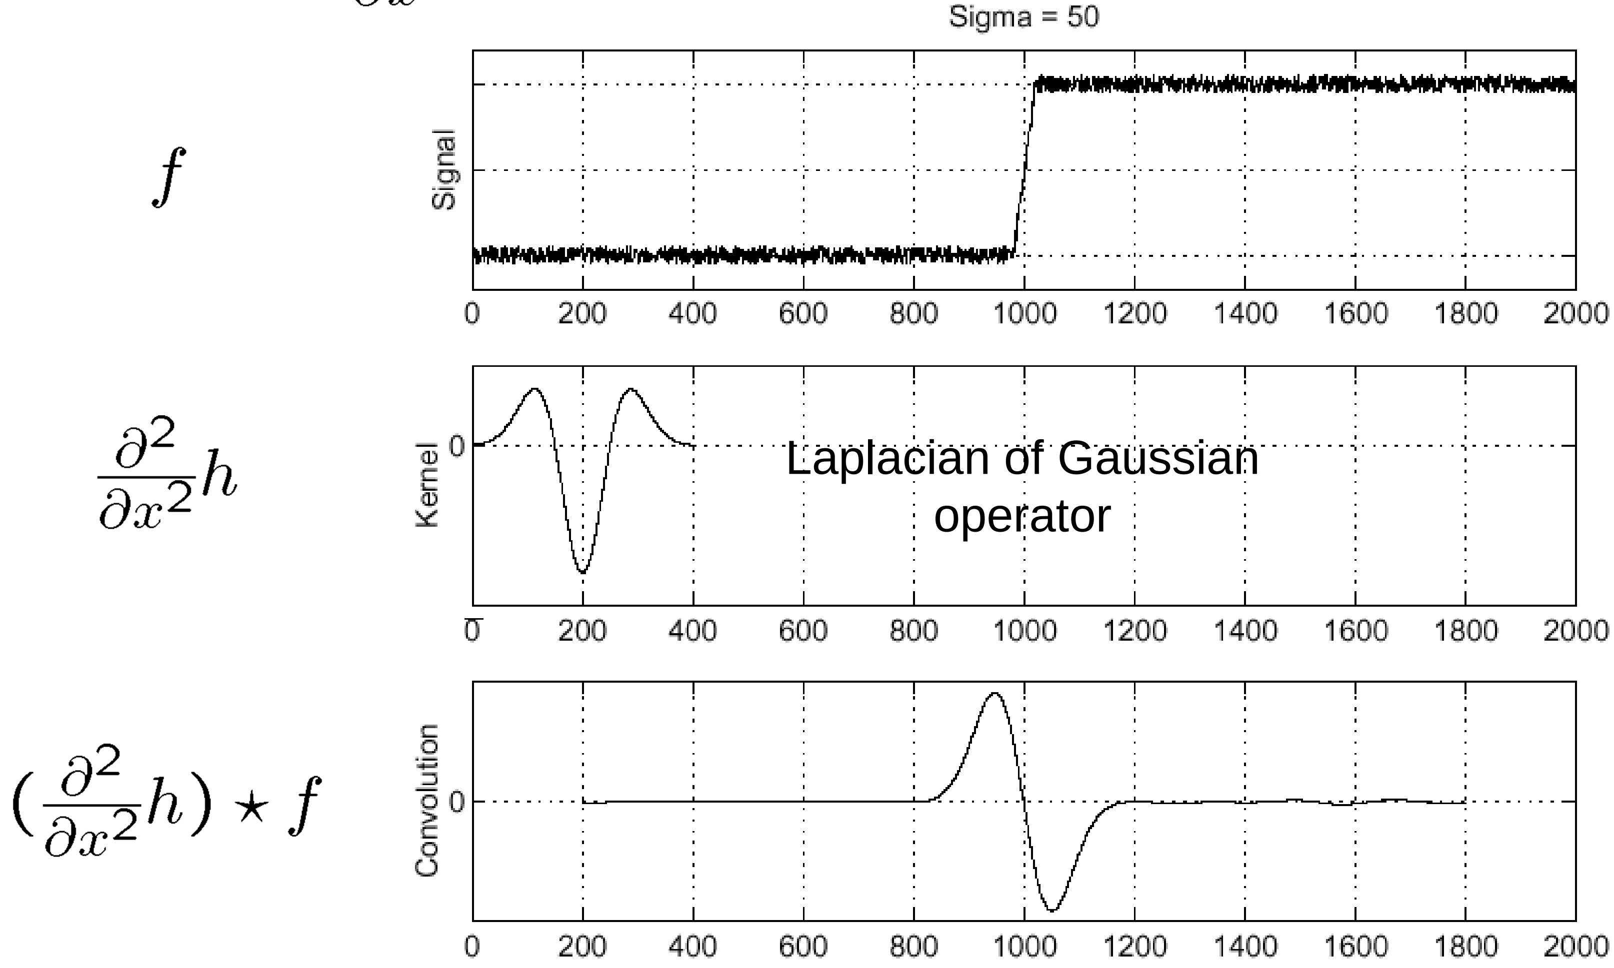

Where is the edge? Zero-crossings of bottom graph

2D edge detection filters

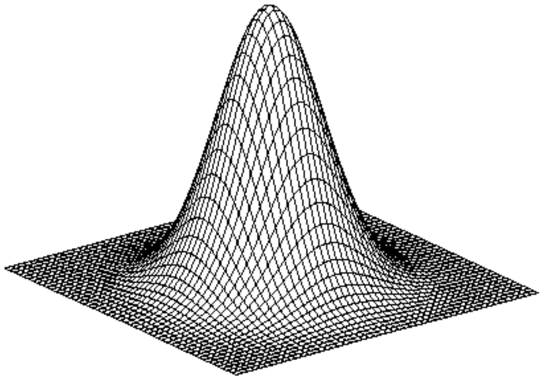

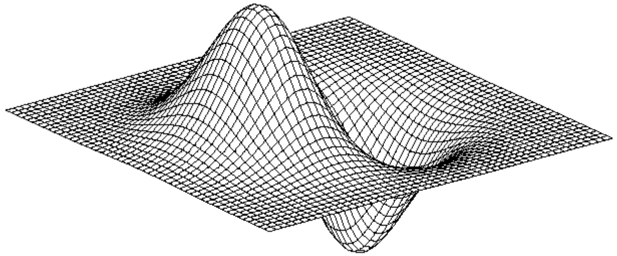

Laplacian of Gaussian

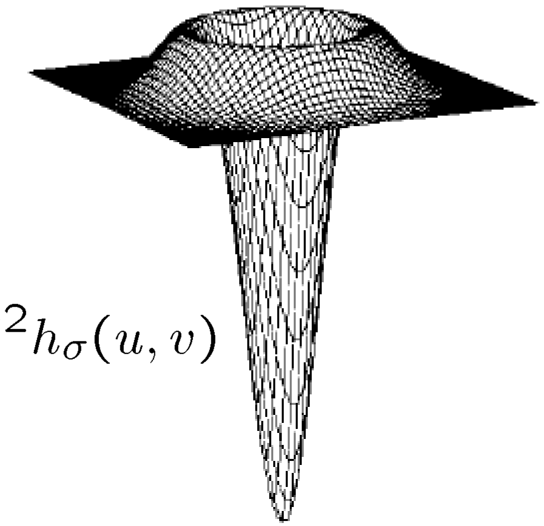

Gaussian

$$h_\sigma(u,v) = \frac{1}{2\pi\sigma^2} e^{-\frac{u^2+v^2}{2\sigma^2}}$$

derivative of Gaussian

$$\frac{\partial}{\partial x} h_\sigma(u,v)$$

$$\nabla^2 h_\sigma(u,v)$$

∇^2 is the **Laplacian** operator:

$$\nabla^2 f = \frac{\partial^2 f}{\partial x^2} + \frac{\partial^2 f}{\partial y^2}$$

Try this in MATLAB

```
g = fspecial('gaussian',15,2);
imagesc(g);  colormap(gray);
surfl(g)
gclown = conv2(clown,g,'same');
imagesc(conv2(clown,[-1 1],'same'));
imagesc(conv2(gclown,[-1 1],'same'));
dx = conv2(g,[-1 1],'same');
imagesc(conv2(clown,dx,'same'));
lg = fspecial('log',15,2);
lclown = conv2(clown,lg,'same');
imagesc(lclown)
imagesc(clown + .2*lclown)
```