Image Processing and Representations

Prepared by Behzad Sajadi

Borrowed from Frédo Durand's
Lectures at MIT

Image processing

- Filtering, Convolution, and our friend Joseph Fourier

What is an image?

- We can think of an **image** as a function, *f*,
- from R^2 to R:
  - $f(x, y)$ gives the **intensity** at position (x, y)
  - Realistically, we expect the image only to be defined over a rectangle, with a finite range:
    - $f: [a,b] \times [c,d] \rightarrow [0,1]$
- A color image is just three functions pasted together. We can write this as a "vector-valued" function:

$$f(x, y) = \begin{bmatrix} r(x, y) \\ g(x, y) \\ b(x, y) \end{bmatrix}$$

Images as functions

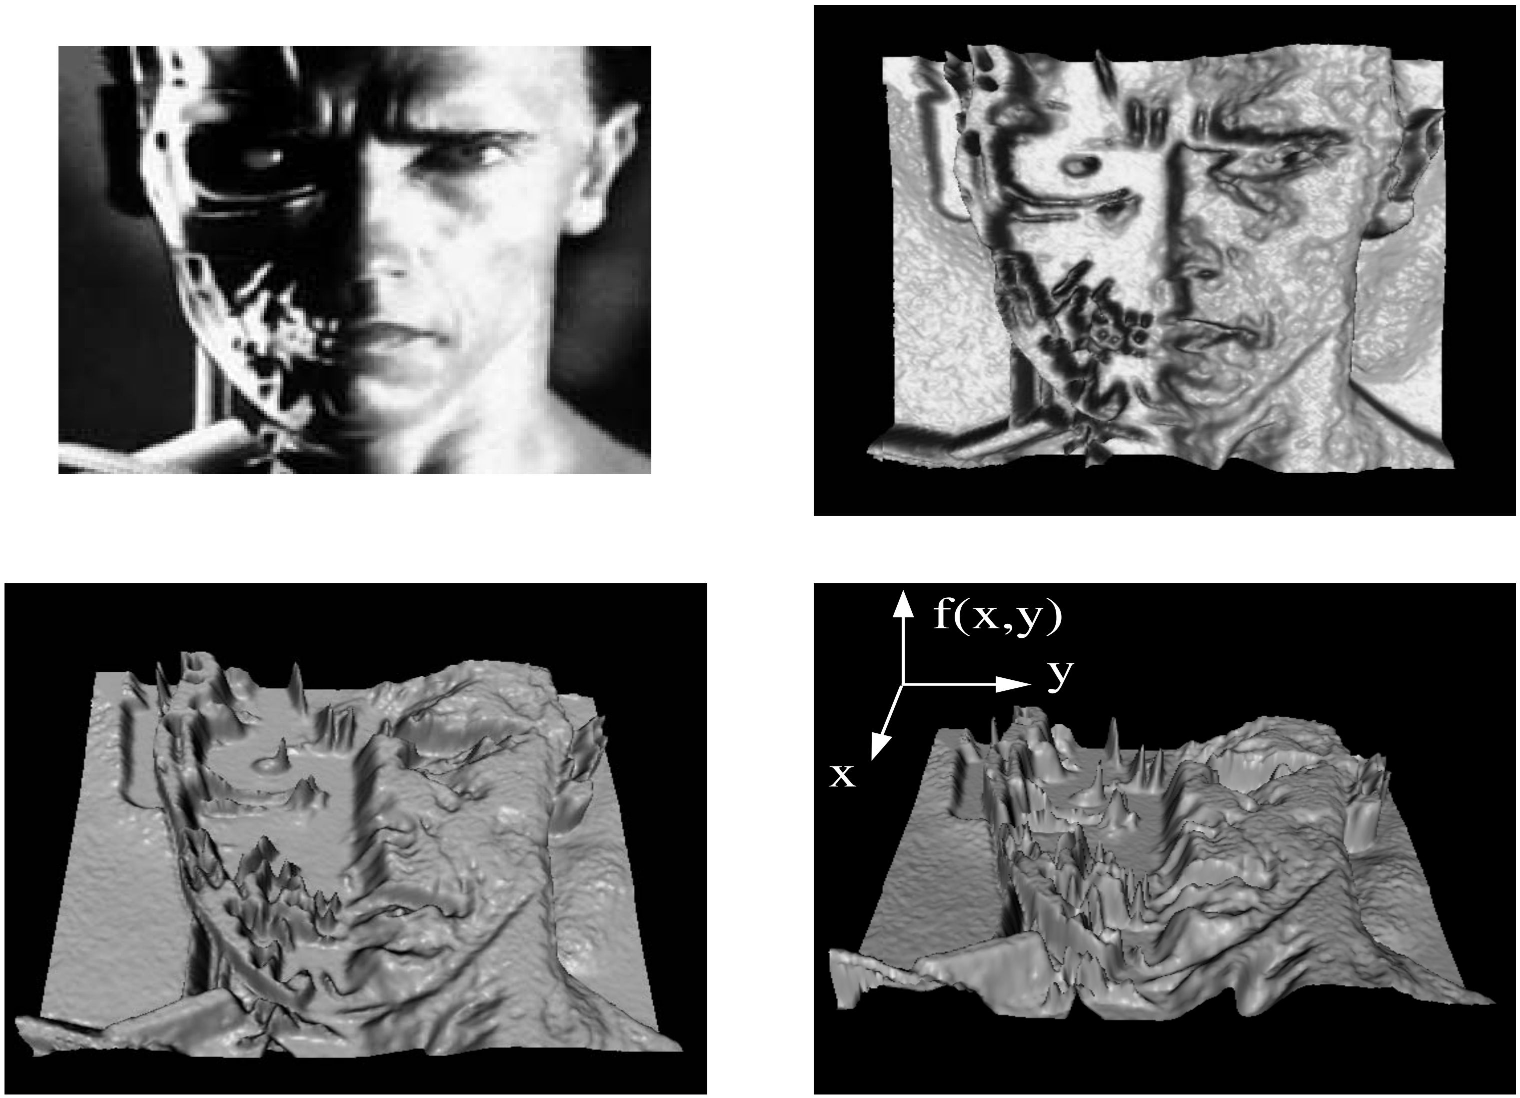

Image Processing

- image filtering: change ***range*** of image
  - $g(x) = h(f(x))$

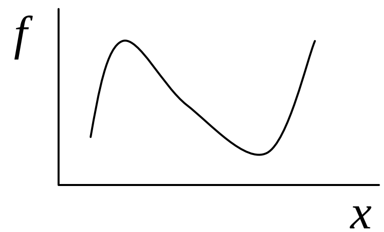

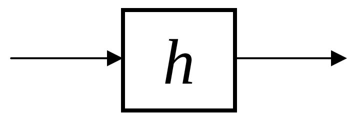

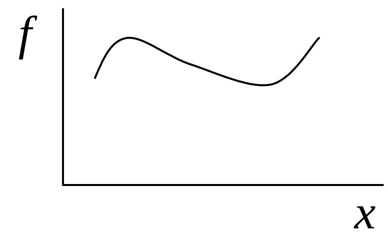

- image warping: change ***domain*** of image
  - $g(x) = f(h(x))$

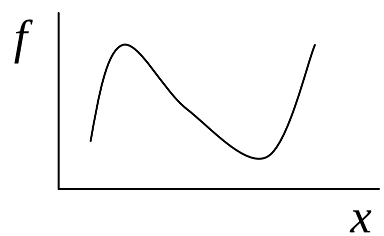

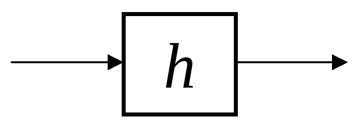

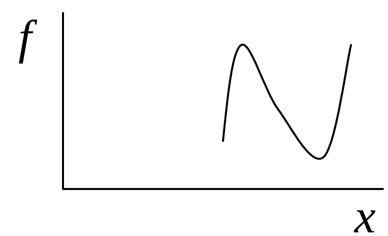

Image Processing

- image filtering: change ***range*** of image
  - *g(x) = h(f(x))*


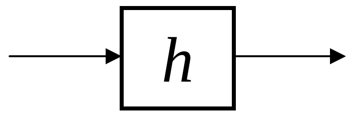

- image warping: change ***domain*** of image
  - *g(x) = f(h(x))*


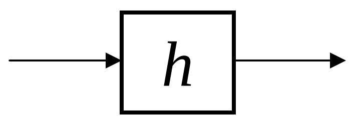

Point Processing

- The simplest kind of range transformations are those independent of position x,y:
  - $g = t(f)$
- This is called point processing.

- **Important:** every pixel for himself – spatial information completely lost!

Negative

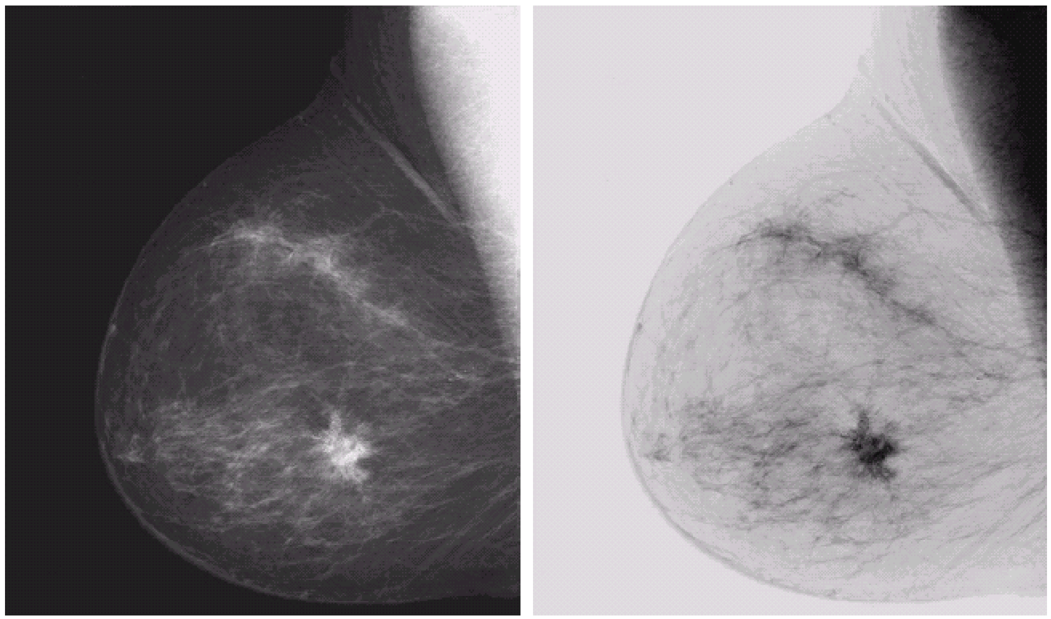

a b

FIGURE 3.4 (a) Original digital mammogram. (b) Negative image obtained using the negative transformation in Eq. (3.2-1). (Courtesy of G.E. Medical Systems.)

Contrast Stretching

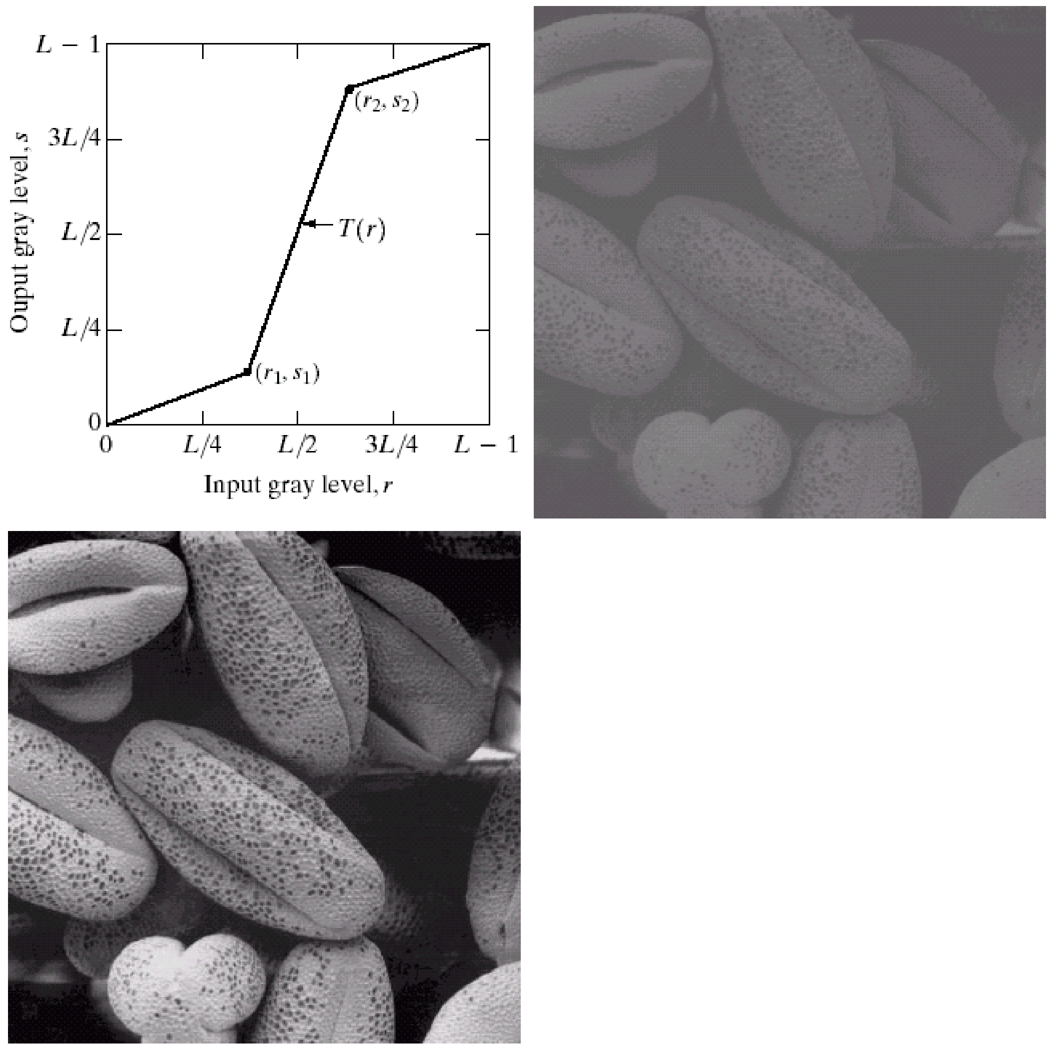

a b
c d

FIGURE 3.10 Contrast stretching. (a) Form of transformation function. (b) A low-contrast image. (c) Result of contrast stretching. (d) Result of thresholding. (Original image courtesy of Dr. Roger Heady, Research School of Biological Sciences, Australian National University, Canberra, Australia.)

Image Histograms

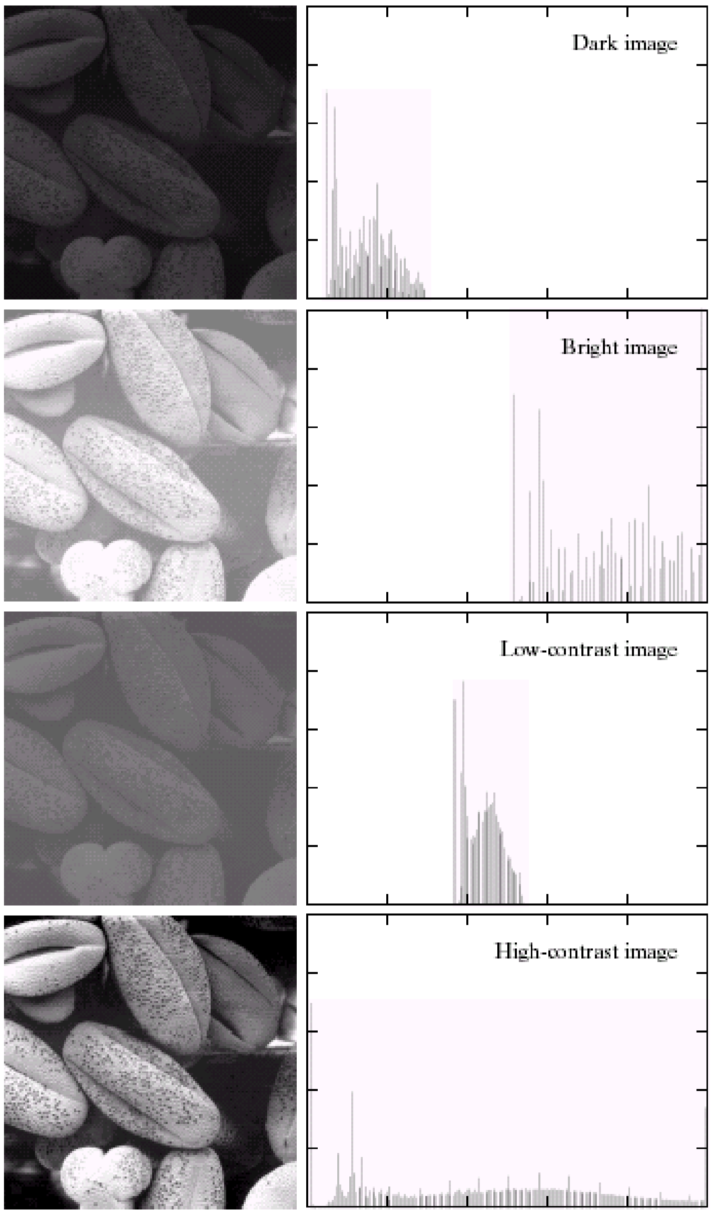

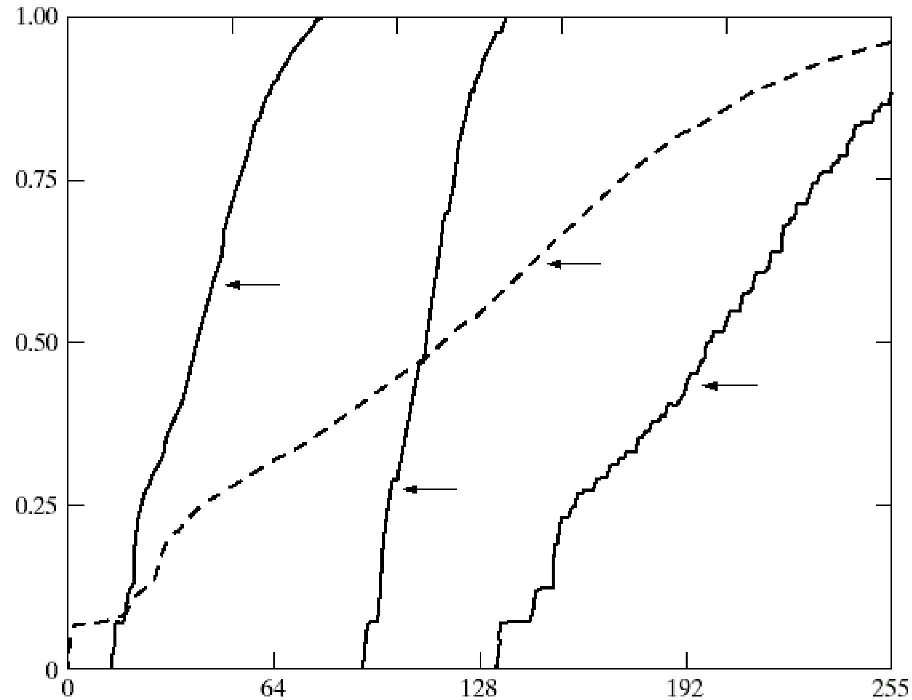

Cumulative Histograms

s = T(r)

a b

FIGURE 3.15 Four basic image types: dark, light, low contrast, high contrast, and their corresponding histograms. (Original image courtesy of Dr. Roger Heady, Research School of Biological Sciences, Australian National University, Canberra, Australia.)

Histogram Equalization

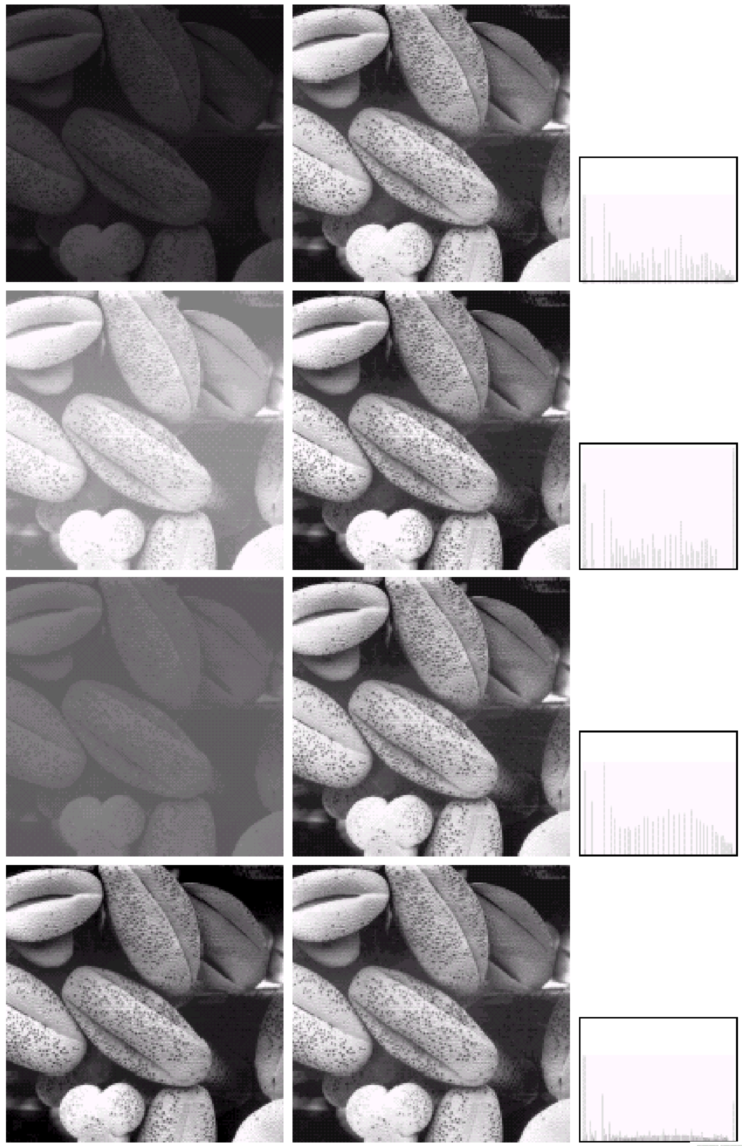

FIGURE 3.17 (a) Images from Fig. 3.15. (b) Results of histogram equalization. (c) Corresponding histograms.

Questions?

Filtering

- So far we have looked at range-only and domain-only transformation
- But other transforms need to change the range according to the spatial neighborhood
  - Linear filtering in particular

Linear filtering

- Replace each pixel by a linear combination of its neighbors.
- The prescription for the linear combination is called the “convolution kernel”.

10	5	3
4	5	1
1	1	7

Local image data

0	0	0
0	0.5	0
0	1	0.5

kernel

	7	

Modified image data
(shown at one pixel)

More formally: Convolution

$$f[m,n] = I \otimes g = \sum_{k,l} I[m-k, n-l]\, g[k,l]$$

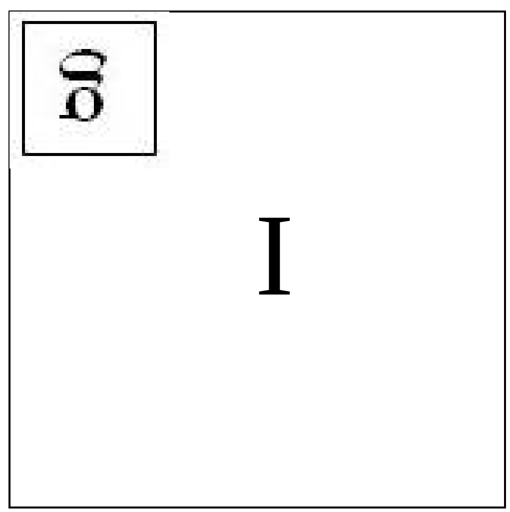

Linear filtering (warm-up slide)

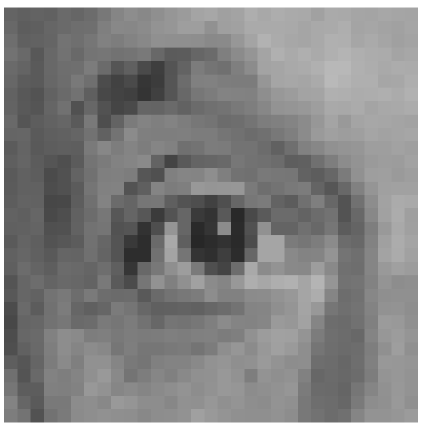

original

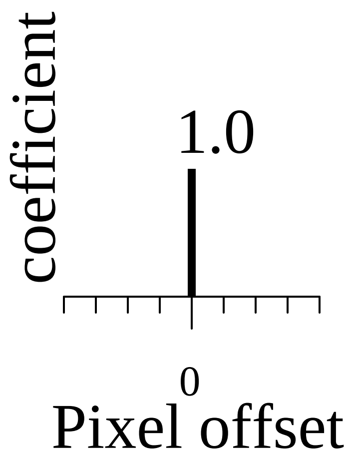

?

Linear filtering (warm-up slide)

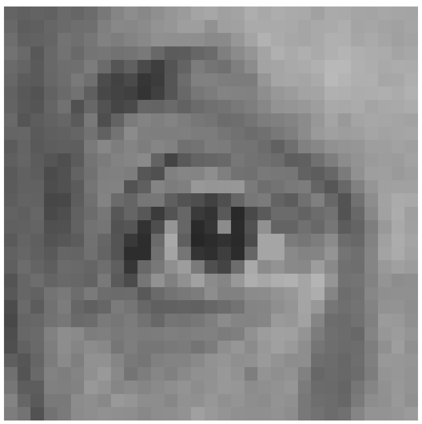

original

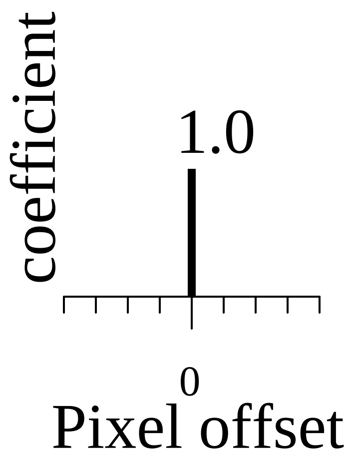

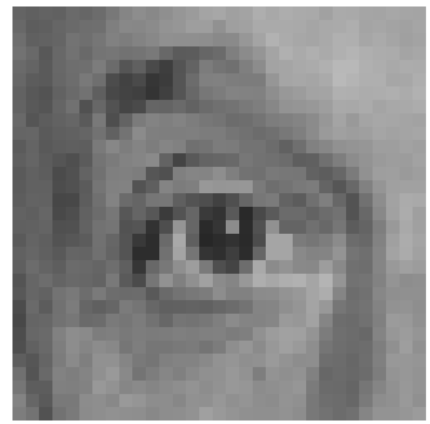

Filtered
(no change)

Linear filtering

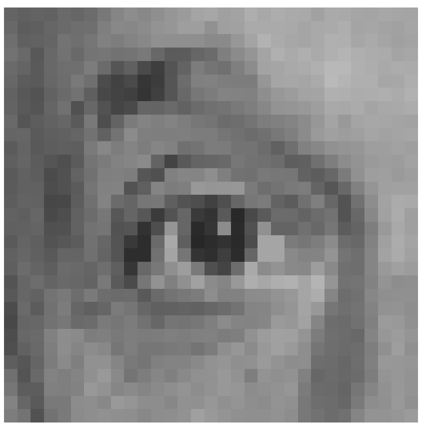

original

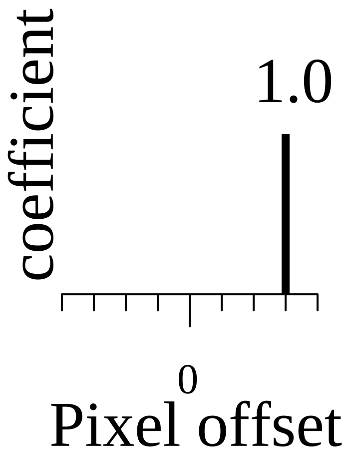

?

shift

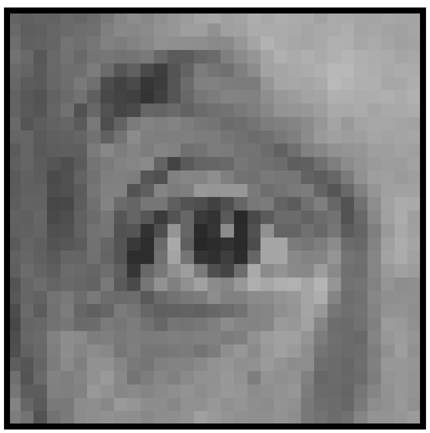

original

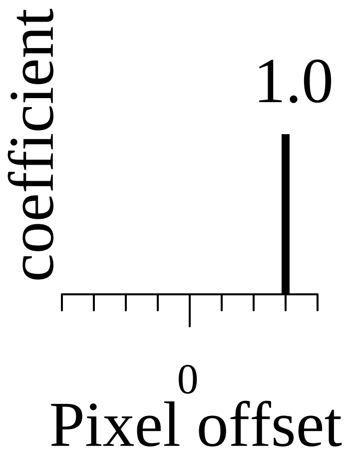

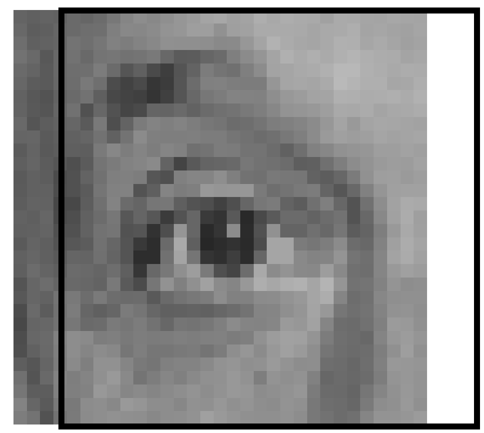

shifted

Linear filtering

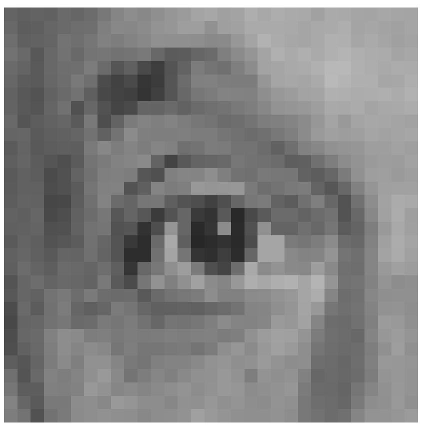

original

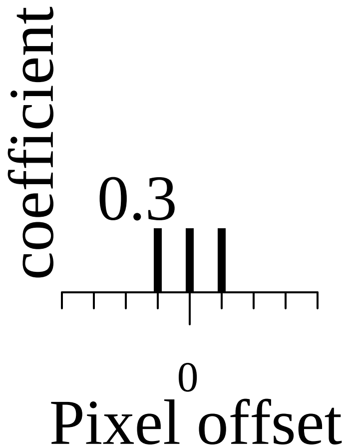

?

Blurring

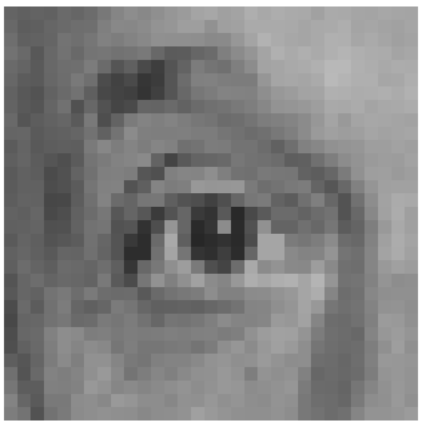

original

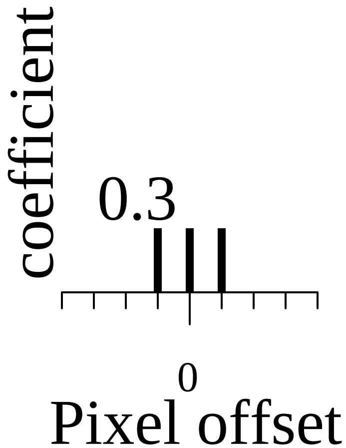

Blurred (filter applied in both dimensions).

Blur examples

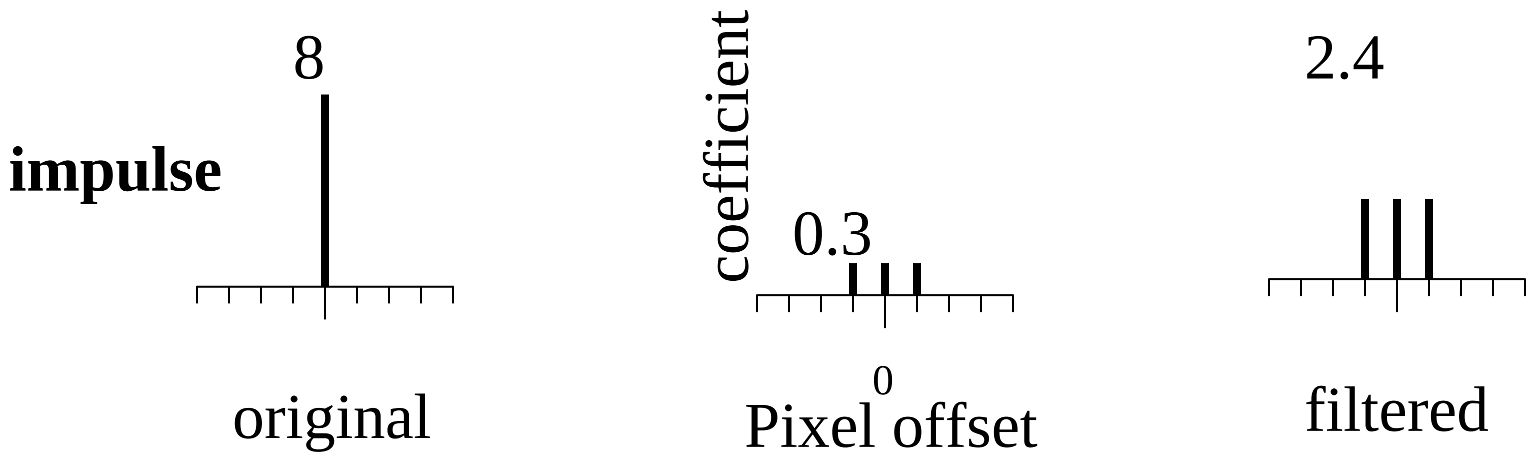

Blur examples

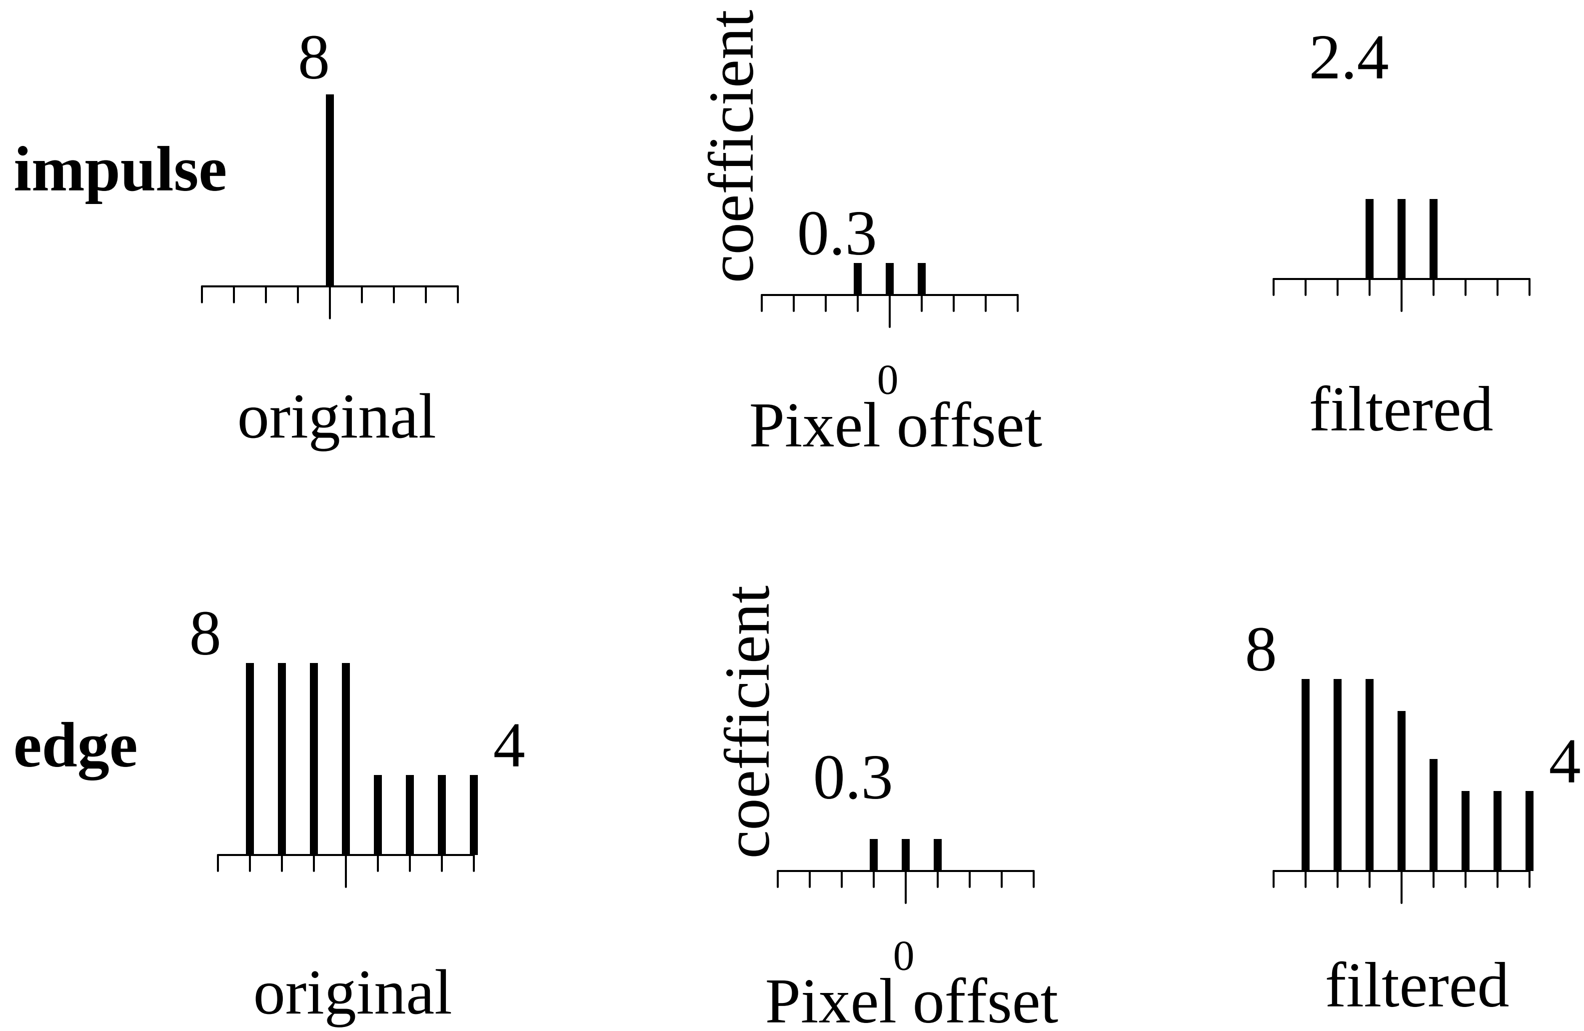

Questions?

Linear filtering (warm-up slide)

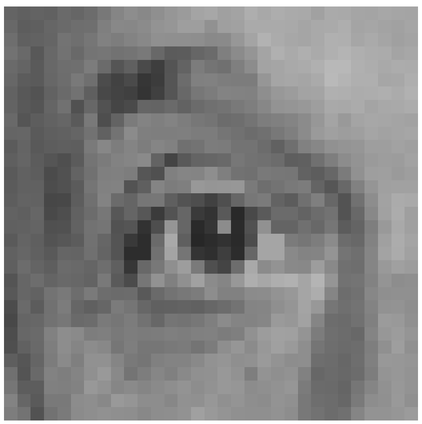

original

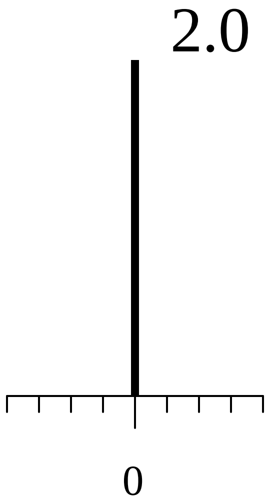

−

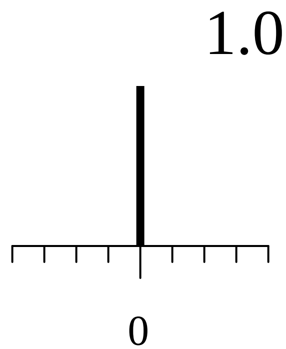

?

Linear filtering (no change)

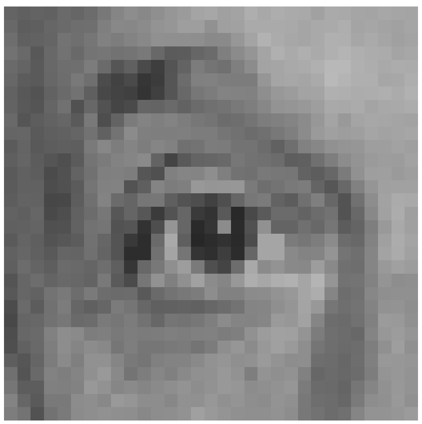

original

2.0

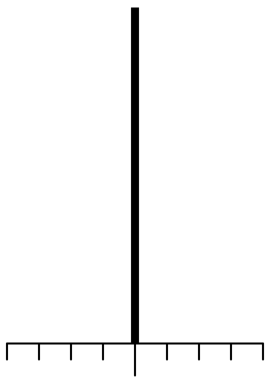

0

−

1.0

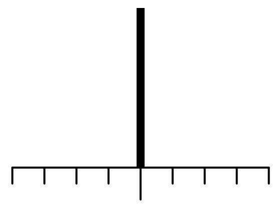

0

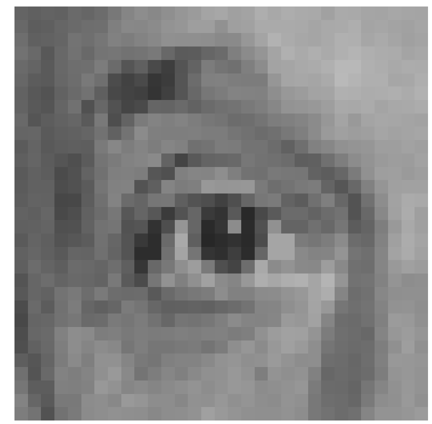

Filtered
(no change)

Linear filtering

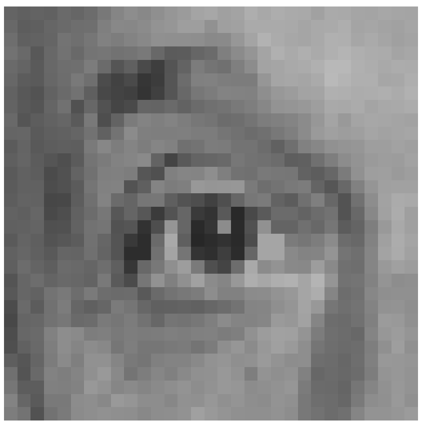

original

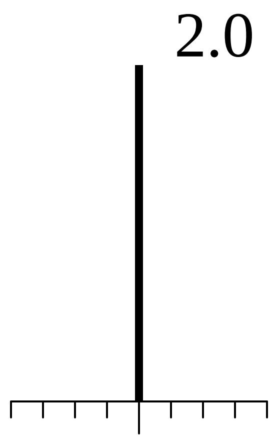

−

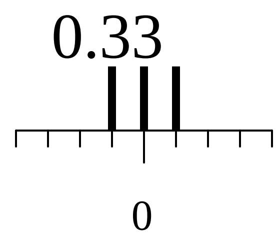

?

(remember blurring)

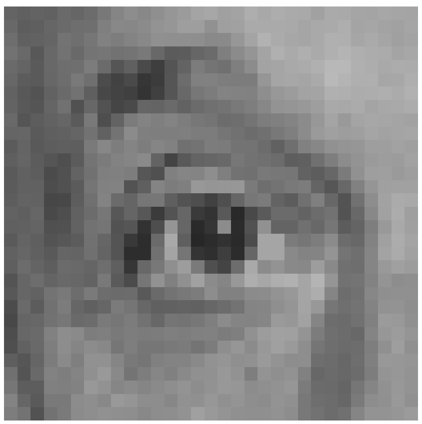

original

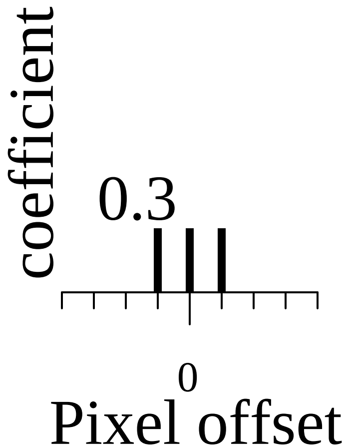

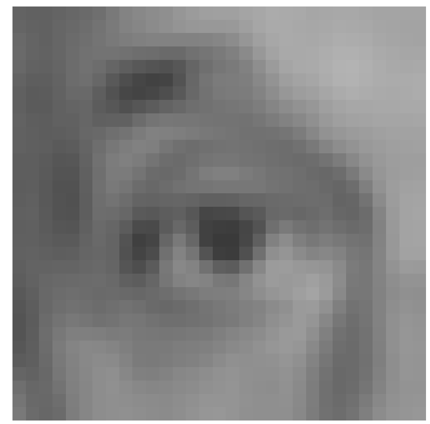

Blurred (filter applied in both dimensions).

Sharpening

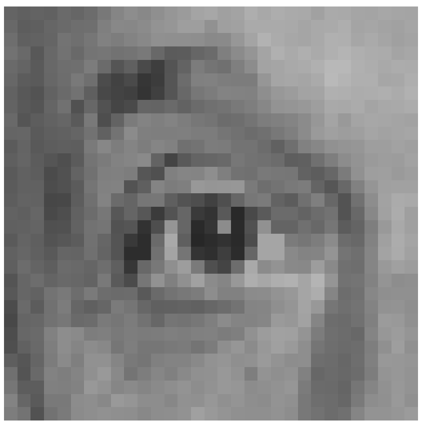

original

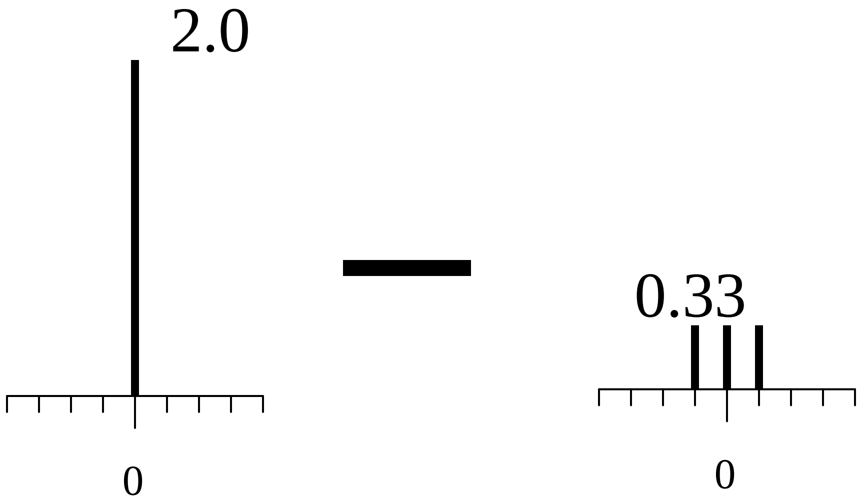

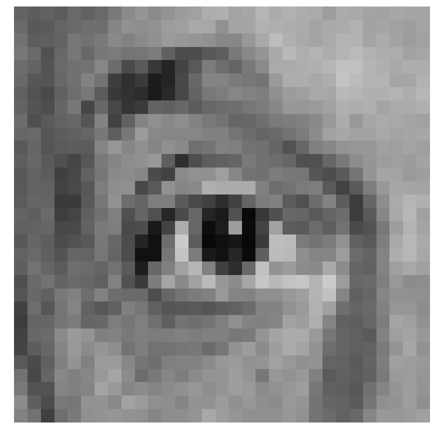

Sharpened
original

Sharpening example

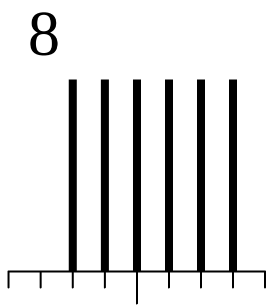

original

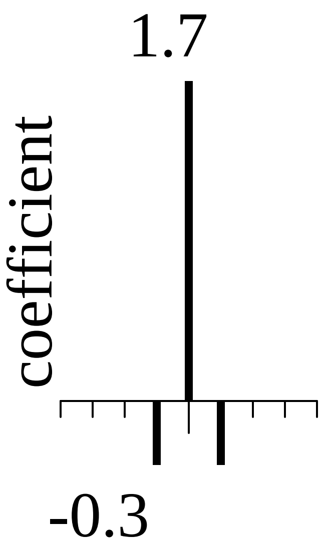

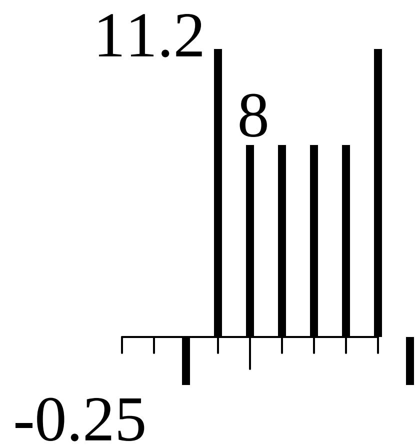

Sharpened
(differences are accentuated; constant areas are left untouched).

Sharpening

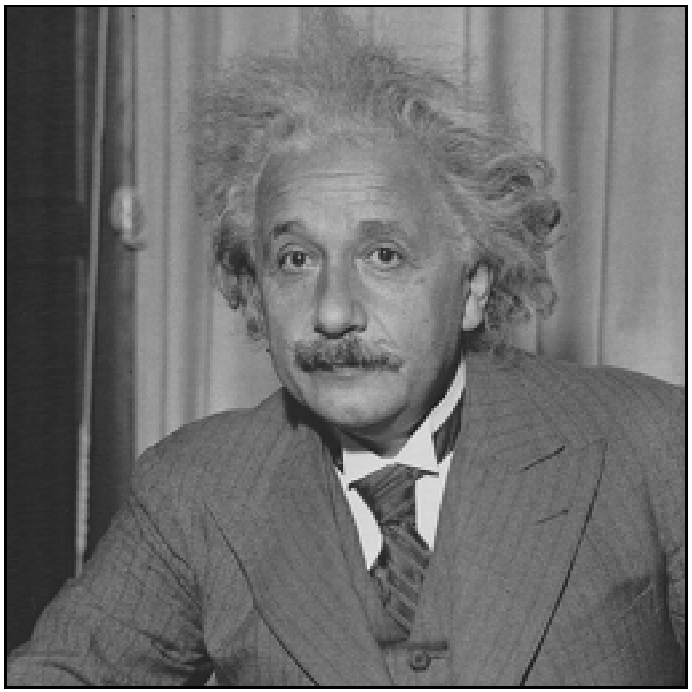

before

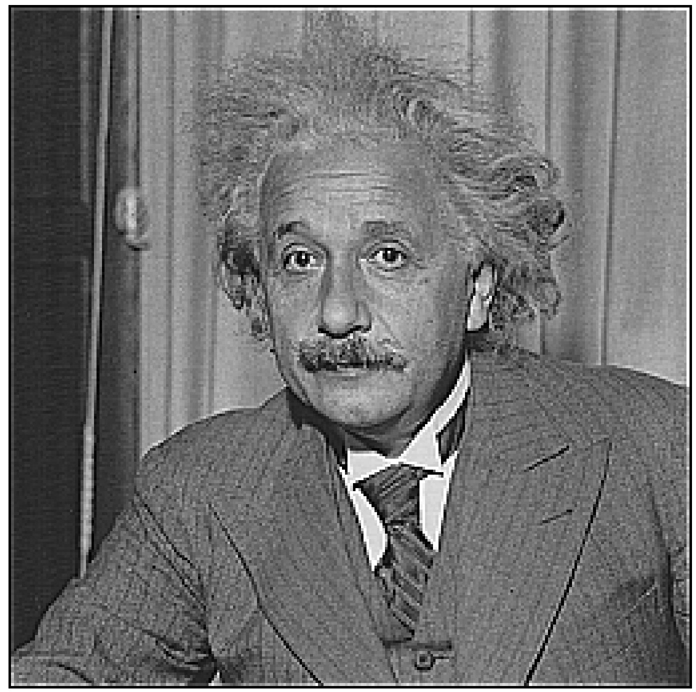

after

Questions?

Studying convolutions

- Convolution is complicated
  - But at least it’s linear
    (f+kg)- h = f-h +kg
- We want to find a better expression
  - Let’s study function whose behavior is simple under convolution

Blurring: convolution

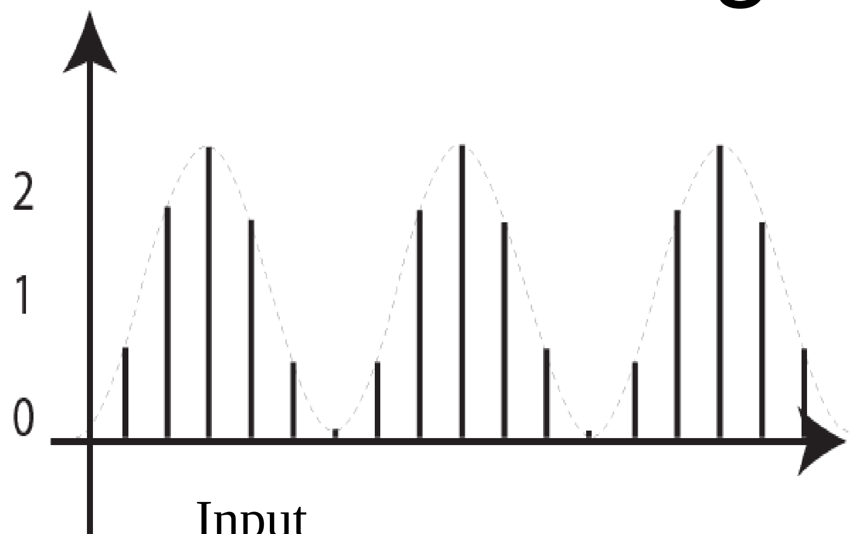

Input

$\otimes$

Convolution sign

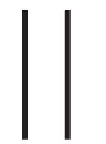

Kernel

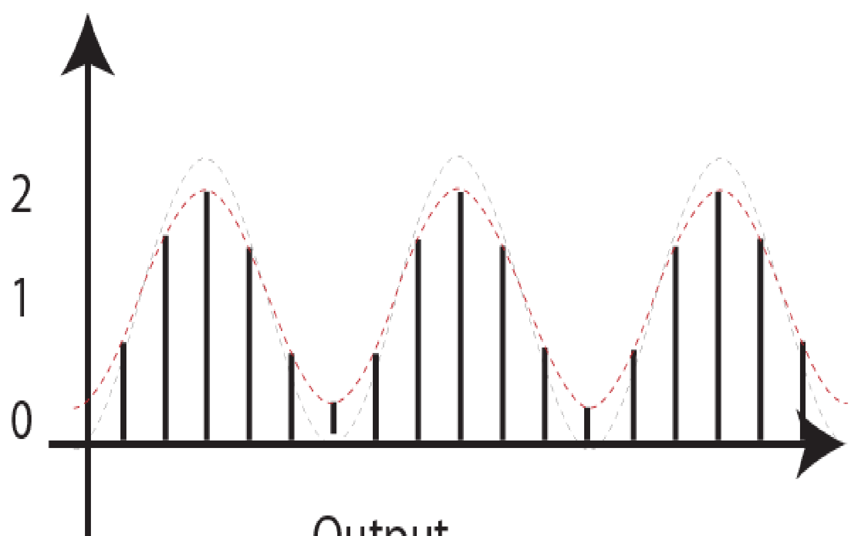

Output

Same shape, just reduced contrast!!!

This is an eigenvector
(output is the input multiplied by a constant)

Big Motivation for Fourier analysis

- Sine waves are eigenvectors of the convolution operator

Other motivation for Fourier analysis: sampling

- The sampling grid is a periodic structure
  - Fourier is pretty good at handling that
  - A sine wave can have serious problems with sampling
- Sampling is a linear process

Sampling Density

- If we’re lucky, sampling density is enough

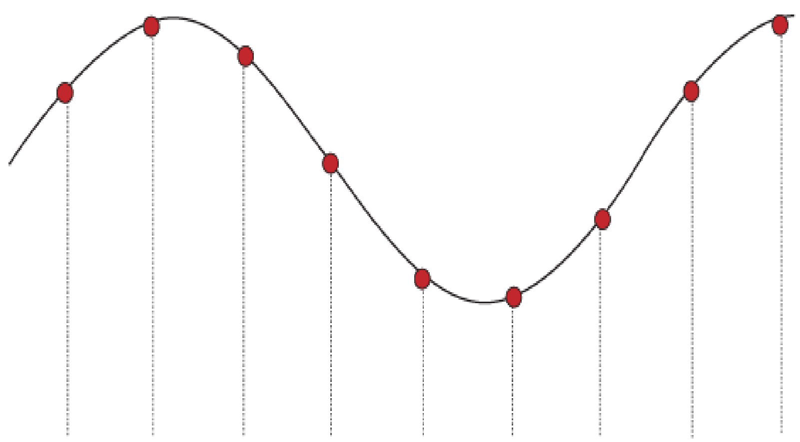

Input

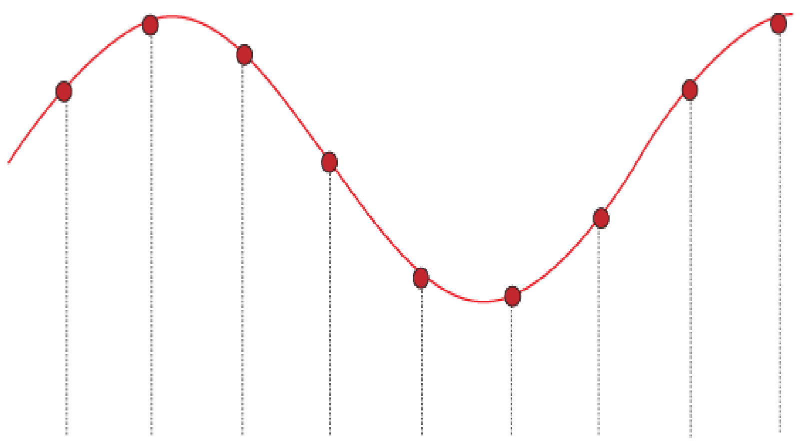

Reconstructed

Sampling Density

- If we insufficiently sample the signal, it may be mistaken for something simpler during reconstruction (that's aliasing!)

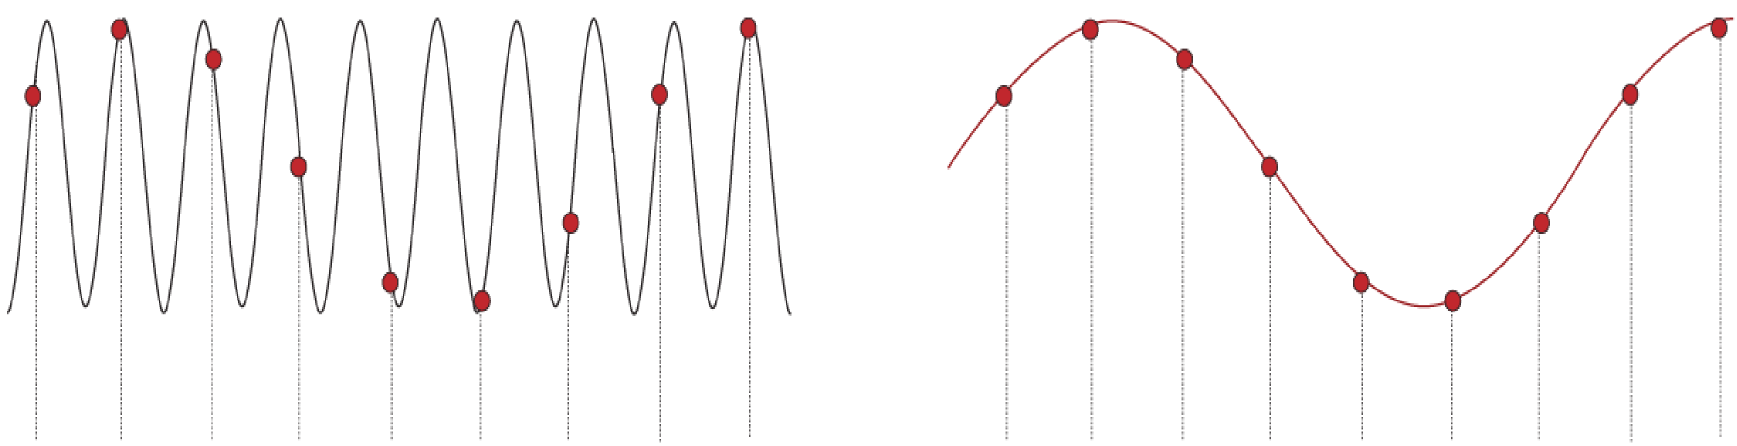

Motivation for sine waves

- Blurring sine waves is simple
  - You get the same sine wave, just scaled down
  - The sine functions are the eigenvectors of the convolution operator
- Sampling sine waves is interesting
  - Get another sine wave
  - Not necessarily the same one! (aliasing)

If we represent functions (or images) with a sum of sine waves, convolution and sampling are easy to study

Questions?

Fourier as a change of basis

- Discrete Fourier Transform: just a big matrix

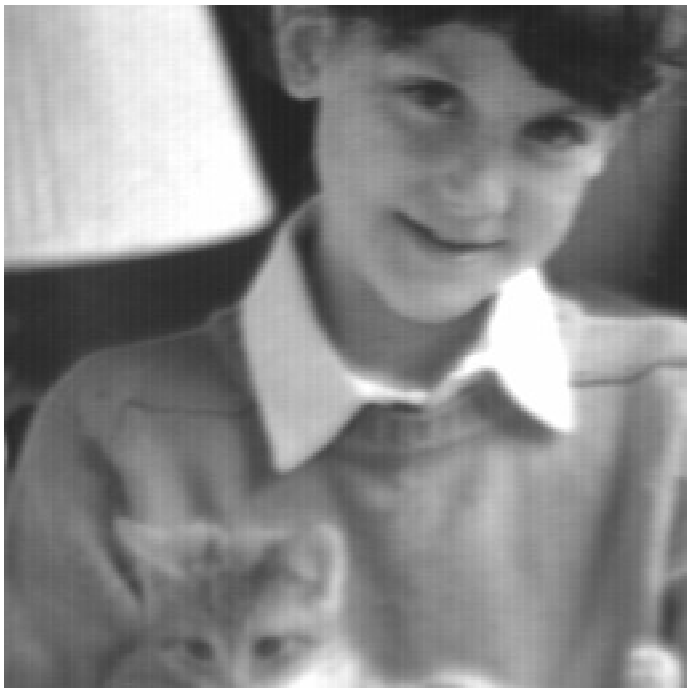

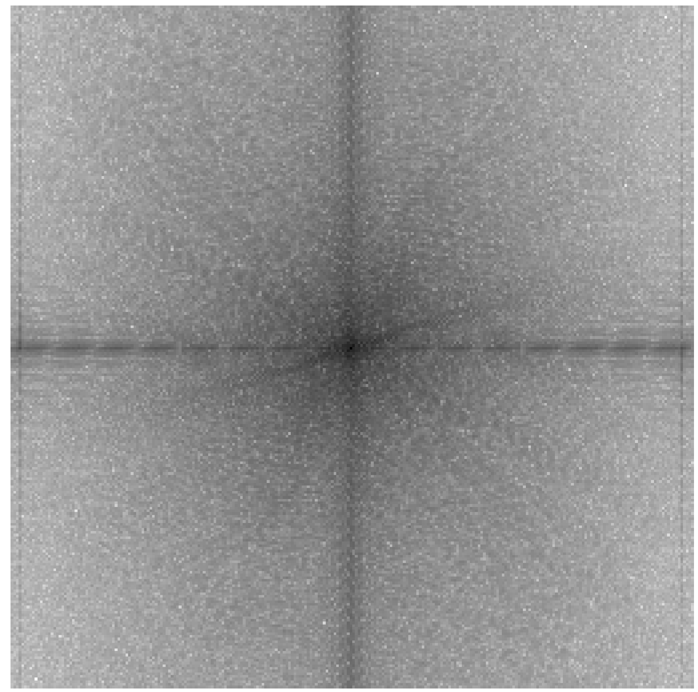

http://www.reindeergraphics.com

To get some sense of what basis elements look like, we plot a basis element --- or rather, its real part --- as a function of x,y for some fixed u, v. We get a function that is constant when (ux+vy) is constant. The magnitude of the vector (u, v) gives a frequency, and its direction gives an orientation. The function is a sinusoid with this frequency along the direction, and constant perpendicular to the direction.

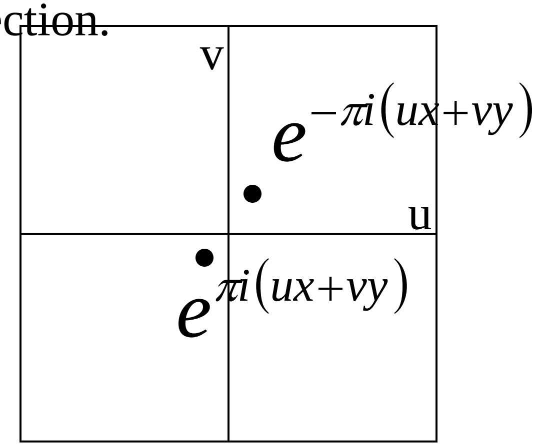

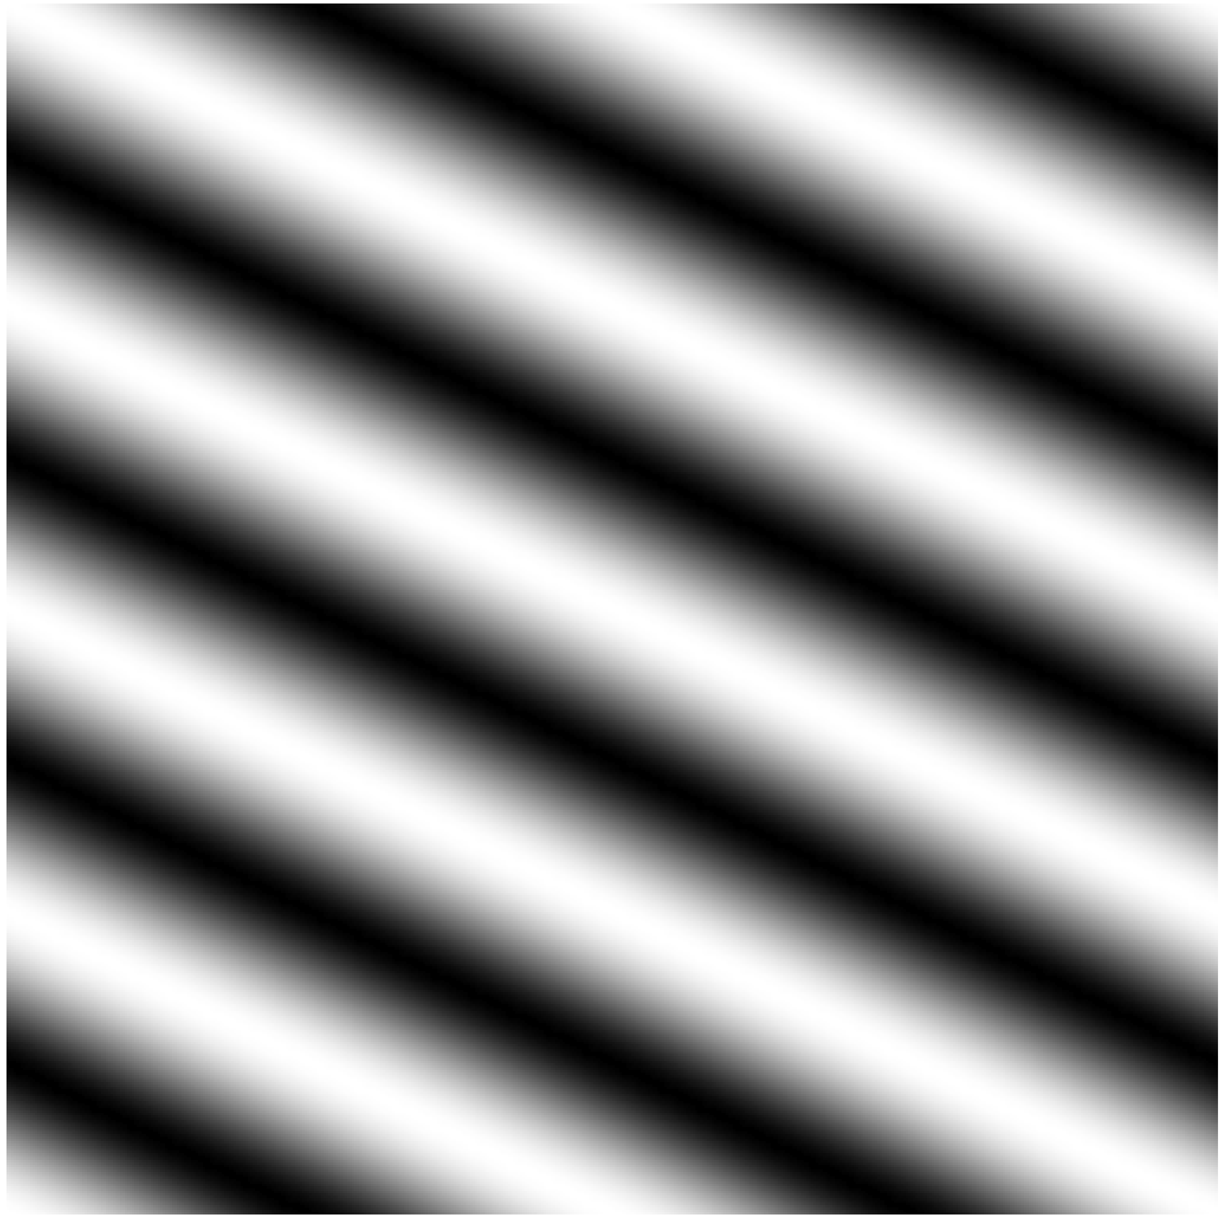

Here u and v are larger than in the previous slide.

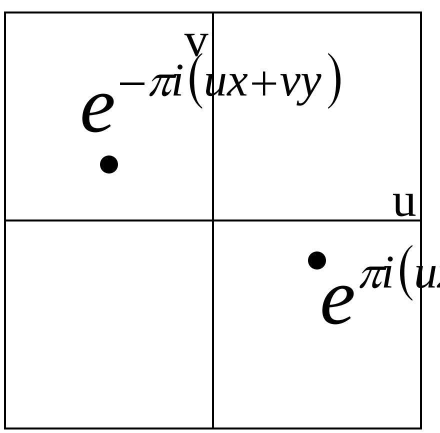

$$e^{-\pi i(ux+vy)}$$

$$e^{\pi i(ux+vy)}$$

And larger still...

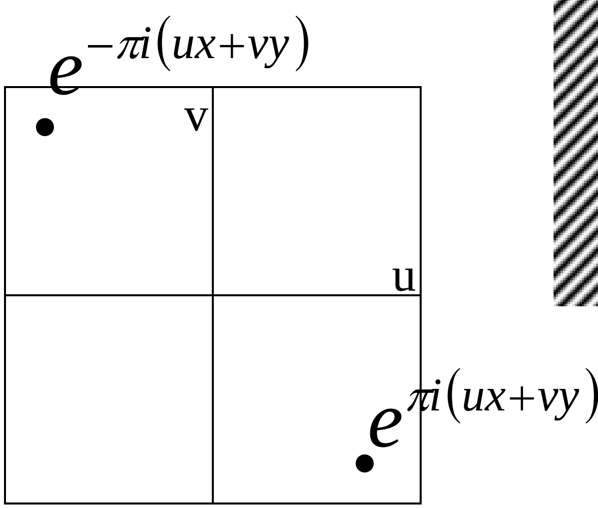

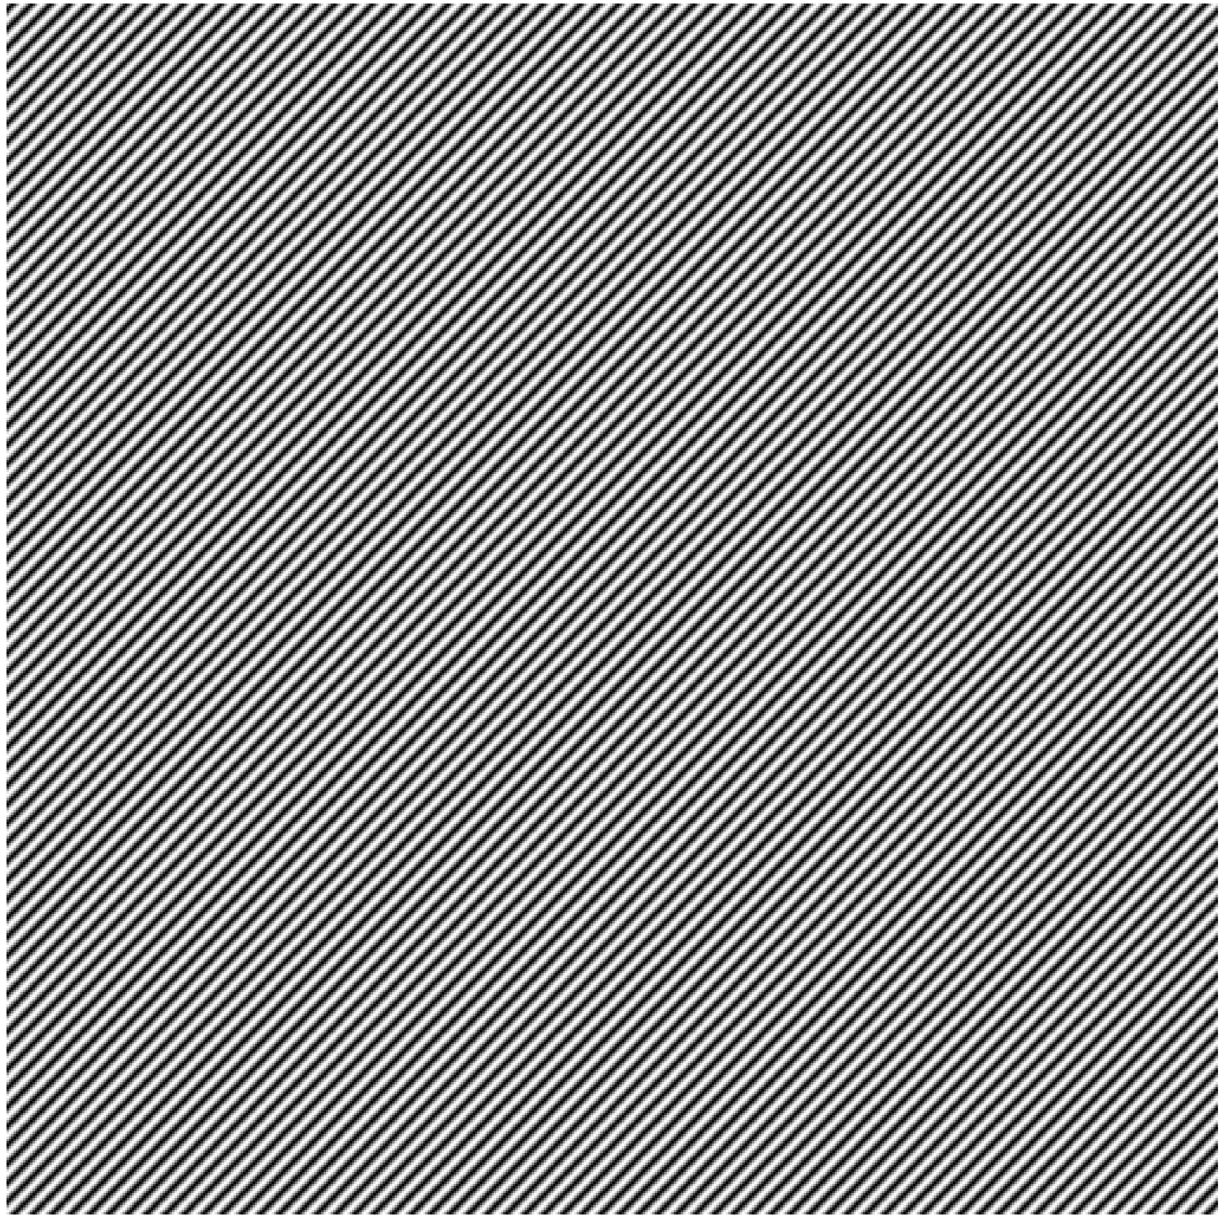

Motivations

- Computation bases
  - E.g. fast filtering
- Sampling rate and filtering bandwidth
- Optics: wave nature of light & diffraction
- Insights

Questions?

Fourier Series

- Consider the family of complex exponentials

$$\psi_n(t) = e^{jn\omega t} \quad \text{where} \quad n \in Z$$

- Properties
  - Periodic with period $T = 2\pi / \omega$
  - Orthogonal on any interval $[T] = [t_0, t_0 + T]$
- Hence, we can write a *periodic* signal *x*(*t*) with period *T* as

$$x(t) = \sum_k a_k \psi_k(t) \quad \text{where} \quad a_k = \frac{\langle \psi_k | x \rangle_{[T]}}{\langle \psi_k | \psi_k \rangle_{[T]}}$$

The Fourier Transform

- Defined for infinite, aperiodic signals
- Derived from the Fourier series by "extending the period of the signal to infinity"
- The Fourier transform is defined as

$$X(\omega) = \frac{1}{\sqrt{2\pi}} \int x(t) e^{-j\omega t} dt$$

- $X(\omega)$ is called the spectrum of $x(t)$
- It contains the magnitude and phase of each complex exponential of frequency ω in $x(t)$

The Fourier Transform

- The inverse Fourier transform is defined as

$$x(t) = \frac{1}{\sqrt{2\pi}} \int X(\omega) e^{j\omega t} d\omega$$

- Fourier transform pair

$$x(t) \xleftrightarrow{F} X(\omega)$$

- $x(t)$ is called the *spatial domain* representation
- $X(\omega)$ is called the *frequency domain* representation

Duality

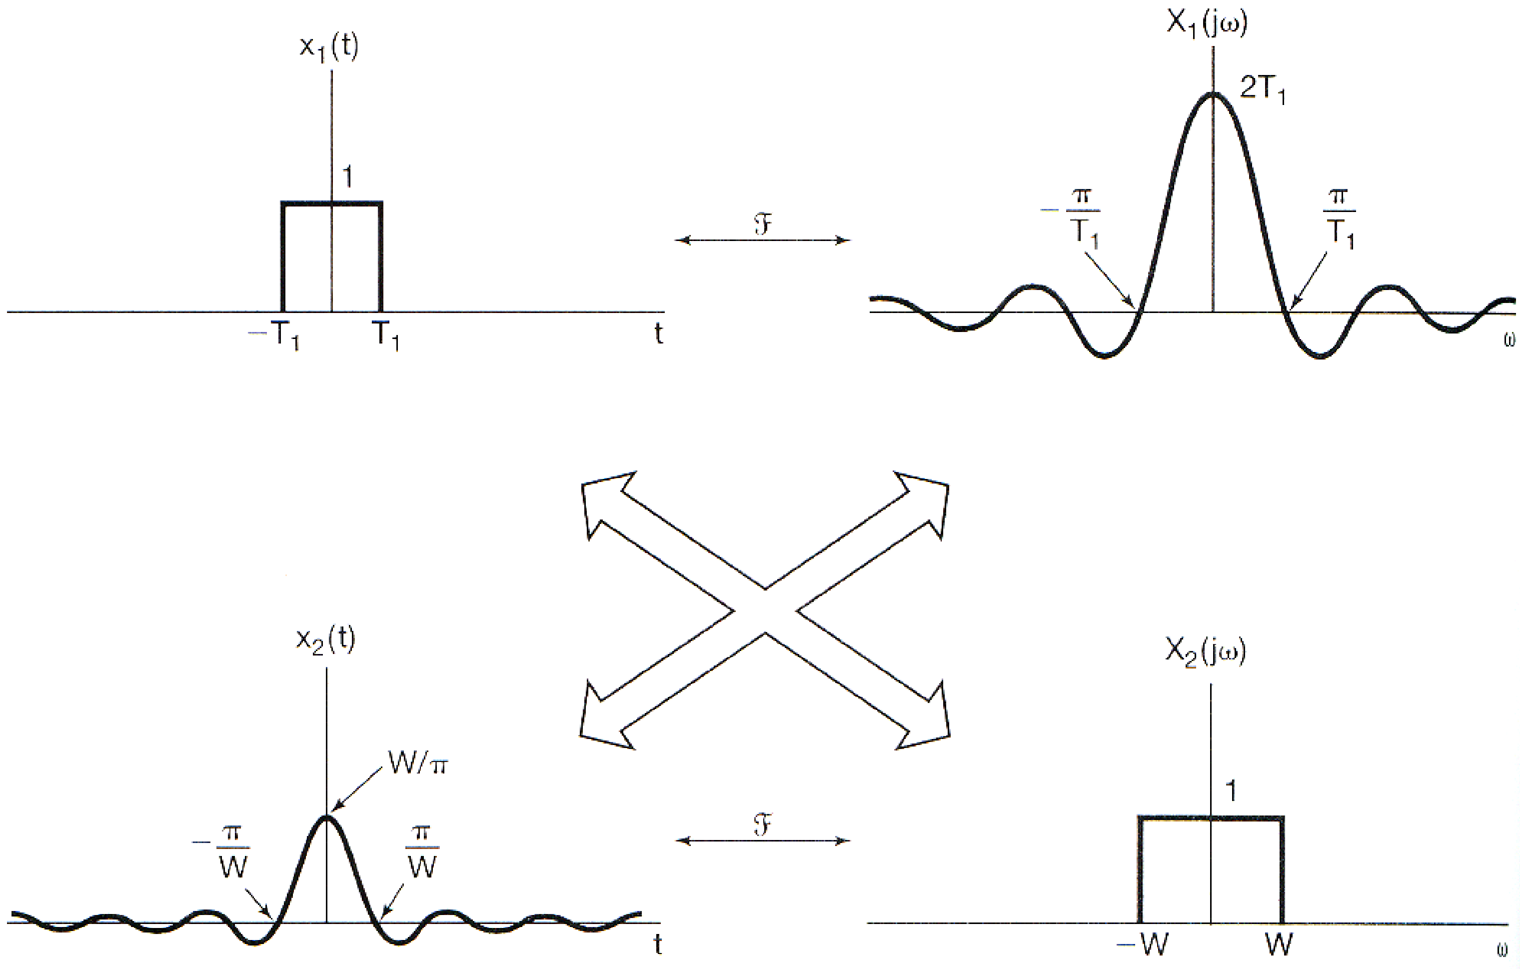

$x_1(t)$
1
$-T_1$
T_1
t
$\mathfrak{F}$
$X_1(j\omega)$
$2T_1$
$-\frac{\pi}{T_1}$
$\frac{\pi}{T_1}$
ω
$x_2(t)$
W/π
$-\frac{\pi}{W}$
$\frac{\pi}{W}$
t
$\mathfrak{F}$
$X_2(j\omega)$
1
$-W$
W
ω

Beware of differences

- Different definitions of Fourier transform
- We use

$$X(\omega) = \frac{1}{\sqrt{2\pi}} \int x(t) e^{-j\omega t} dt$$

- Other people might exclude normalization or include 2π in the frequency
- X might take ω or $j\omega$ as argument
- Physicist use *j*, mathematicians use *i*

Questions?

Phase

- Don’t forget the phase! Fourier transform results in complex numbers

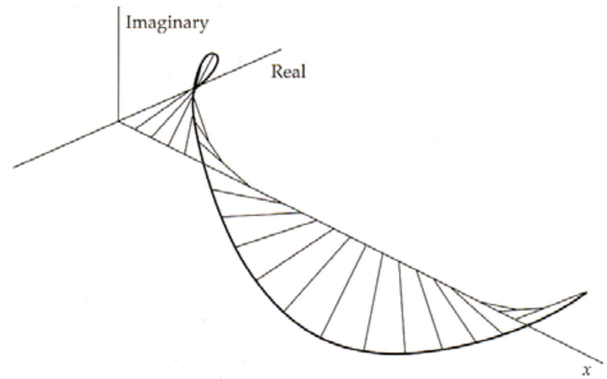

- Can be seen as sum of sines and cosines

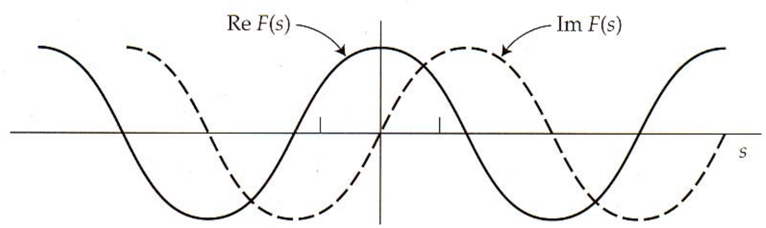

- Or modulus/phase

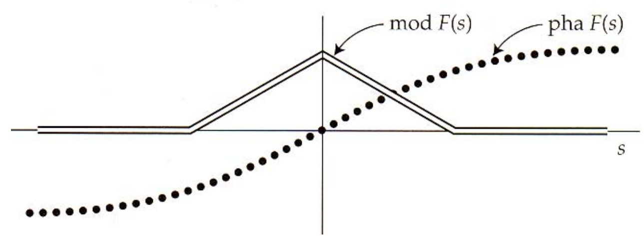

Phase is important!

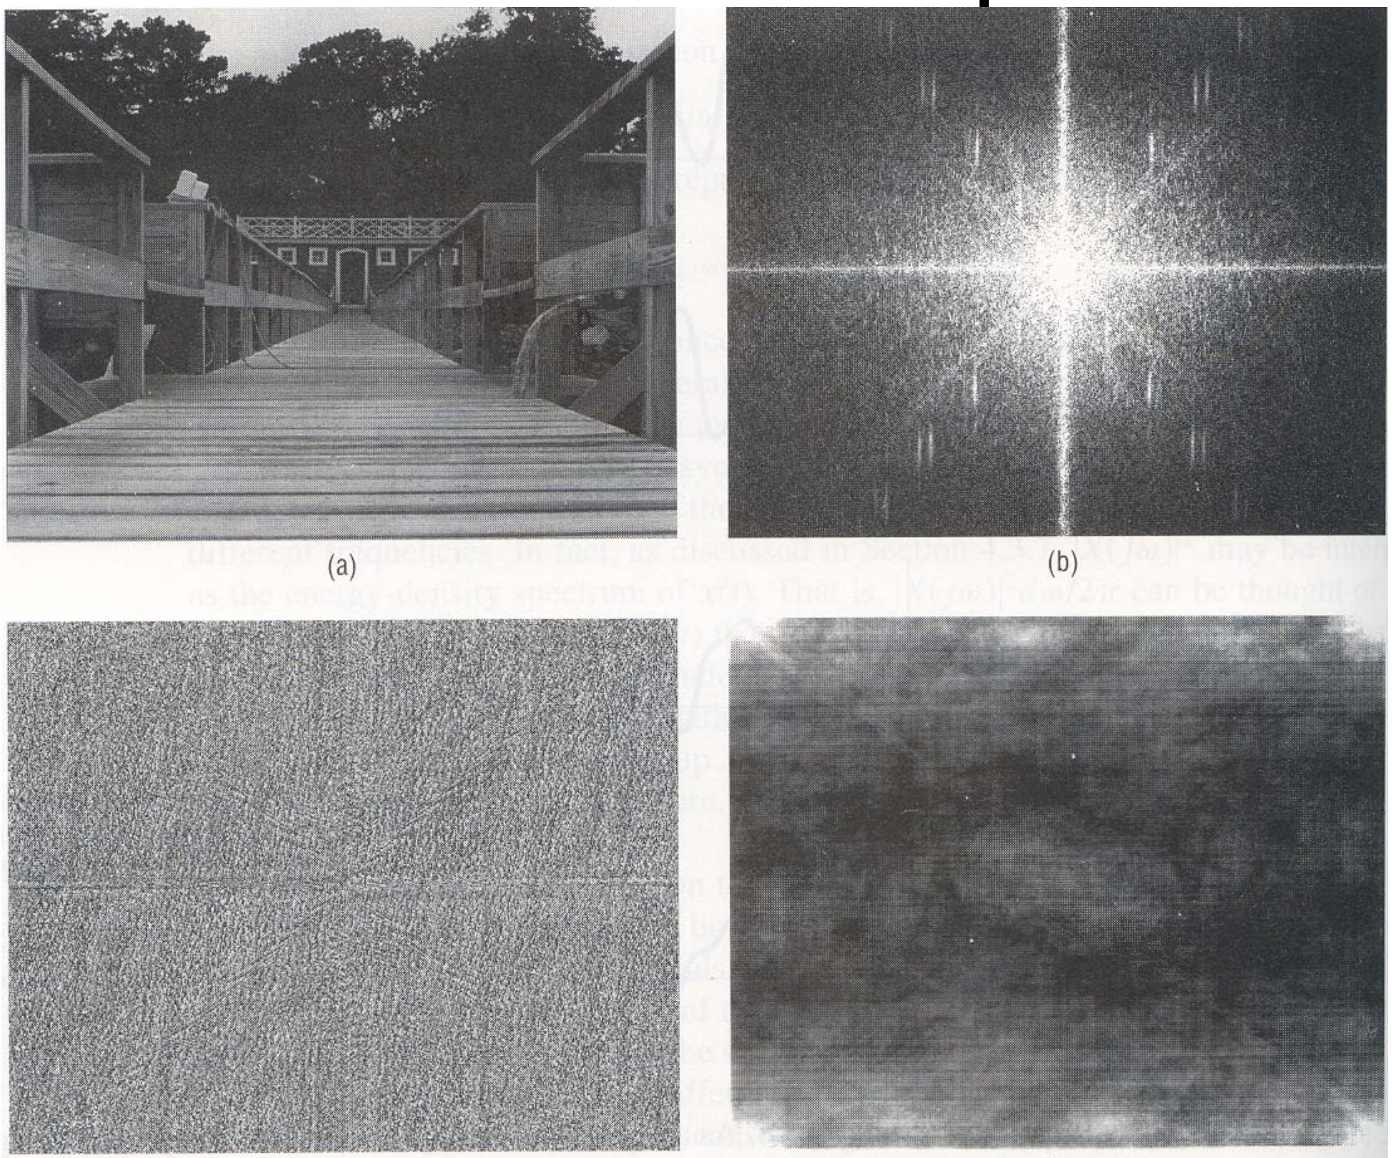

(a) (b)

(c) (d)

Phase is important!

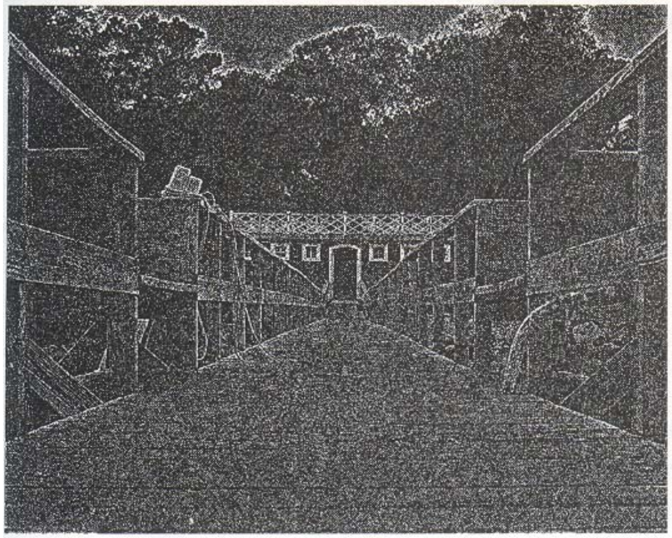

(e)

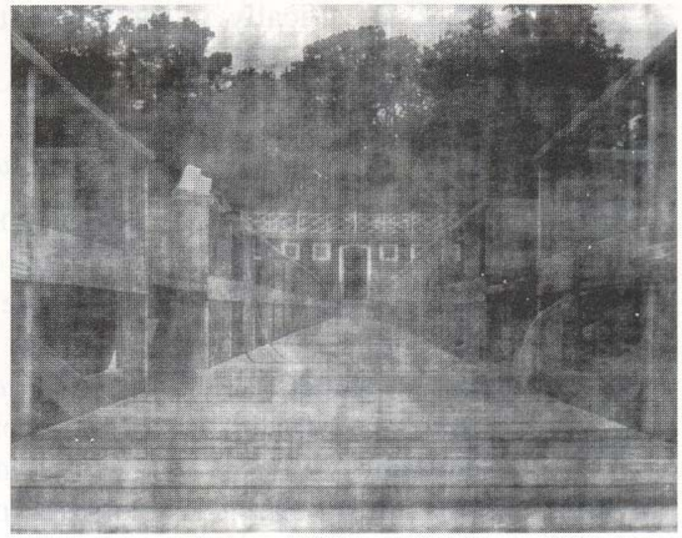

(f)

(g)

Figure 6.2 (a) The image shown in Figure 1.4; (b) magnitude of the two-dimensional Fourier transform of (a); (c) phase of the Fourier transform of (a); (d) picture whose Fourier transform has magnitude as in (b) and phase equal to zero; (e) picture whose Fourier transform has magnitude equal to 1 and phase as in (c); (f) picture whose Fourier transform has phase as in (c) and magnitude equal to that of the transform of the picture shown in (g).

Questions?

Convolution

- Sliding window

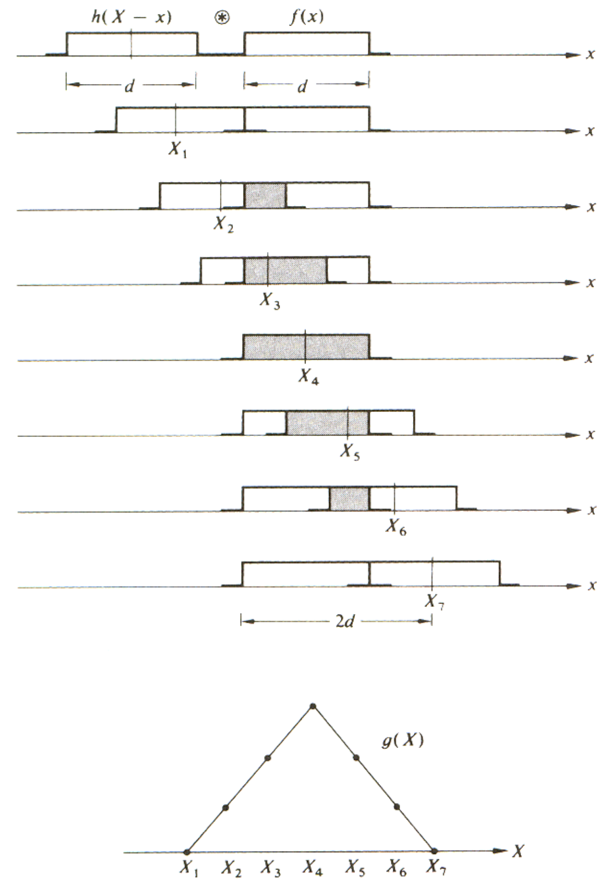

Convolution

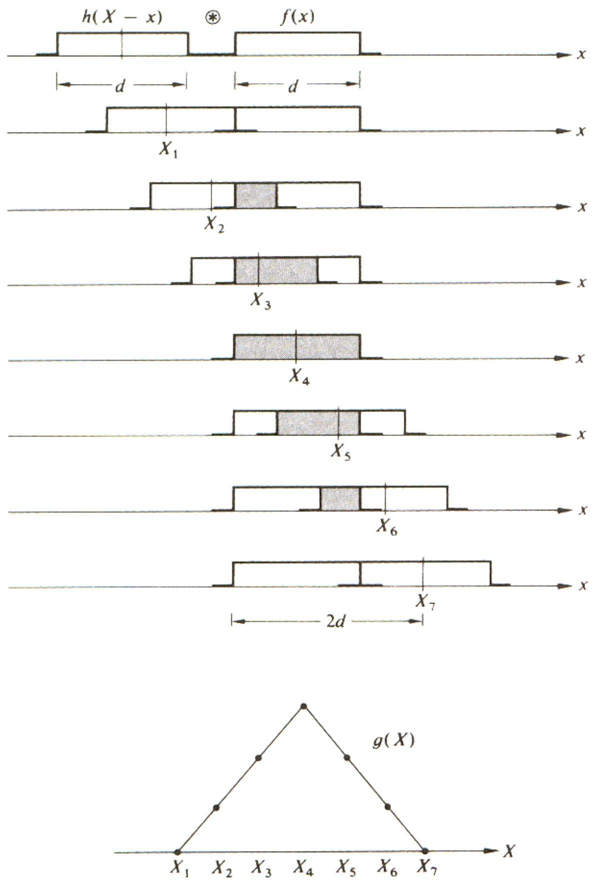

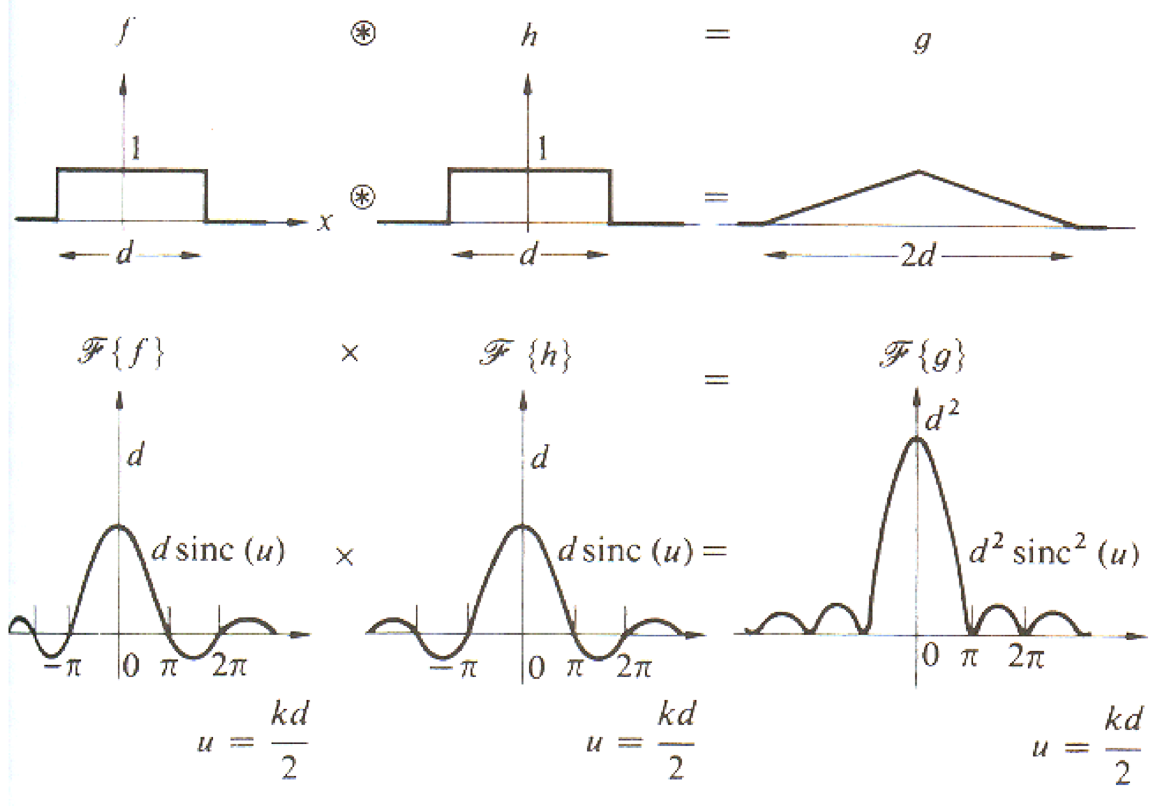

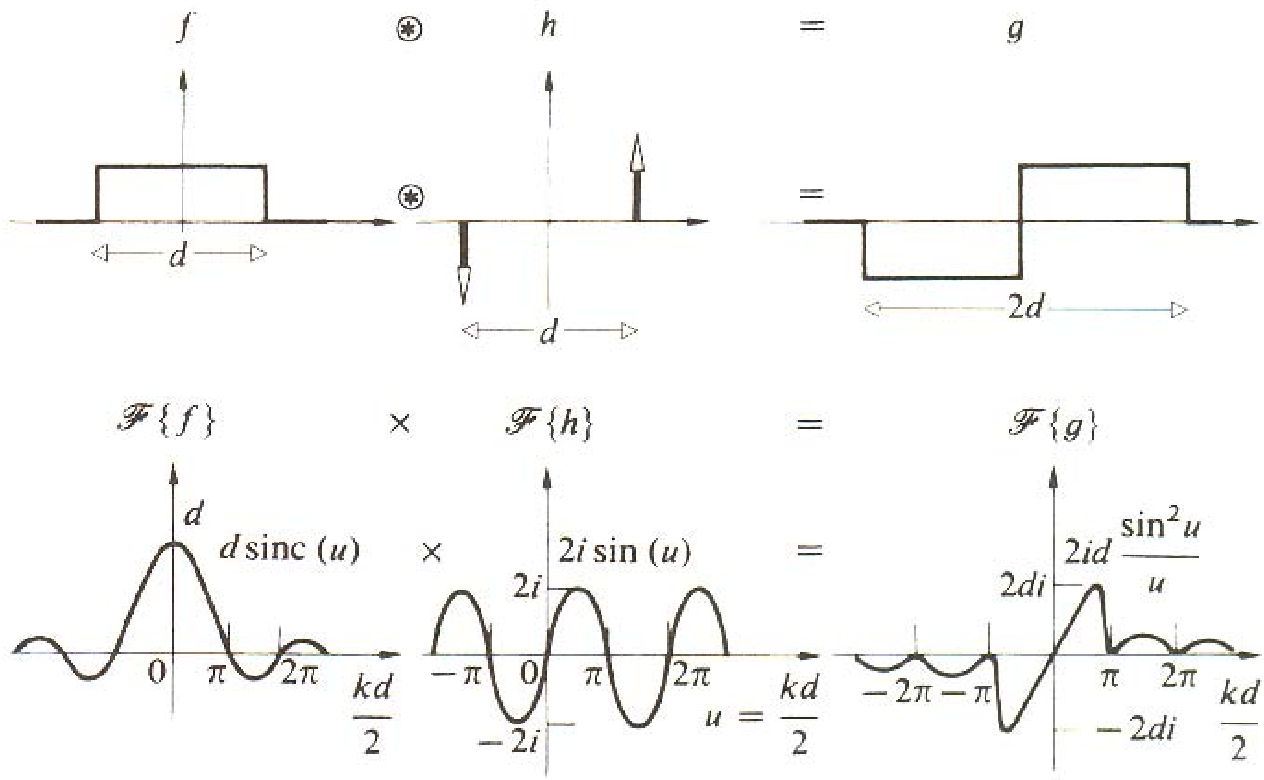

f
⊛
h
=
g
⊛
=
d
d
2d
ℱ{f}
×
ℱ{h}
=
ℱ{g}
d
d sinc (u)
×
2i sin (u)
=
2id sin²u / u
2i
2di
0
π
2π
kd/2
−π
0
π
2π
u = kd/2
−2i
−2π −π
π
2π
−2di
kd/2

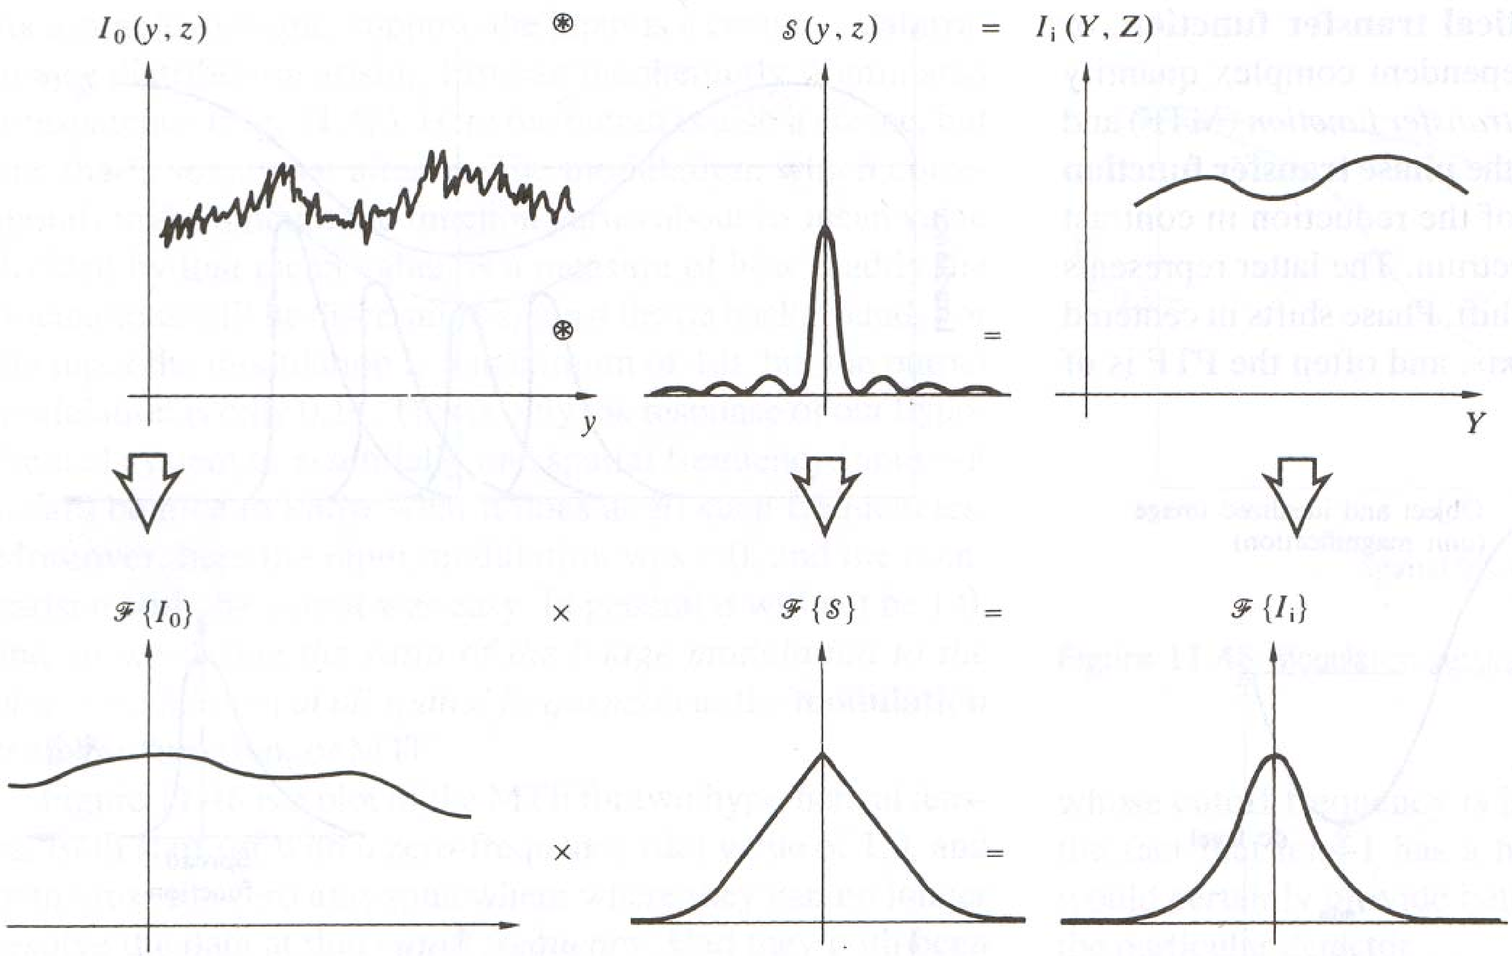

$I_0(y, z)$
⊛
$\mathcal{S}(y, z)$
=
$I_i(Y, Z)$
⊛
=
y
Y
$\mathcal{F}\{I_0\}$
×
$\mathcal{F}\{\mathcal{S}\}$
=
$\mathcal{F}\{I_i\}$
×
=

Low pass

http://www.reindeergraphics.com

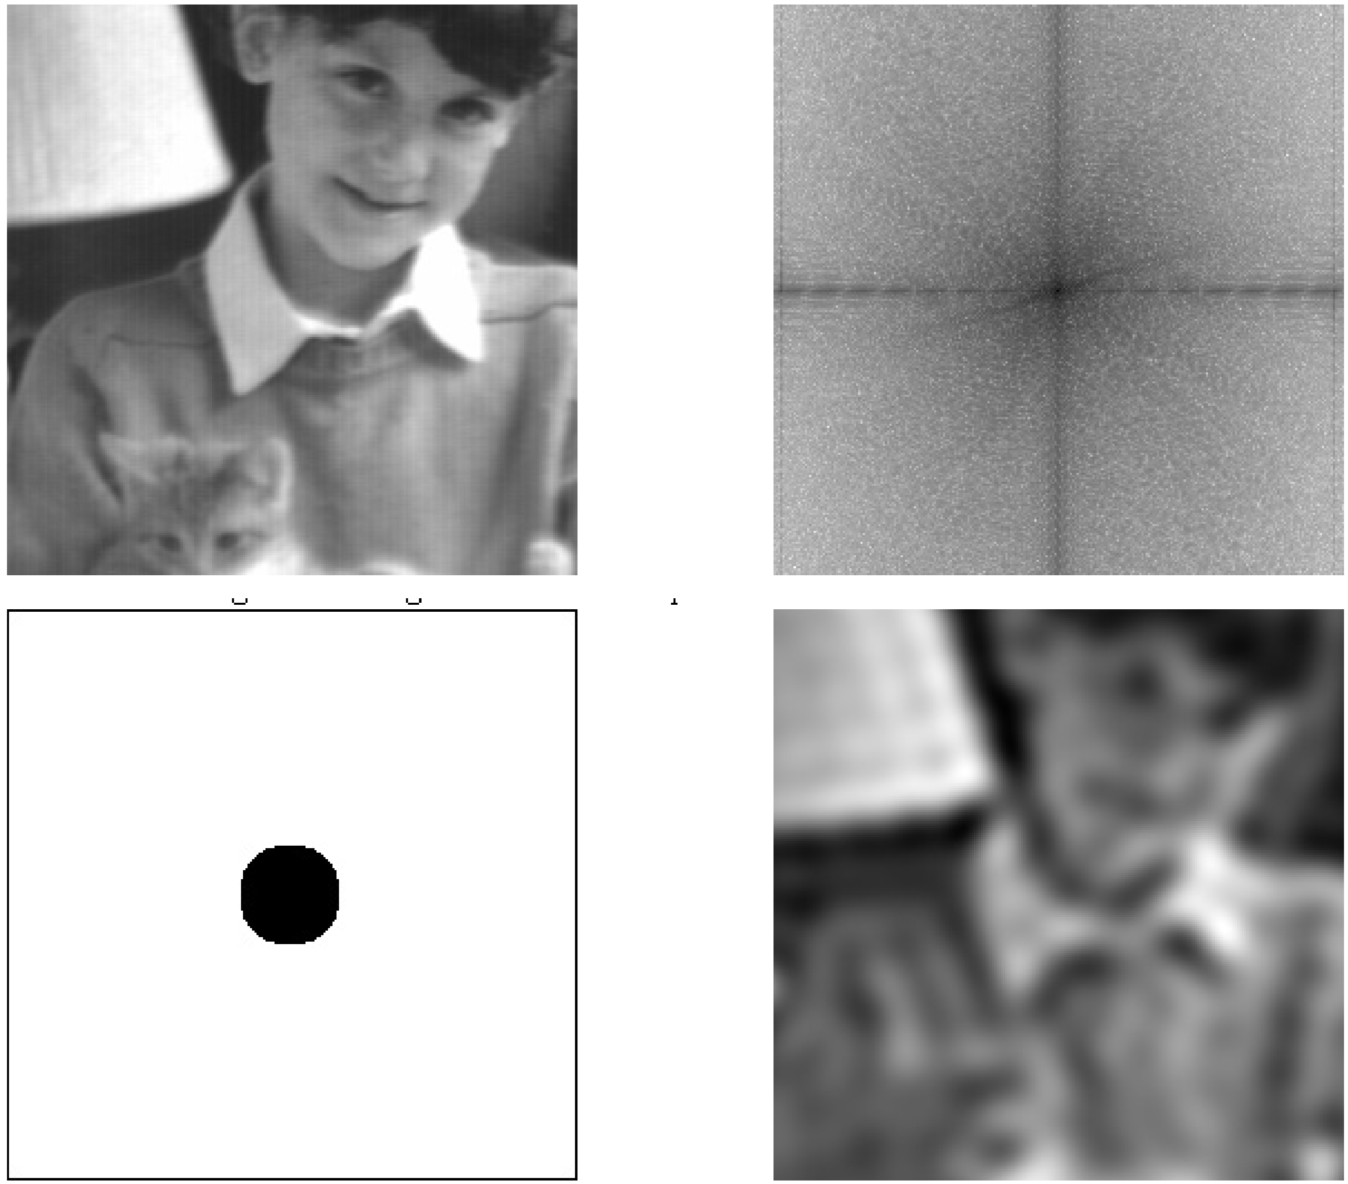

High pass

http://www.reindeergraphics.com

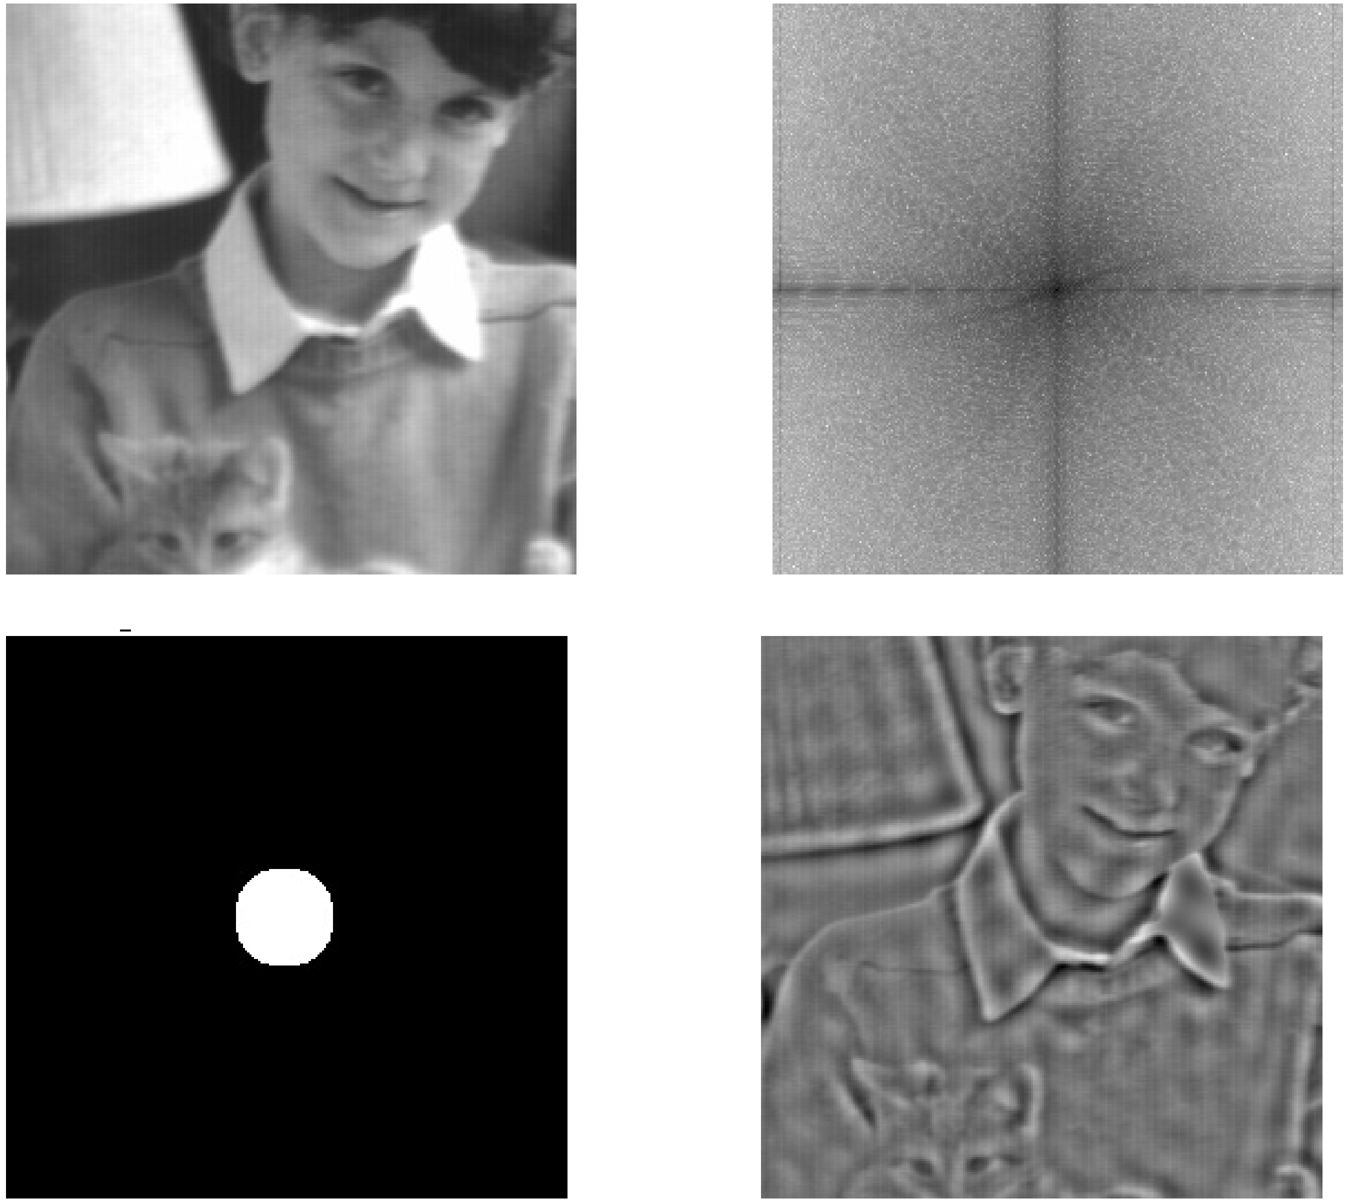

Analysis of our simple filters

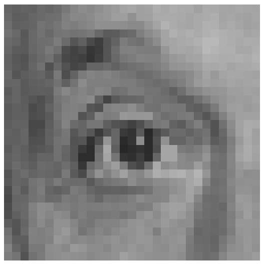

original

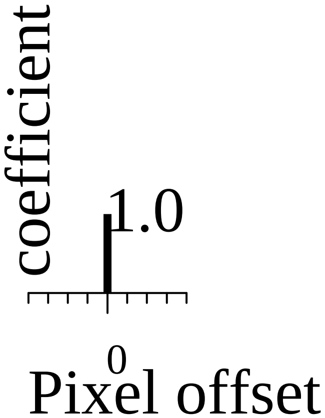

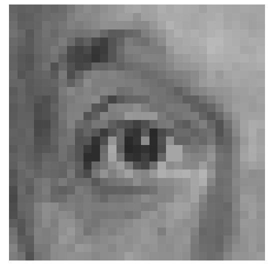

Filtered
(no change)

$$F[m] = \sum_{k=0}^{M-1} f[k] e^{-\pi i \left(\frac{km}{M}\right)}$$

$$= 1$$

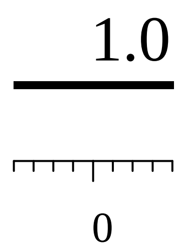

constant

Analysis of our simple filters

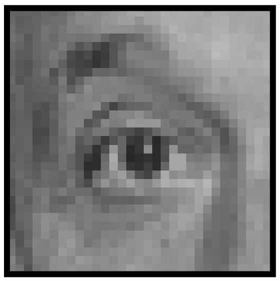

original

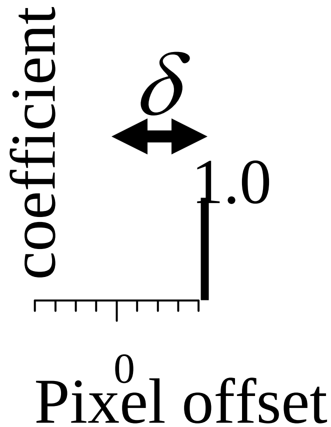

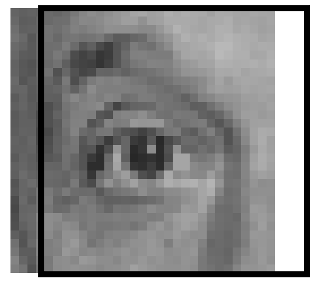

shifted

$$F[m] = \sum_{k=0}^{M-1} f[k] e^{-\pi i \left(\frac{km}{M}\right)}$$

$$= e^{-\pi i \frac{\delta m}{M}}$$

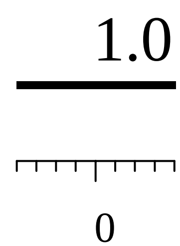

Constant magnitude, linearly shifted phase

Analysis of our simple filters

original

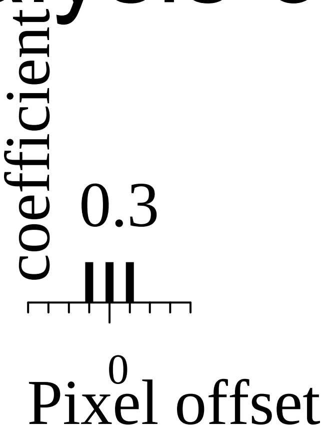

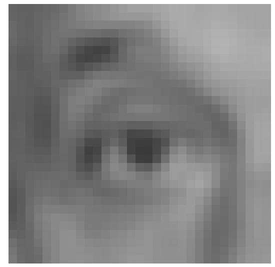

blurred

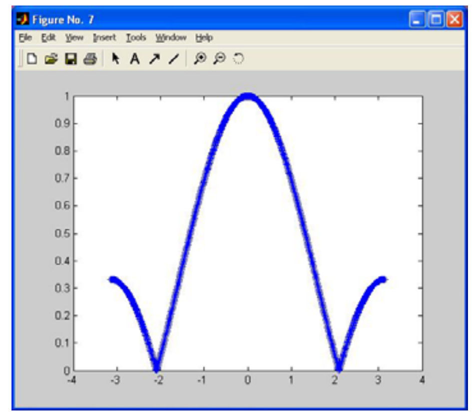

Low-pass filter

$$F[m]=\sum_{k=0}^{M-1} f[k]e^{-\pi i\left(\frac{km}{M}\right)}$$

$$=\frac{1}{3}\left(1+2\cos\left(\frac{\pi m}{M}\right)\right)$$

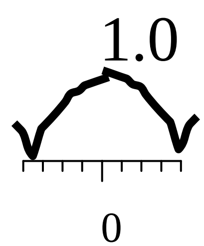

Analysis of our simple filters

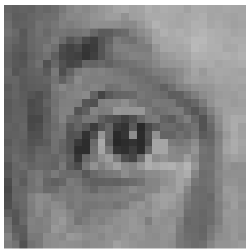

original

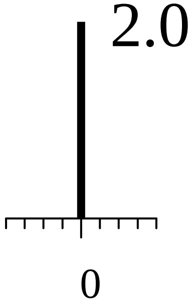

−

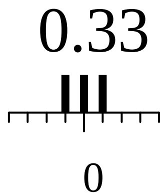

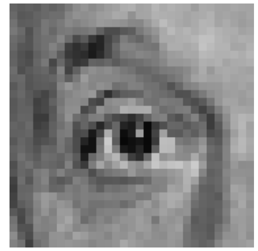

sharpened

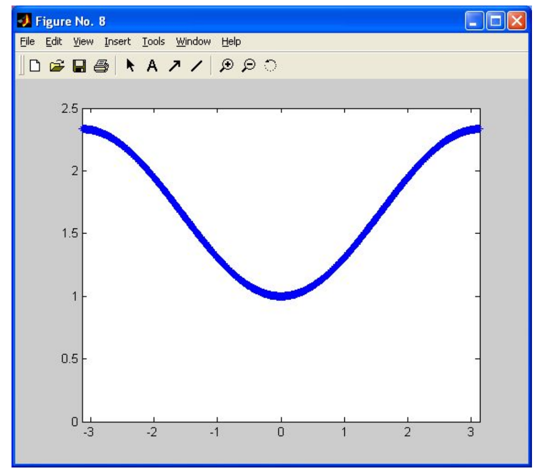

$$F[m]=\sum_{k=0}^{M-1} f[k]e^{-\pi i\left(\frac{km}{M}\right)}$$

$$=2-\frac{1}{3}\left(1+2\cos\left(\frac{\pi m}{M}\right)\right)$$

high-pass filter

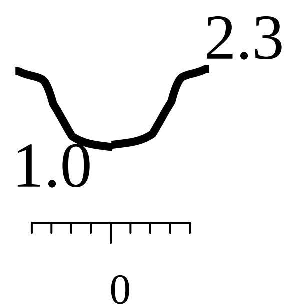

Questions?

Sampling and aliasing

More on Samples

- In signal processing, the process of mapping a continuous function to a discrete one is called *sampling*
- The process of mapping a continuous variable to a discrete one is called *quantization*
- To represent or render an image using a computer, we must both sample and quantize
  - Now we focus on the effects of sampling and how to fight them

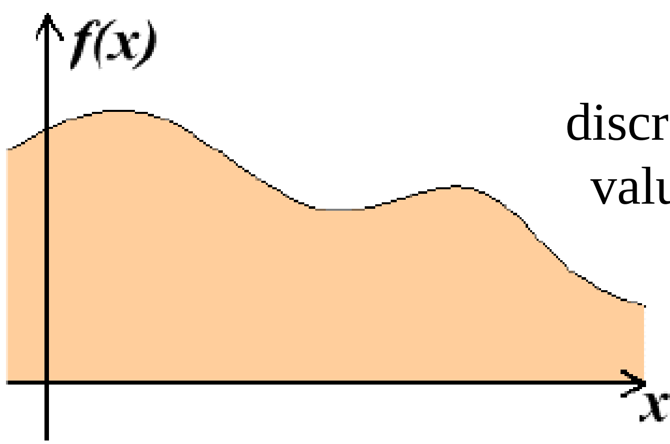

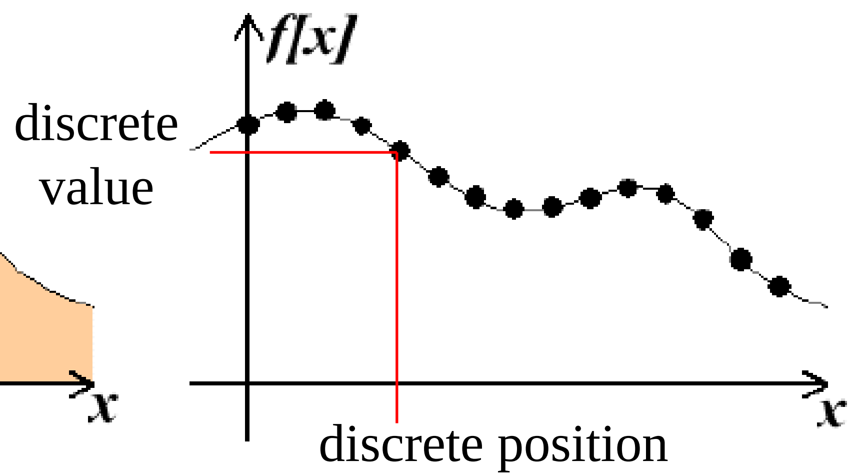

Sampling in the Frequency Domain

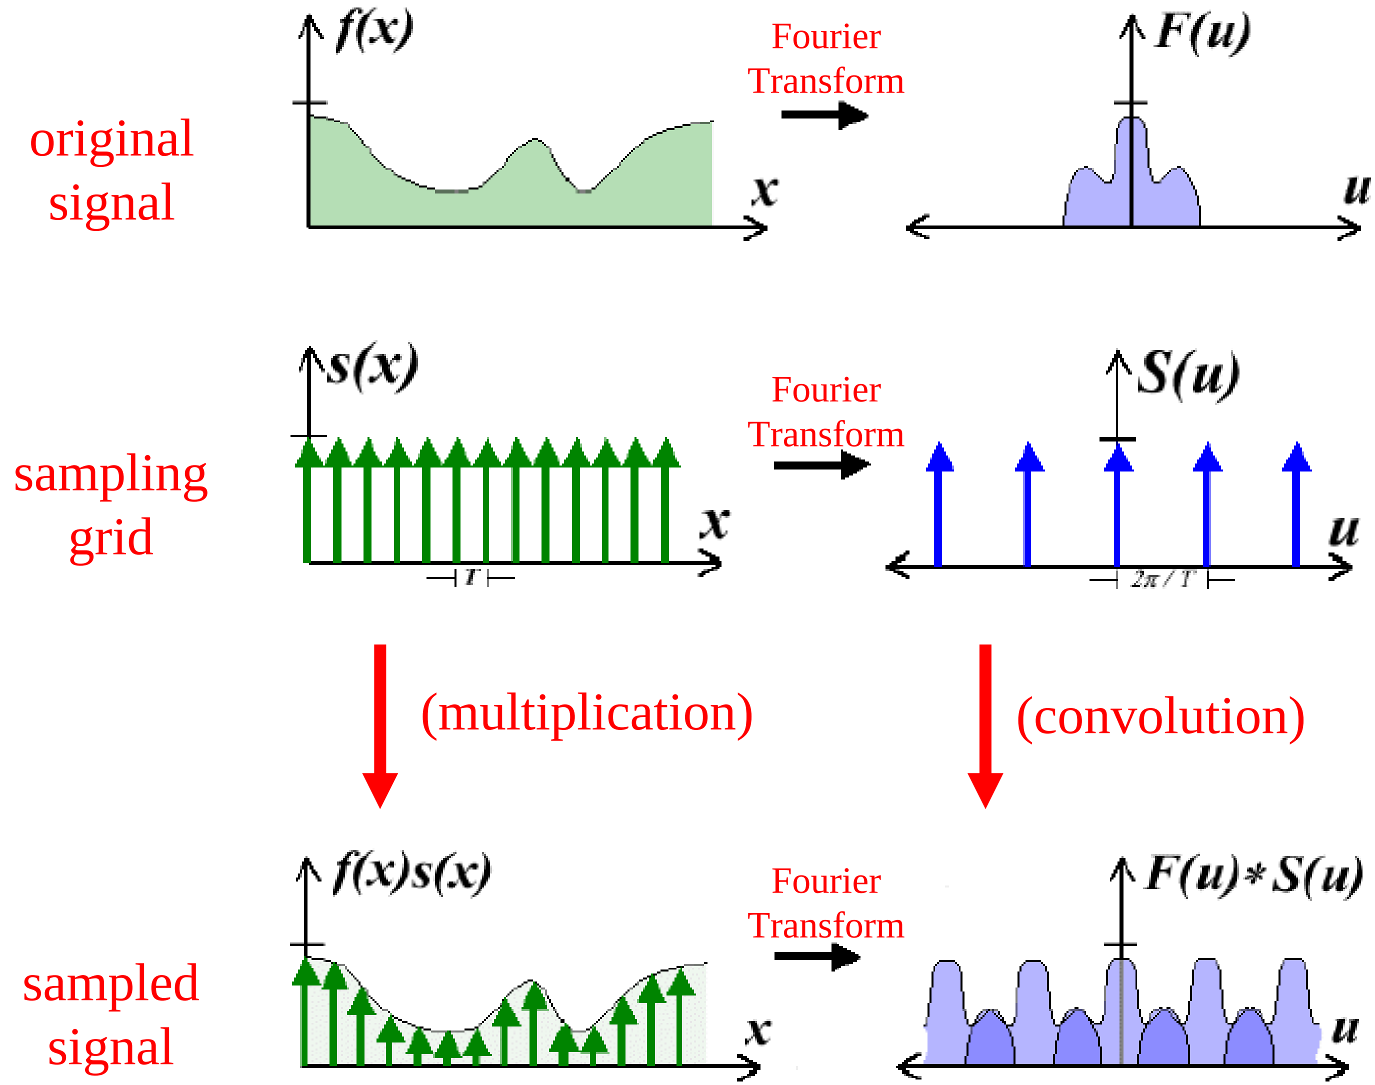

Reconstruction

- If we can extract a copy of the original signal from the frequency domain of the sampled signal, we can reconstruct the original signal!

- But there may be overlap between the copies.

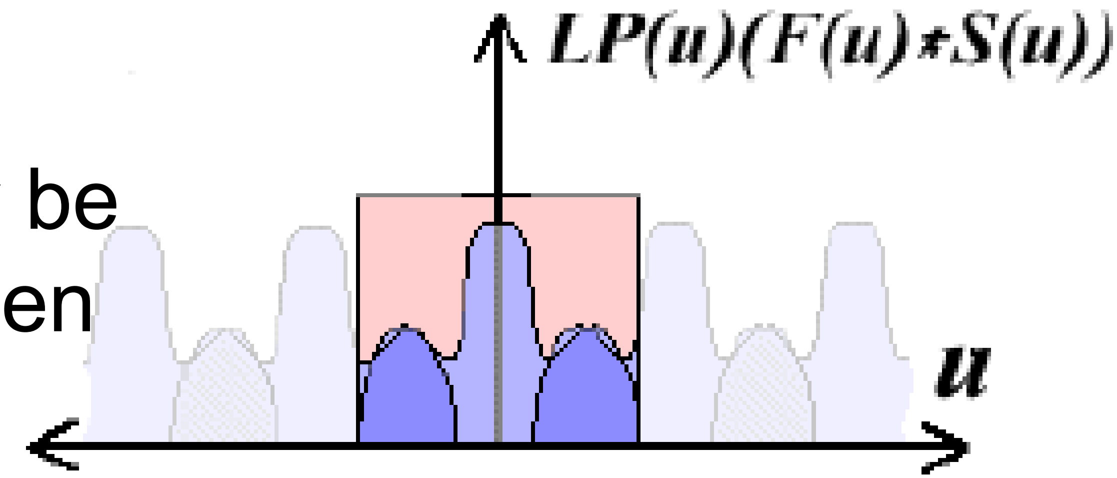

Guaranteeing Proper Reconstruction

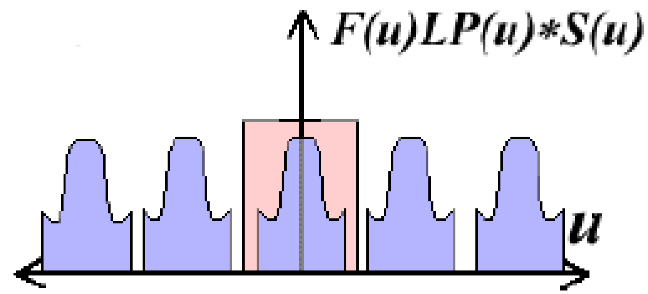

- Separate by removing high frequencies from the original signal (low pass pre-filtering)
- Separate by increasing the sampling density

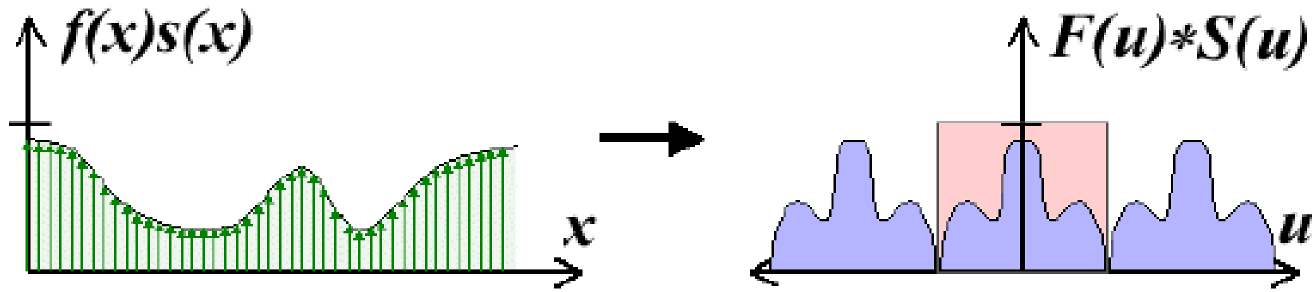

- If we can't separate the copies, we will have overlapping frequency spectrum during reconstruction → *aliasing*.

Sampling Theorem

- When sampling a signal at discrete intervals, the sampling frequency must be *greater than twice* the highest frequency of the input signal in order to be able to reconstruct the original perfectly from the sampled version (Shannon, Nyquist, Whittaker, Kotelnikov)