Filtering in the Fourier Domain

Tolga Tasdizen

ECE

University of Utah

2D convolution

$$f(x,y) \star h(x,y) = \sum_{m=0}^{M-1} \sum_{n=0}^{N-1} f(m,n) h(x-m, y-n)$$

- 2D convolution theorem:

$$f(x,y) \star h(x,y) \Longleftrightarrow F(u,v) H(u,v)$$

- Equivalently:

$$f(x,y) \star h(x,y) = \mathcal{F}^{-1} \left\{ \mathcal{F}\{f(x,y)\} \mathcal{F}\{h(x,y)\} \right\}$$

Impulse response

- Filter H(u,v)
- If the input image $f(x, y) = \delta(x, y)$
- The filtered output will be

$$h(x, y) = \mathcal{F}^{-1}\{H(u, v)\}$$

is called the impulse response of H(u,v)

Frequency filtering

- Example: Sobel derivative filter
  – Take Fourier transform of filter and image, multiply, invert
- But we can also design/ specify filters directly in the frequency domain
- In both cases we need to account for circular convolution via zero padding

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

$h(x,y)$

Circular convolution

- Implementing filtering via DFTs and multiplication in the Frequency domain implies an assumption of periodicity
  – This is due to the periodicity property of the DFT we discussed earlier
- Care must be taken in implementing filtering in the frequency domain

- Non-periodic implementation in spatial domain
- Assumes values outside the domain of ***f*** and ***h*** are 0
- Note also the length of the output sequence is the sum of the lengths of ***f*** and ***h*** minus 1

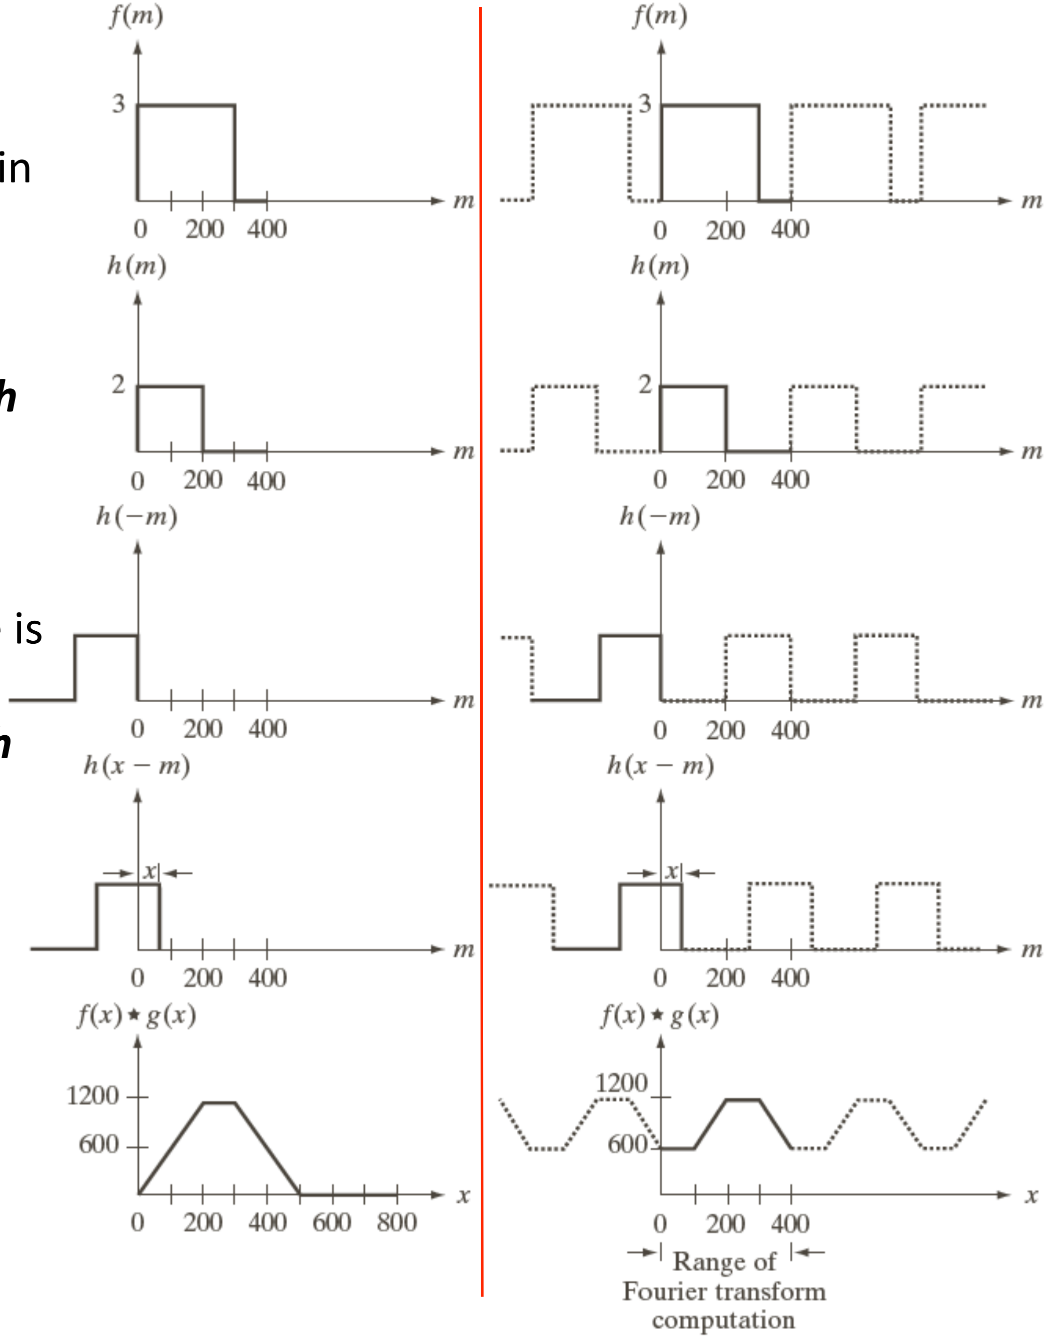

- Periodic implementation
- This is the result that would be obtained if implementing in frequency domain via DFTs
- Incorrect result is obtained because the periods of ***f*** and ***h*** interfere with each other (*wraparound error*)

© 1992–2008 R. C. Gonzalez & R. E. Woods

Zero padding solution

- If f(x) has A samples and h(x) has B samples, pad enough zeros to both functions so they both have length P=A+B-1 or longer
- 2D: f(x,y) size A x B, h(x,y) size C x D
  - Pad both to size P x Q where
  - P is greater than or equal to A + C -1
  - Q is greater than or equal to B + D -1

Frequency filtering steps

Input image f(x,y)
Size M x N

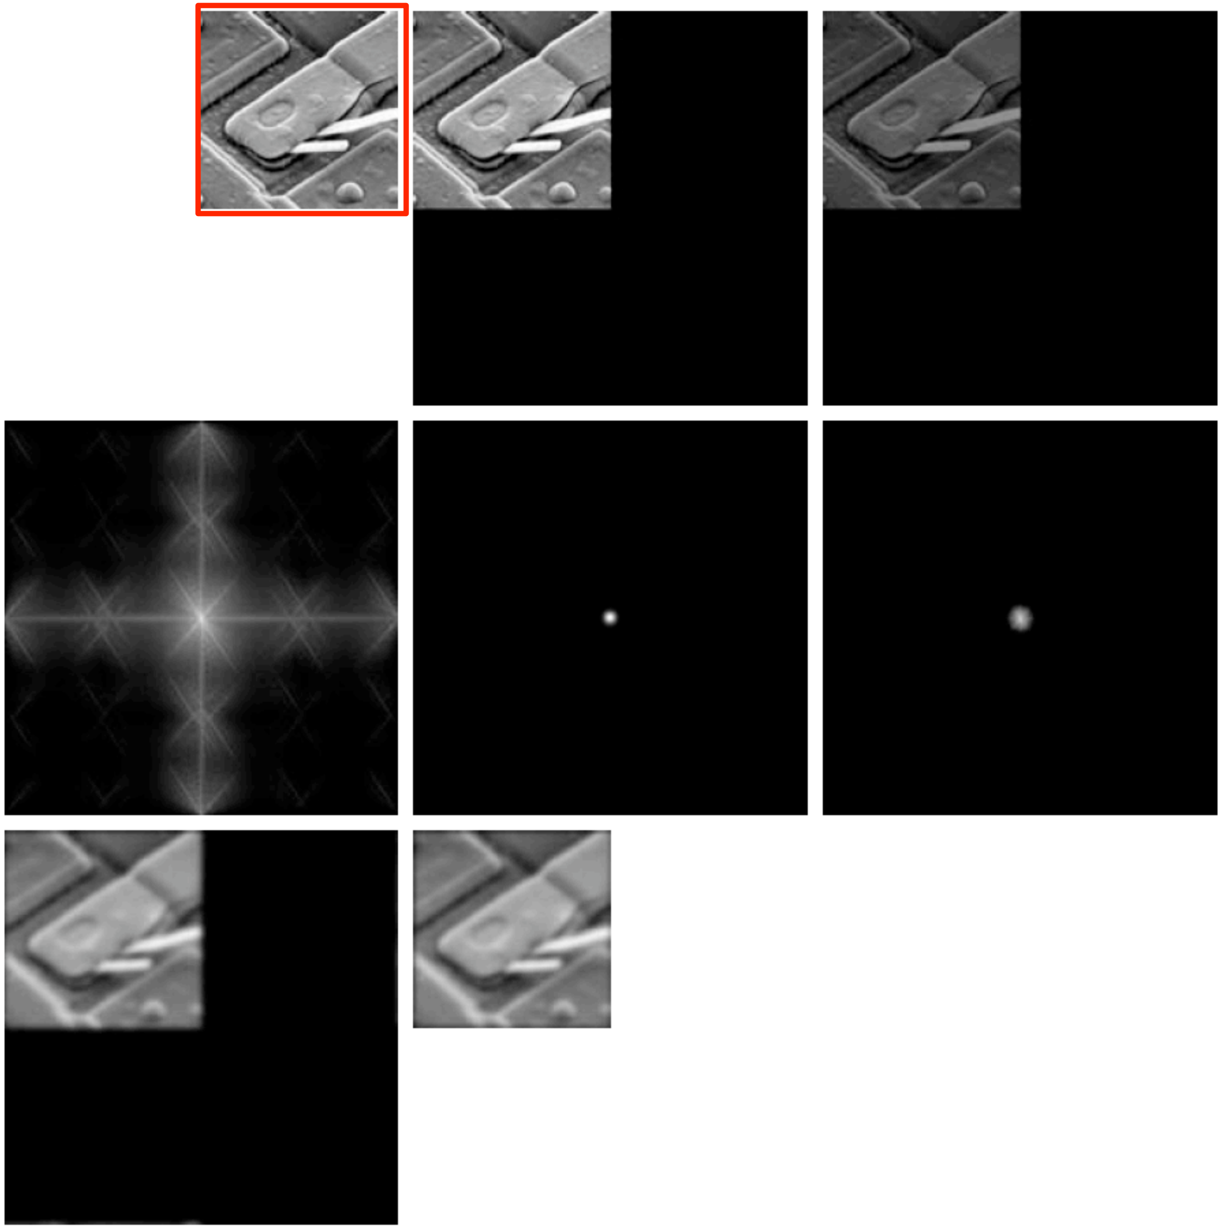

© 1992–2008 R. C. Gonzalez & R. E. Woods

Pad input image $f(x,y)$ to size 2M x 2N. Call result $f_p(x,y)$
Let P=2M, Q=2N

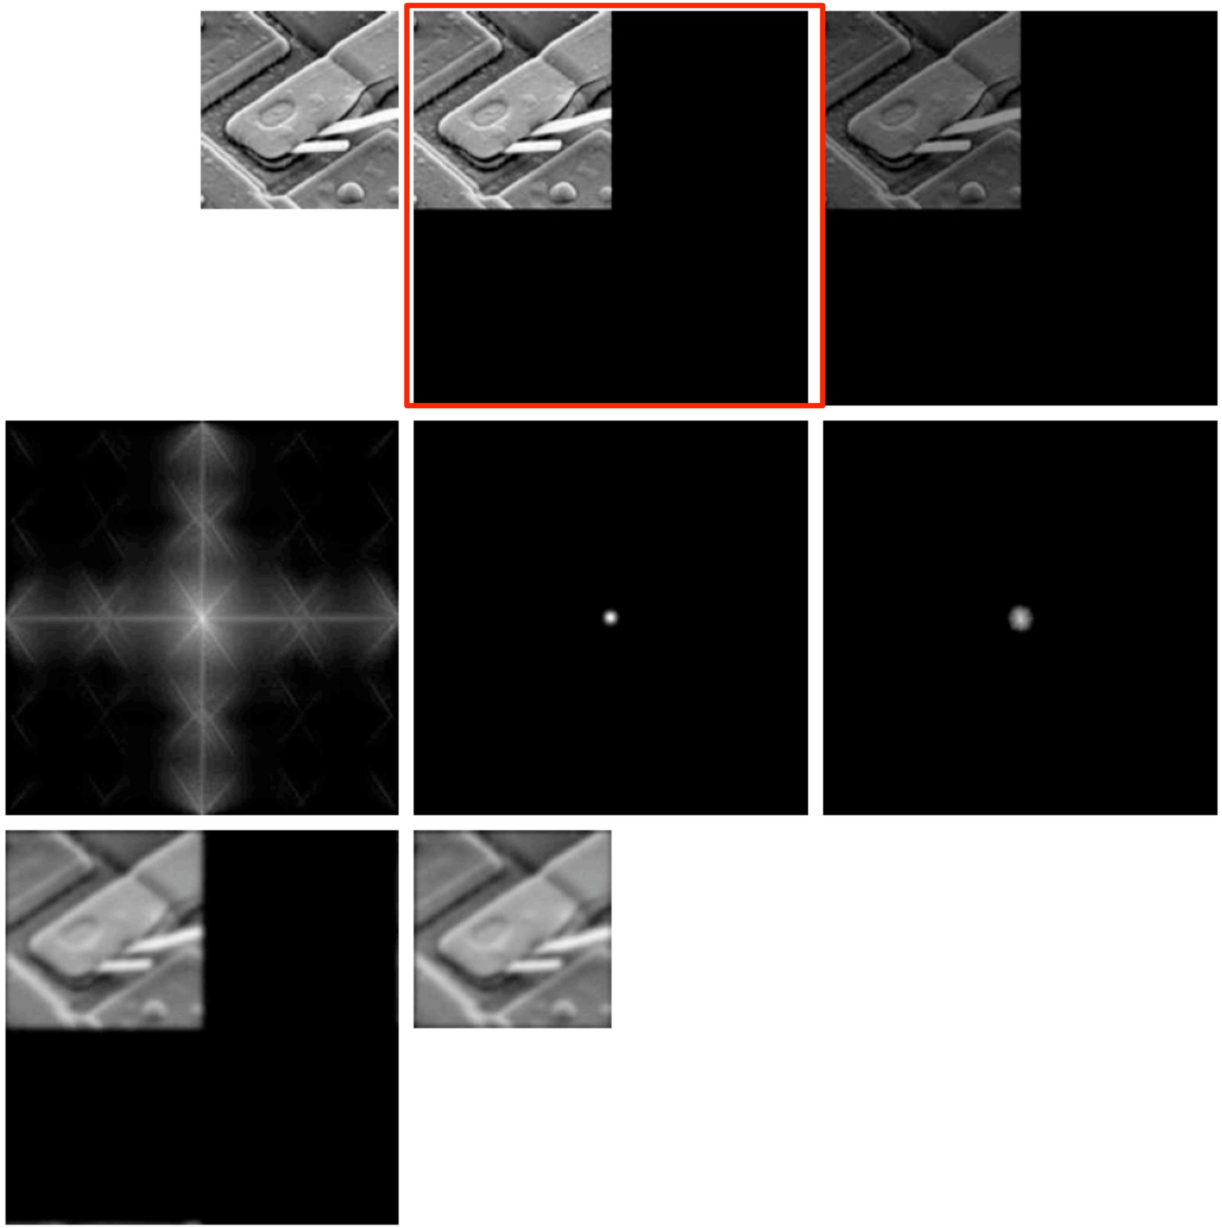

© 1992–2008 R. C. Gonzalez & R. E. Woods

Let $f'_p(x,y)=(-1)^{x+y}f_p(x,y)$
This will center the Fourier transform

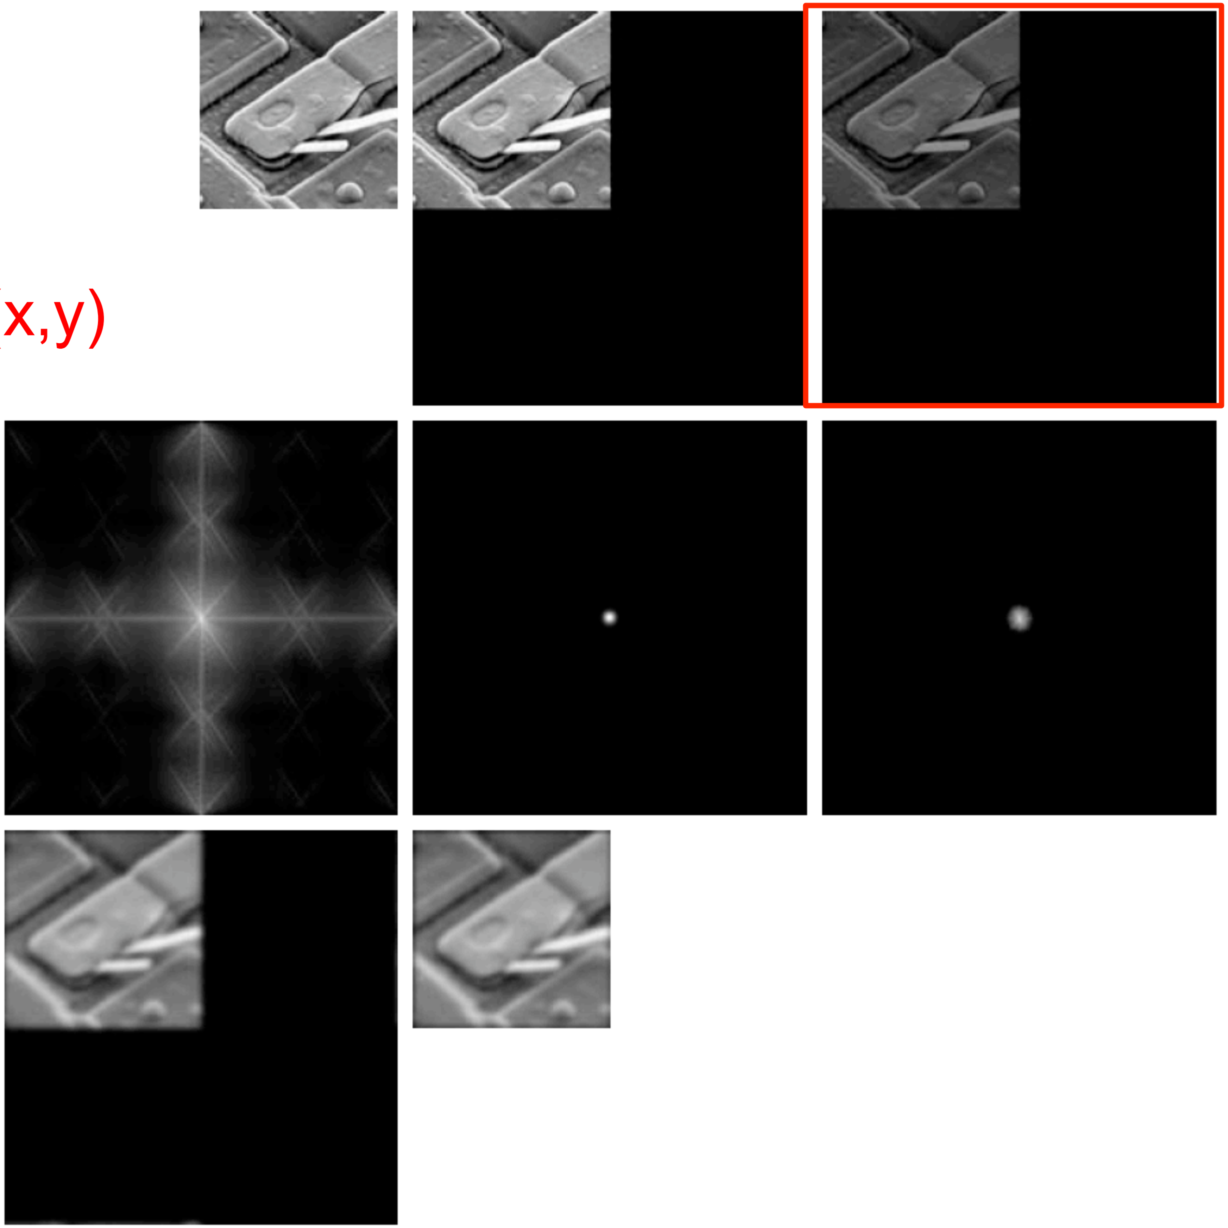

© 1992–2008 R. C. Gonzalez & R. E. Woods

Take DFT of $f'_p(x,y)$ to get $F(u,v)$

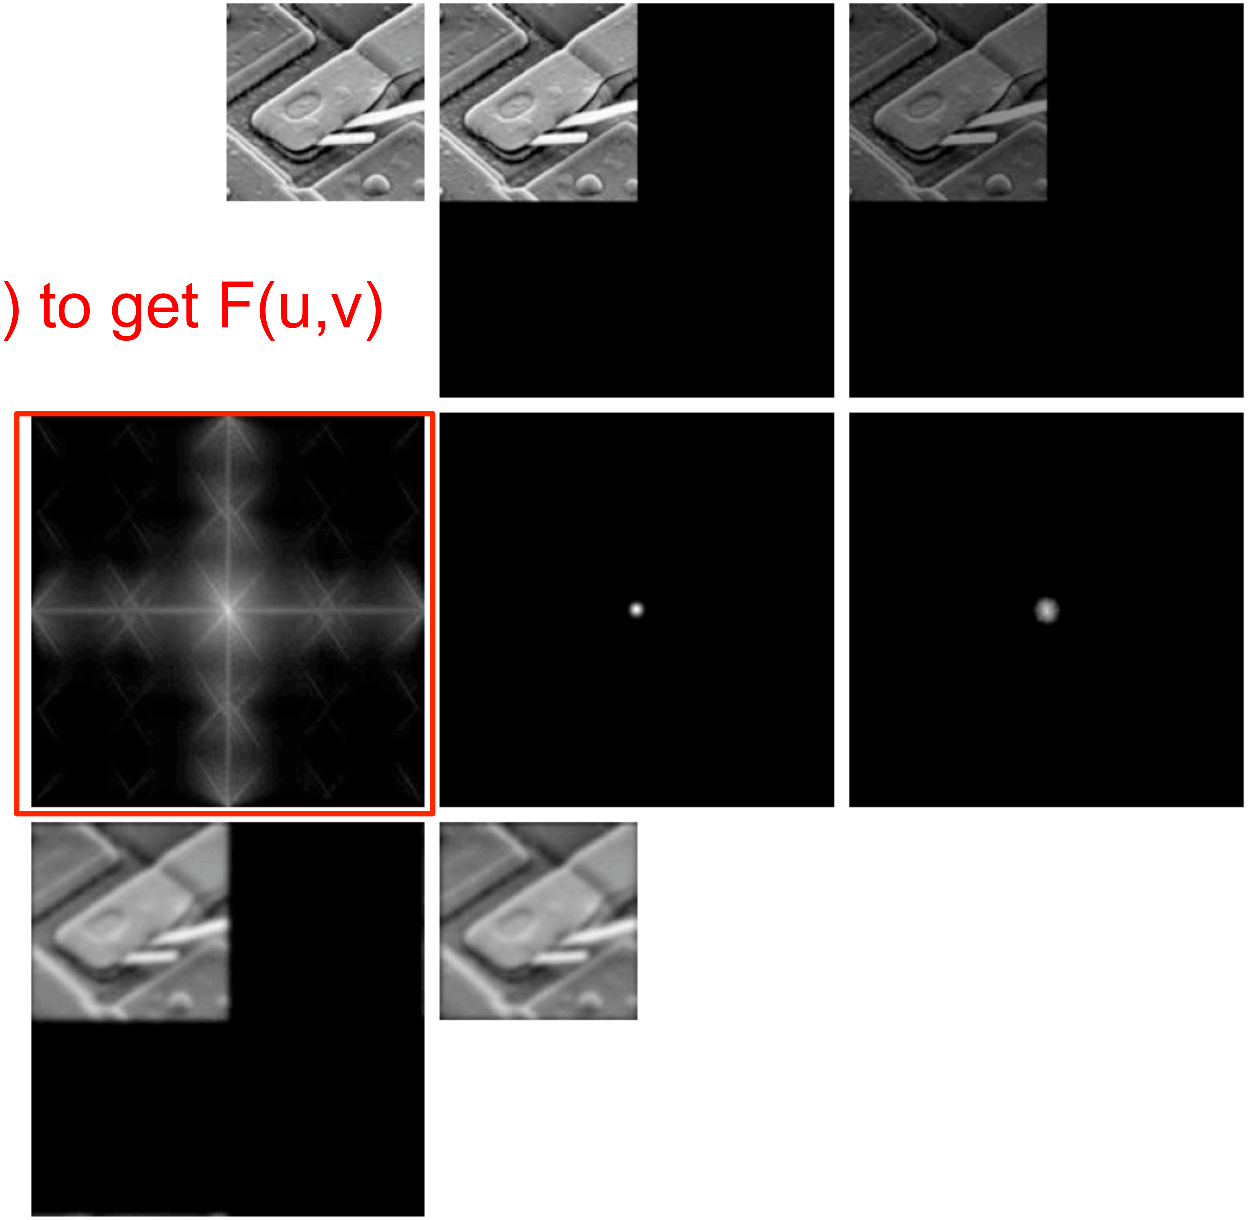

© 1992–2008 R. C. Gonzalez & R. E. Woods

Generate a real, symmetric H(u,v) of size 2M x 2N with center at coordinate M x N

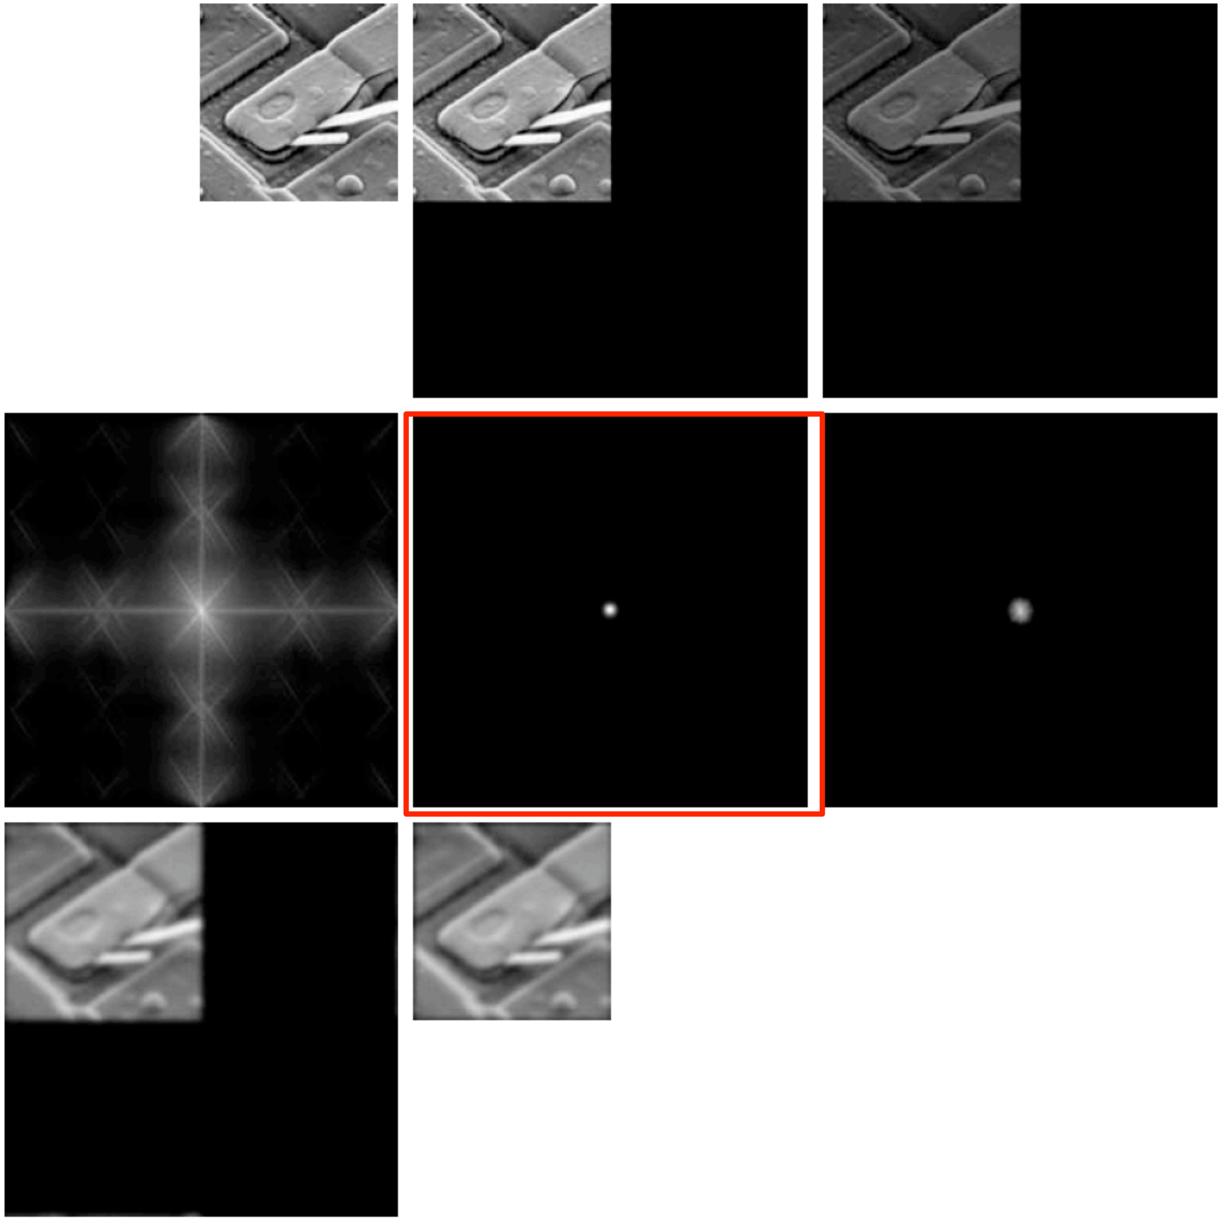

© 1992–2008 R. C. Gonzalez & R. E. Woods

Compute the product
$G(u,v)=F(u,v)H(u,v)$

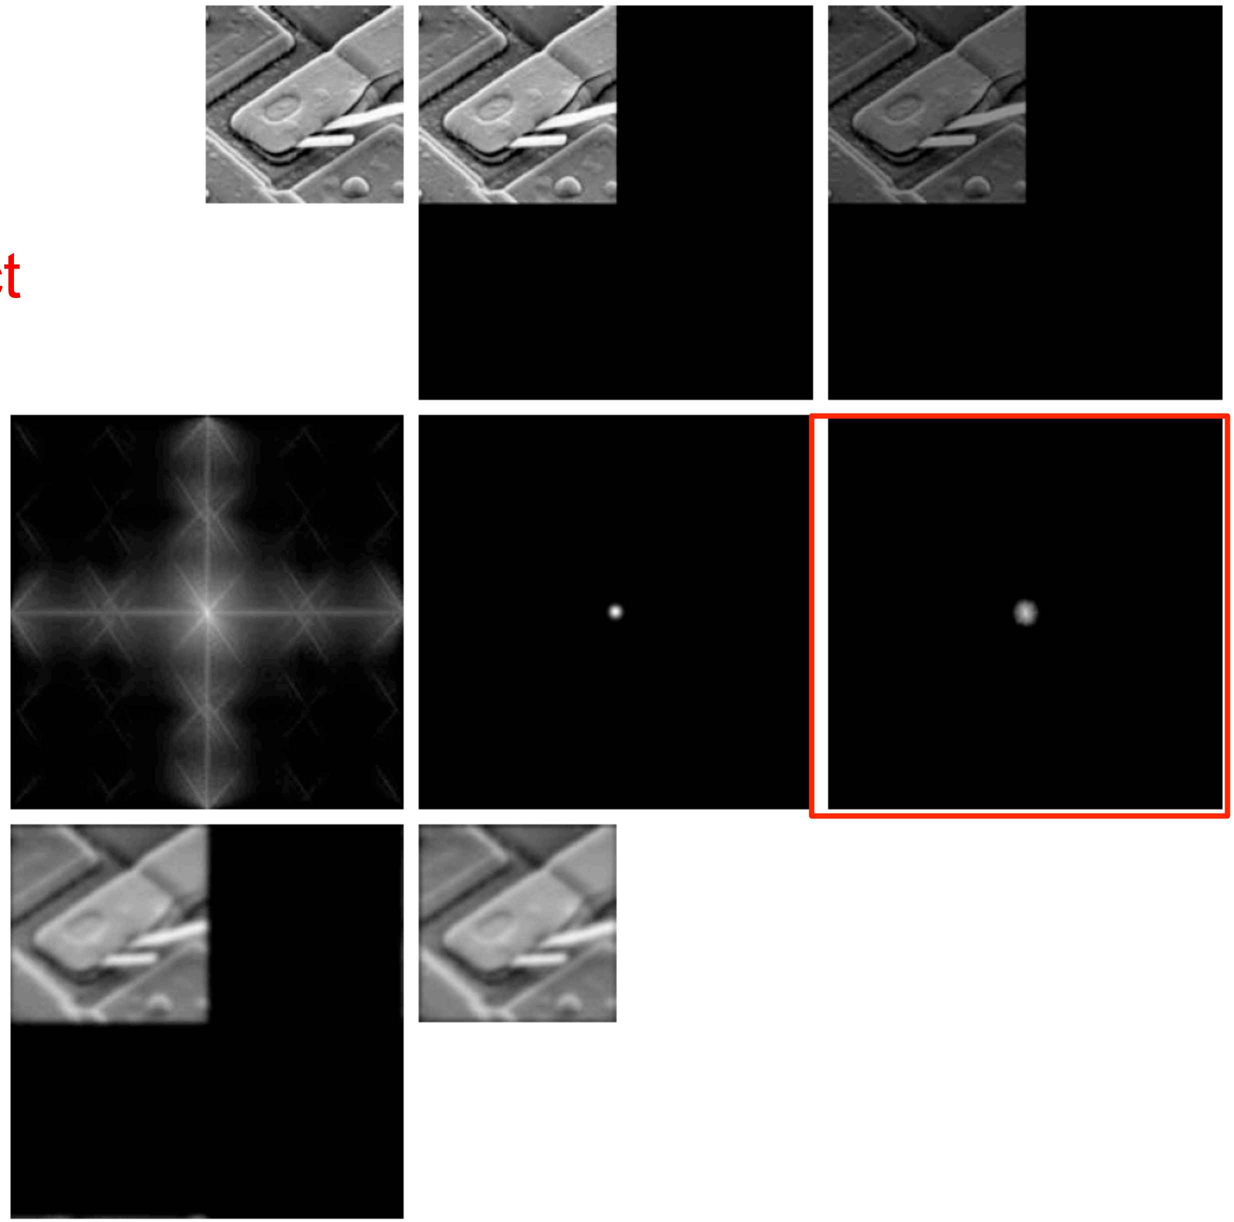

© 1992–2008 R. C. Gonzalez & R. E. Woods

$$g_p(x,y) = \Big(\mathrm{Real}\big(F^{-1}[G(u,v)]\big)\Big)\,(-1)^{x+y}$$

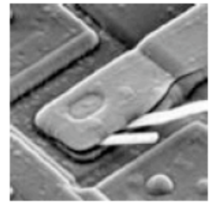

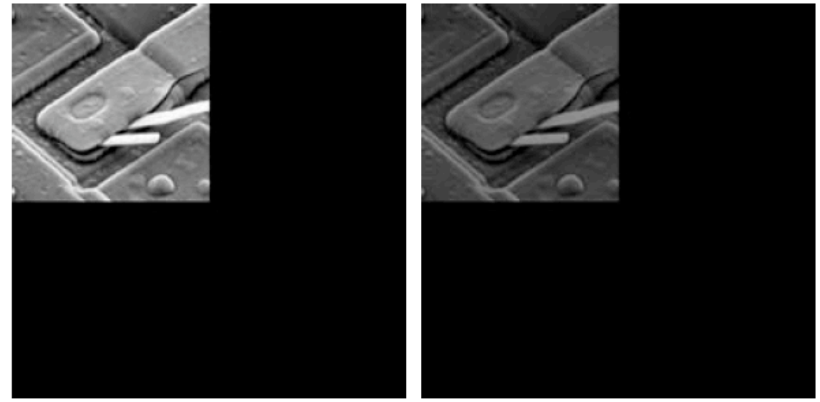

We take the real part for ignoring parasitic complex components resulting from computational inaccuracy

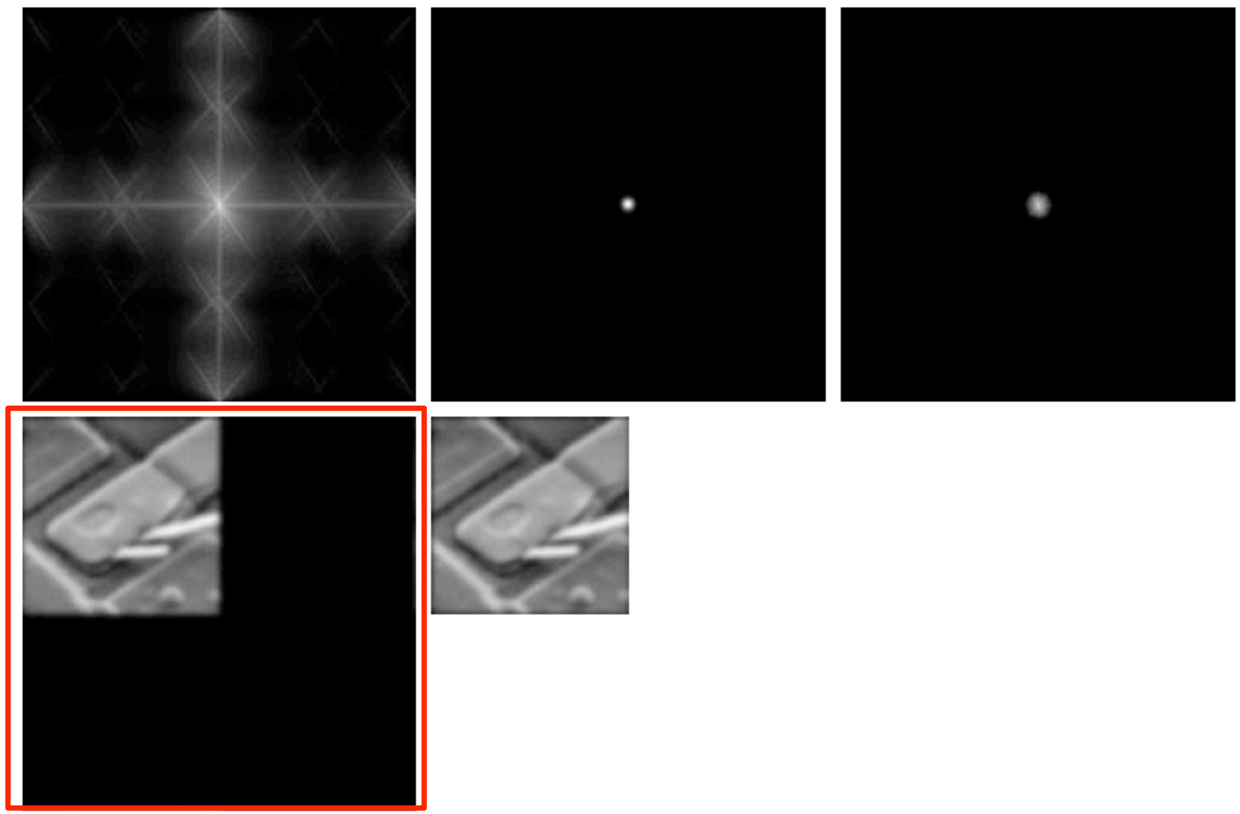

© 1992–2008 R. C. Gonzalez & R. E. Woods

Crop top left MxN image

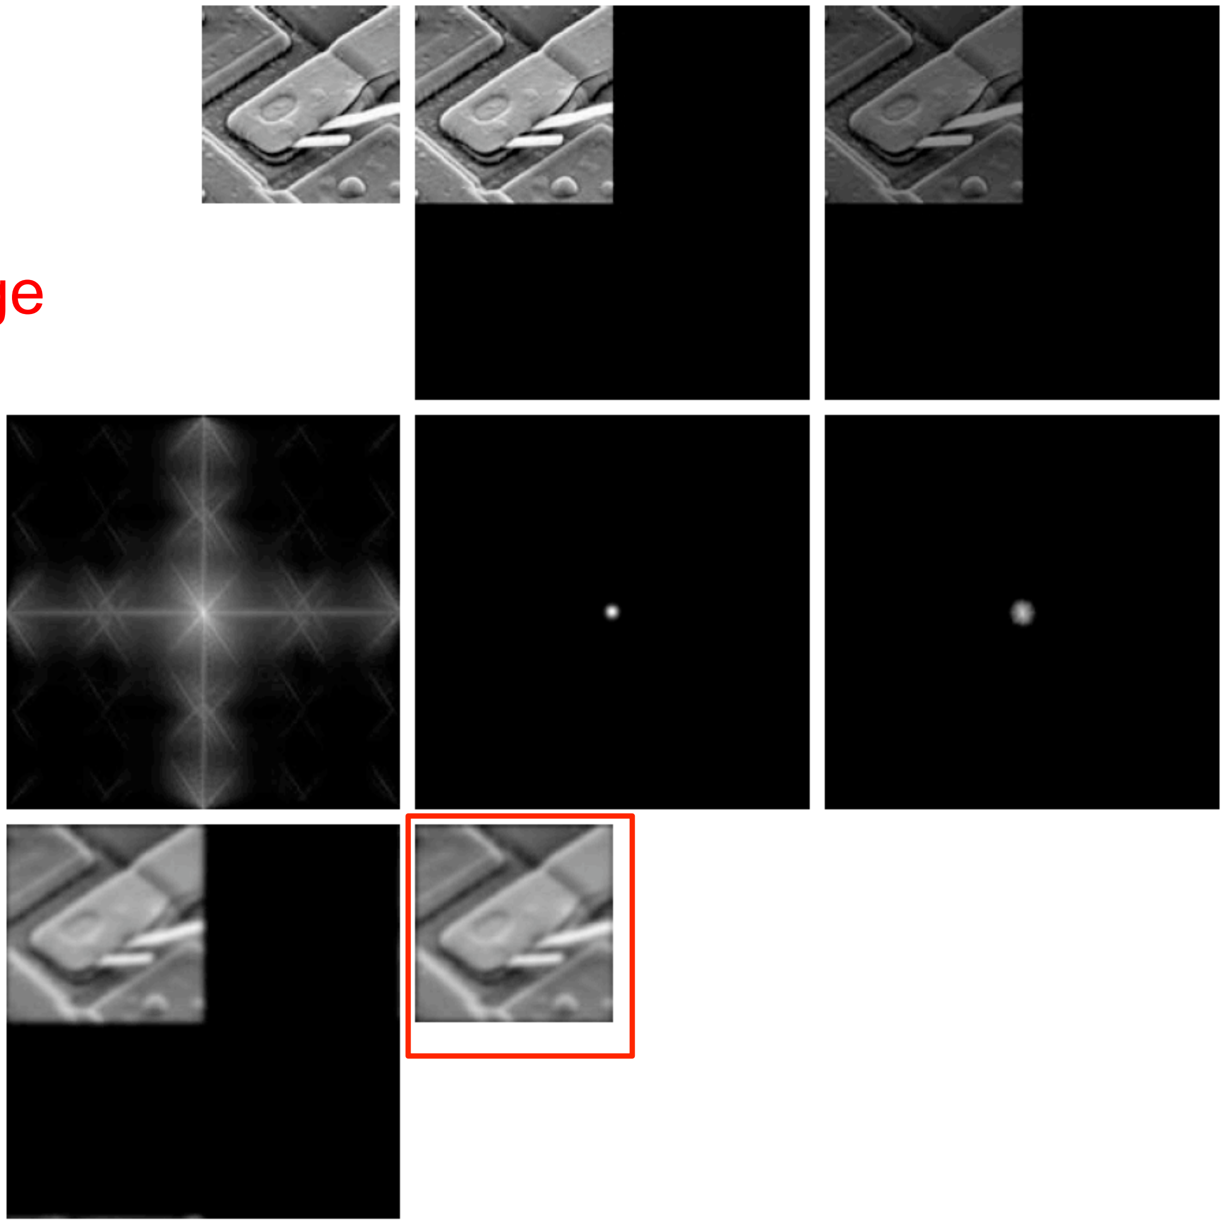

© 1992–2008 R. C. Gonzalez & R. E. Woods

Frequency domain basics

- Low frequencies - Slowly varying spatial intensities
- High frequencies - Abrupt changes in intensity: edges, noise, etc...

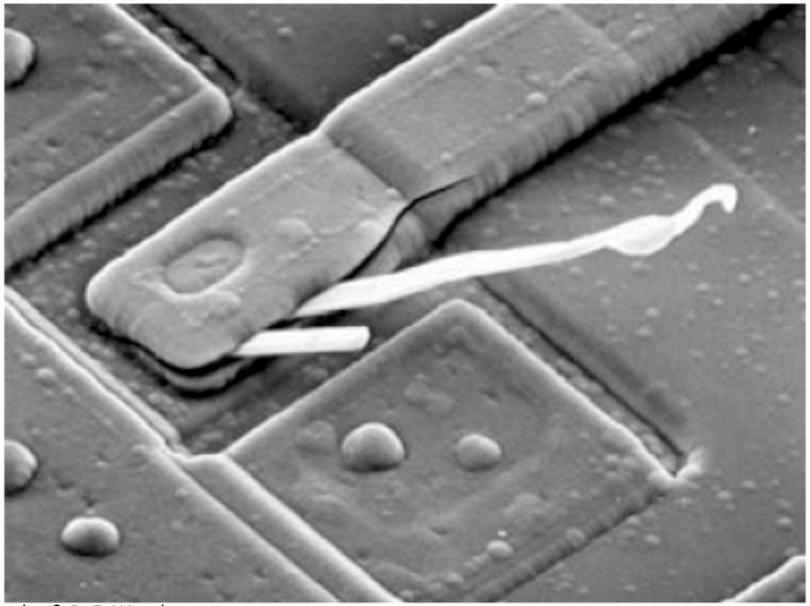

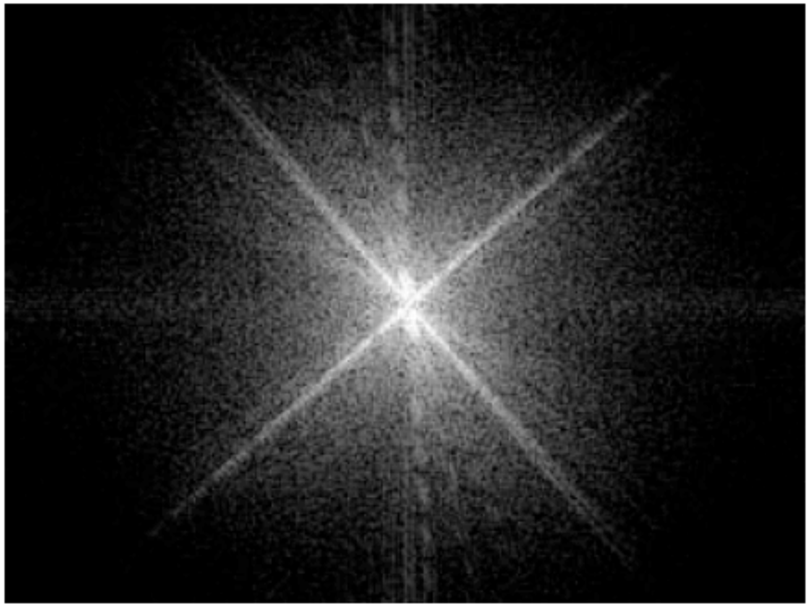

© 1992–2008 R. C. Gonzalez & R. E. Woods

Low-Pass Filter

- Reduce/eliminate high frequencies
- Applications
  - Noise reduction
    - uncorrelated noise is broad band
    - Images have sprectrum that focus on low frequencies

Total image power P_T is the sum of $P(u,v)$ over all u,v

% power retained by ILPF with cutoff D_o

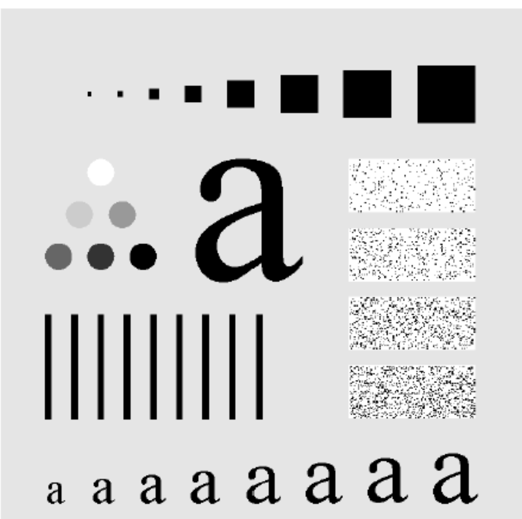

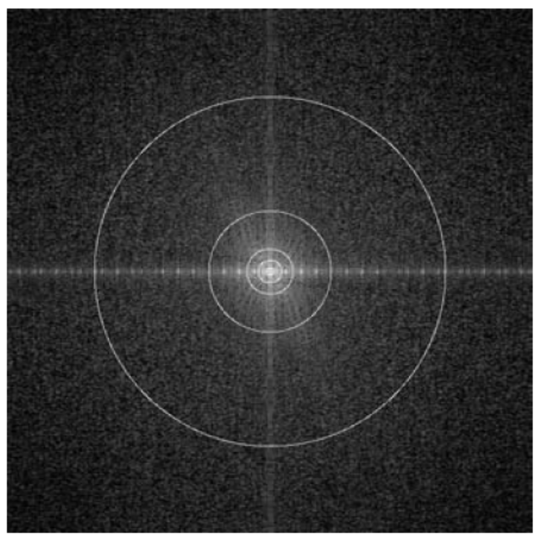

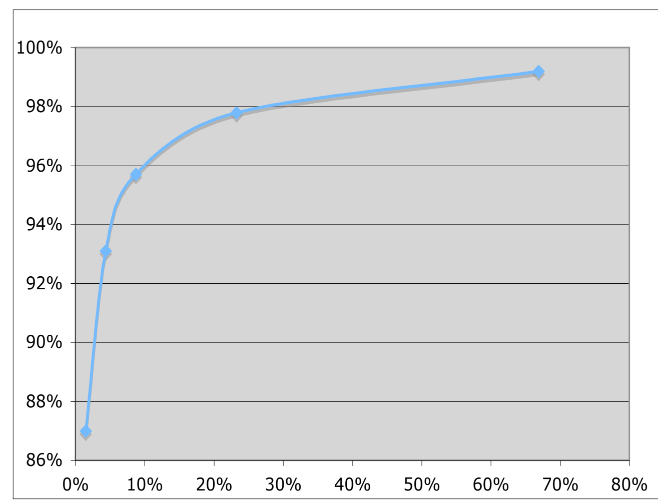

Power spectrum $P(u,v) = | F(u,v) |^2$

Extending Filters to 2D (or higher)

- Two options
  - Separable
    - H(s) -> H(u)H(v)
    - Easy, analysis
  - Rotate
    - H(s) -> H((u2 + v2)1/2)
    - Rotationally invariant

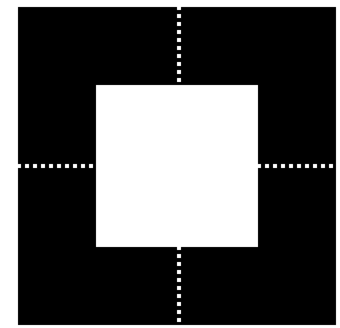

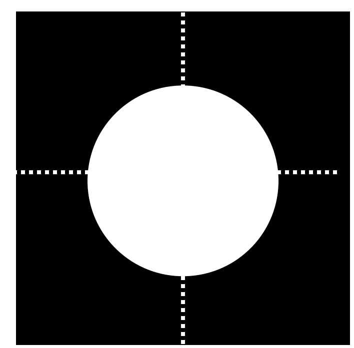

Ideal Lowpass LP - circular

- Passes without attenuation all frequencies within a radius D_o from the origin and completely cuts off the rest

$$H(u,v) = \begin{cases} 1 & if \quad D(u,v) \le D_o \\ 0 & if \quad D(u,v) > D_o \end{cases}$$

$$D(u,v) = \sqrt{(u - P/2)^2 + (v - Q/2)^2}$$

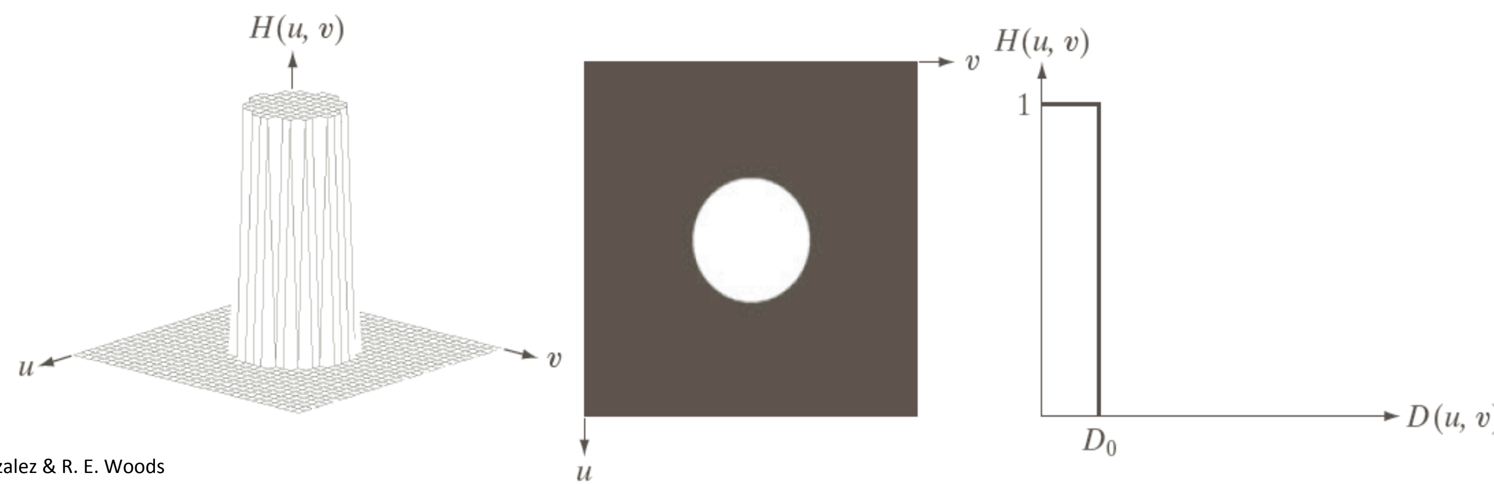

© 1992–2008 R. C. Gonzalez & R. E. Woods

- Retaining less power corresponds to more blurring
- Increasing cutoff frequency increases amount of detail retained
- What else do you see in some of these filtered images?

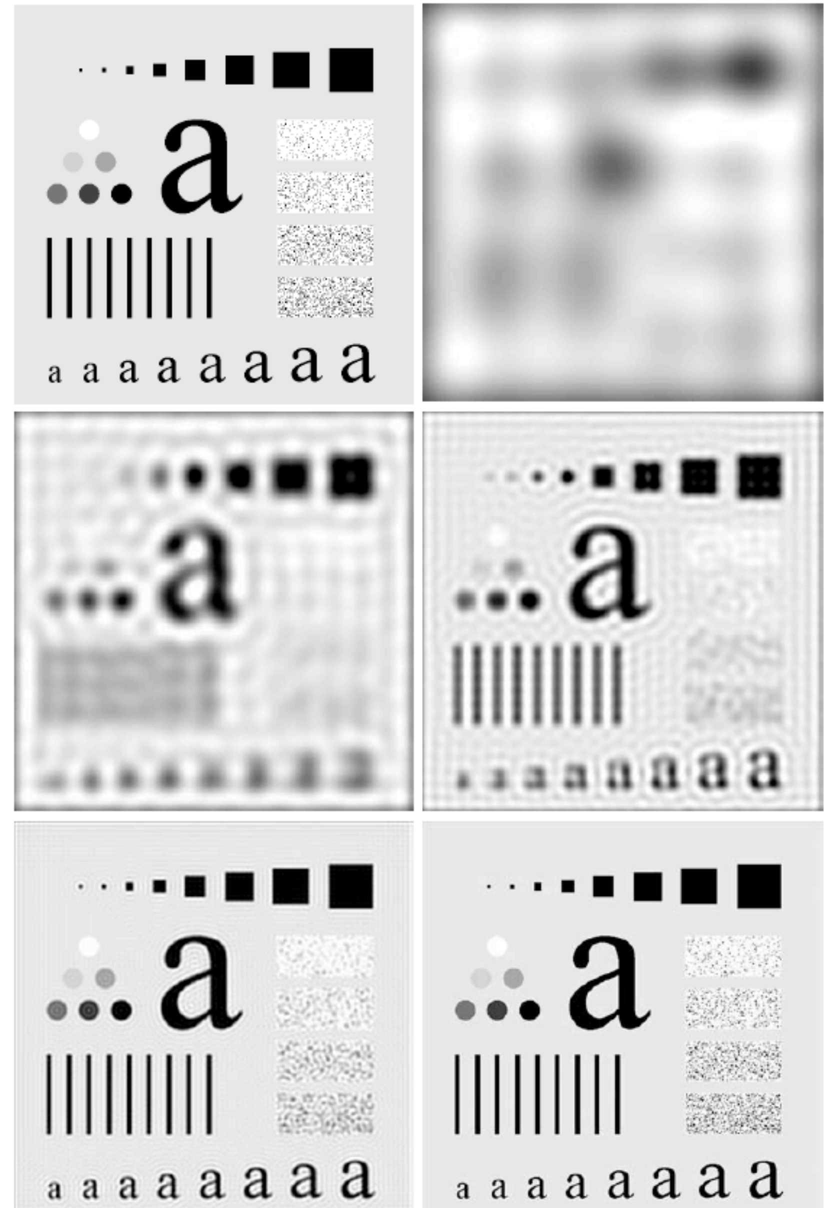

FIGURE 4.42 (a) Original image. (b)–(f) Results of filtering using ILPFs with cutoff frequencies set at radii values 10, 30, 60, 160, and 460, as shown in Fig. 4.41(b). The power removed by these filters was 13, 6.9, 4.3, 2.2, and 0.8% of the total, respectively.

© 1992–2008 R. C. Gonzalez & R. E. Woods

Ringing artifact of ILPF

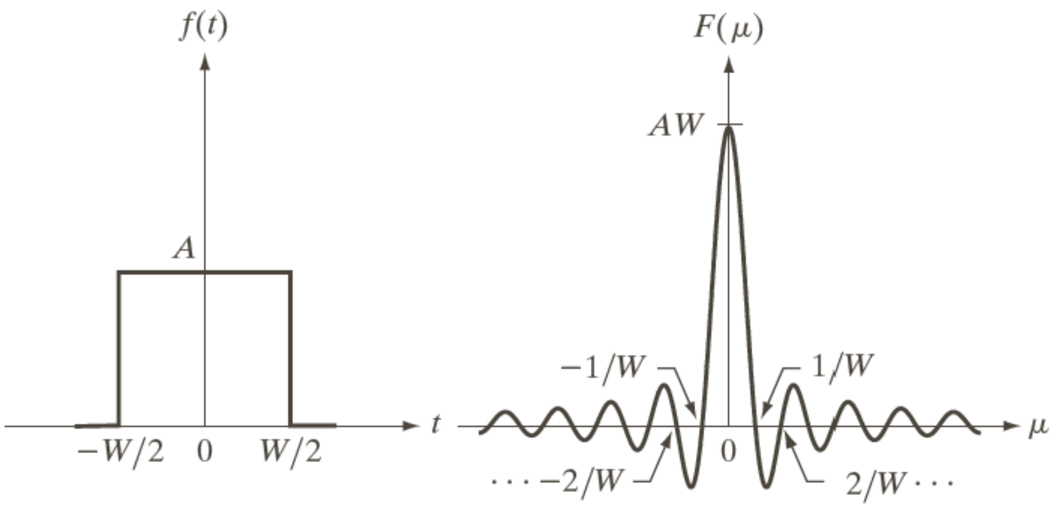

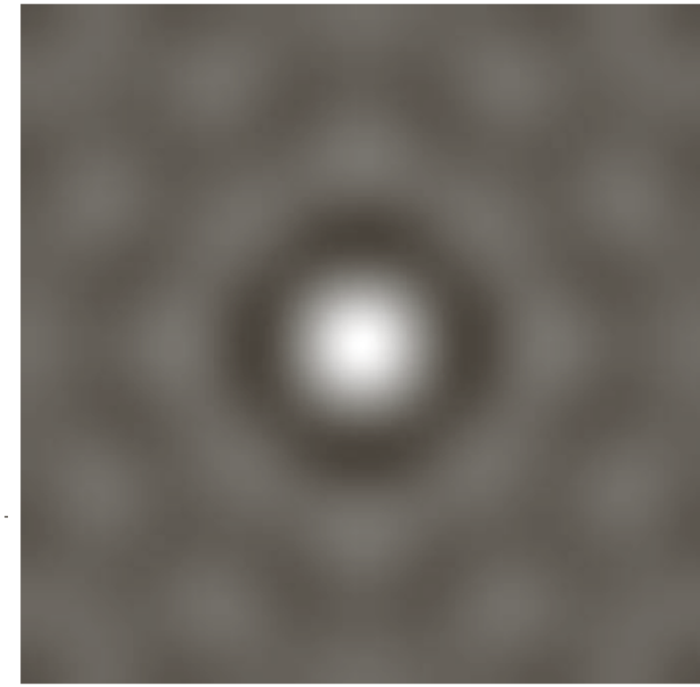

Cutoff freq

Ringing – Gibbs phenomenon

In 2 dimensions

Ideal Low-Pass
Rectangle With Cutoff of 2/3

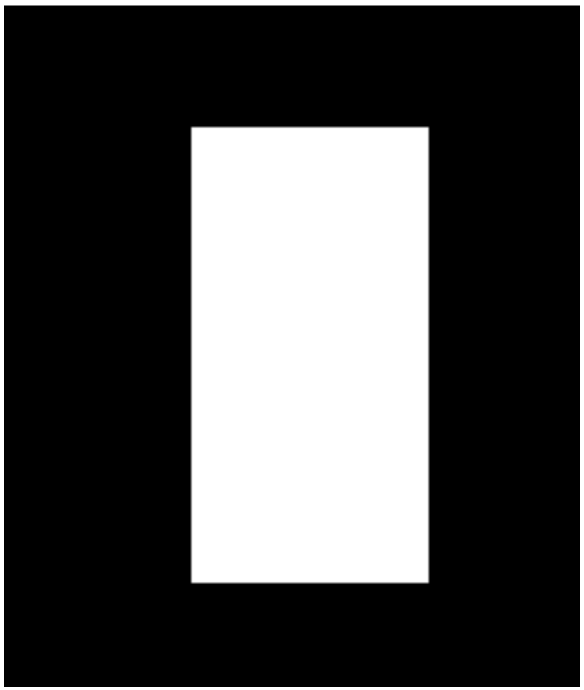

Image

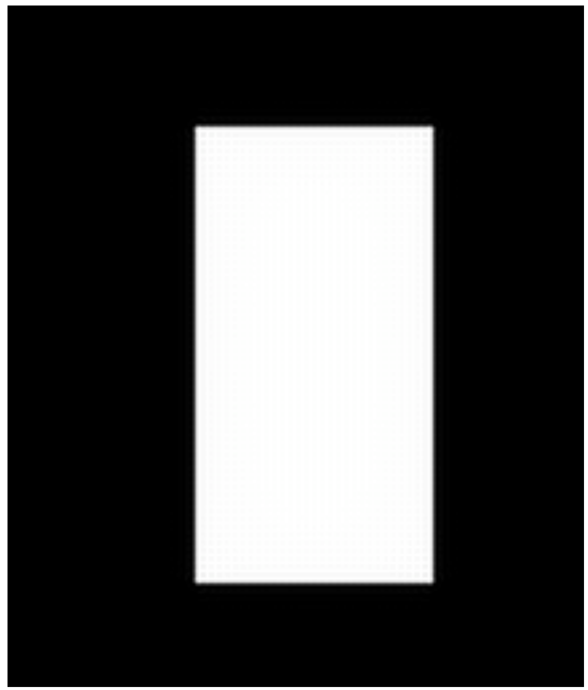

Filtered

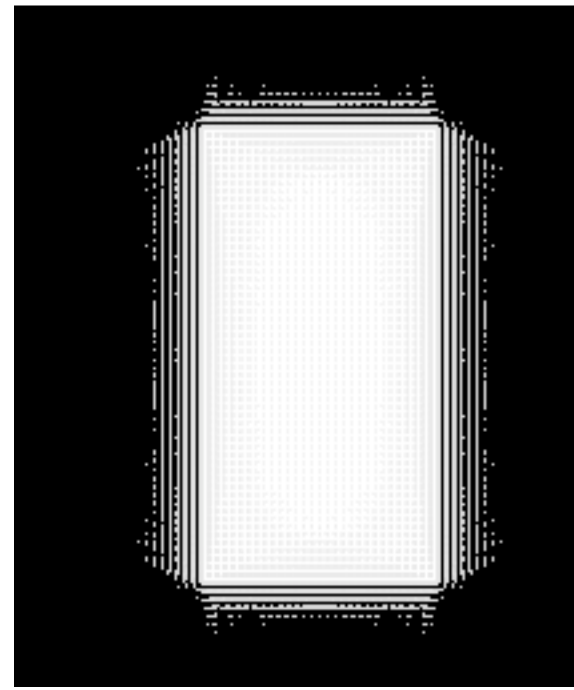

Filtered + HE

Ideal LP – 1/3

Ideal LP – 2/3

Lowpass filter with less ringing

- Avoid sharp discontinuity in H(u,v)
- Butterworth Lowpass Filter

$$H(u,v) = \frac{1}{1 + \left(\frac{D(u,v)}{D_o}\right)^{2n}}$$

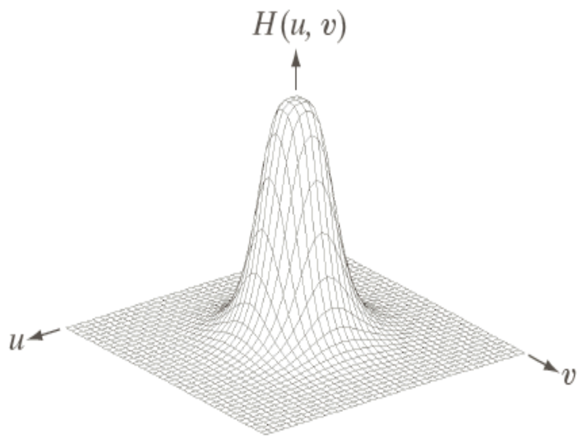

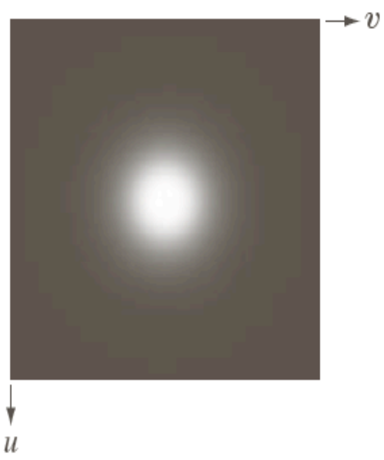

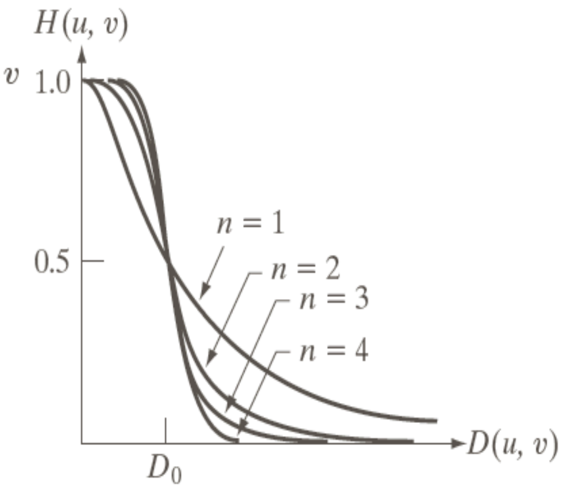

© 1992–2008 R. C. Gonzalez & R. E. Woods

Butterworth properties

- No ringing in spatial domain for n = 1
- Imperceptible ringing for order n = 2
- Ringing can be significant for n > 2
- Limit as n increases, same as ILPF

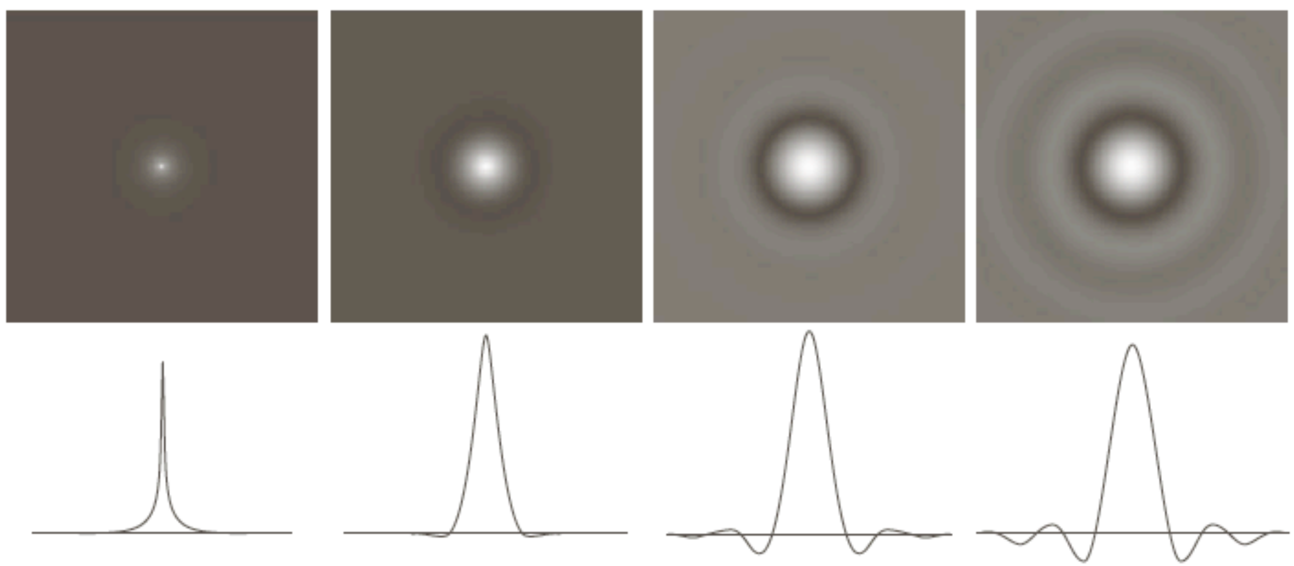

a b c d

FIGURE 4.46 (a)–(d) Spatial representation of BLPFs of order 1, 2, 5, and 20, and corresponding intensity profiles through the center of the filters (the size in all cases is 1000 × 1000 and the cutoff frequency is 5). Observe how ringing increases as a function of filter order.

© 1992–2008 R. C. Gonzalez & R. E. Woods

Butterworth n=2

- No visible ringing artifacts

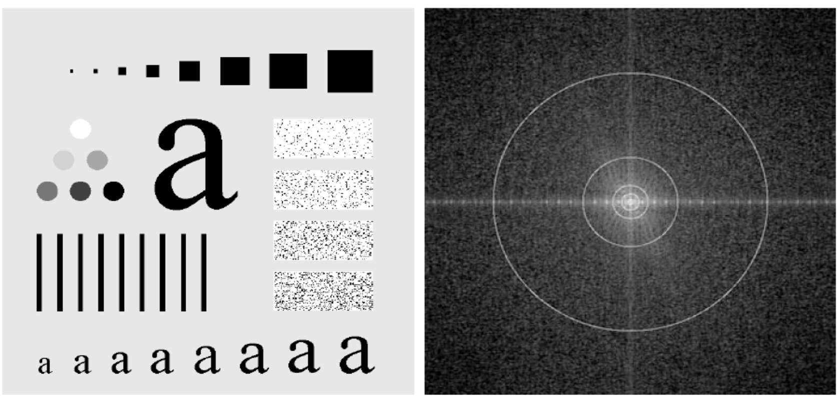

a b

FIGURE 4.41 (a) Test pattern of size 688 × 688 pixels, and (b) its Fourier spectrum. The spectrum is double the image size due to padding but is shown in half size so that it fits in the page. The superimposed circles have radii equal to 10, 30, 60, 160, and 460 with respect to the full-size spectrum image. These radii enclose 87.0, 93.1, 95.7, 97.8, and 99.2% of the padded image power, respectively.

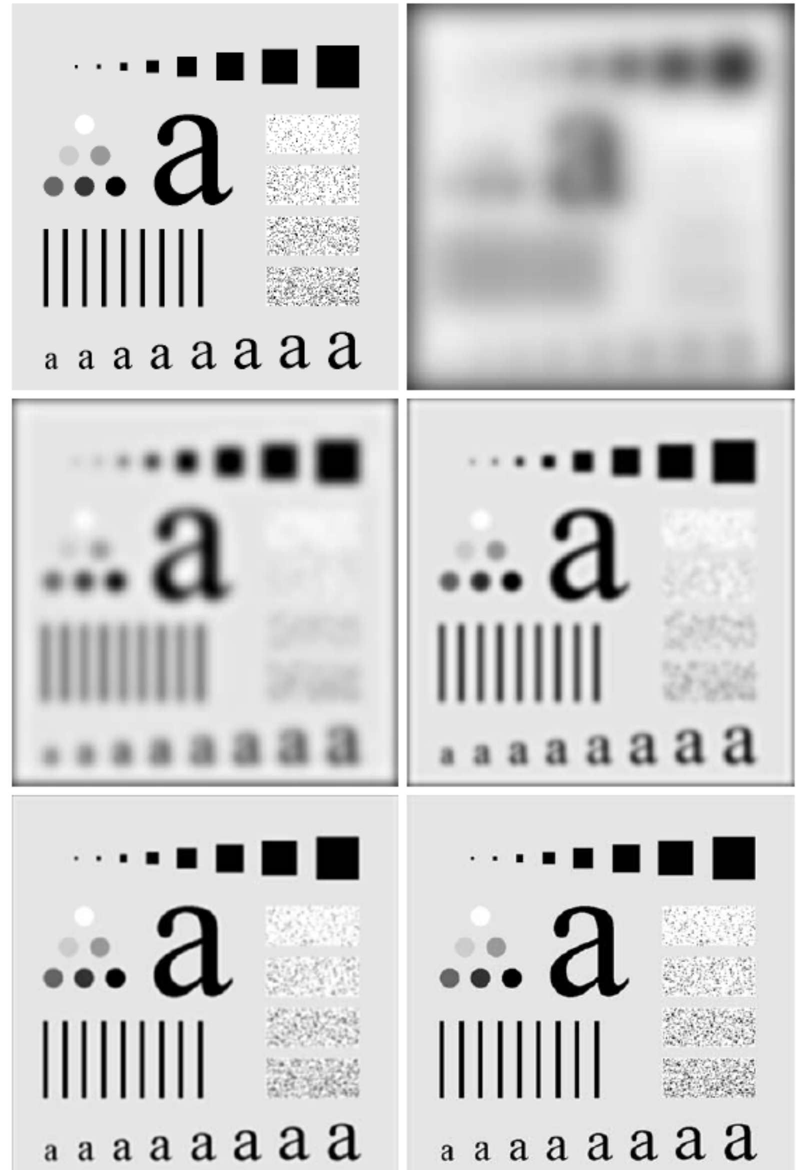

a b
c d
e f

FIGURE 4.45 (a) Original image. (b)–(f) Results of filtering using BLPFs of order 2, with cutoff frequencies at the radii shown in Fig. 4.41. Compare with Fig. 4.42.

© 1992–2008 R. C. Gonzalez & R. E. Woods

Butterworth - 1/3

Butterworth – 2/3

Butterworth vs Ideal LP

Gaussian lowpass filter (GLPF)

- A Gaussian in the frequency domain

$$H(u,v) = e^{-\frac{D^2(u,v)}{2D_o^2}}$$

$$D(u,v) = \sqrt{(u - P/2)^2 + (v - Q/2)^2}$$

- is also a Gaussian in the spatial domain

$$h(x,y) = 2\pi D_o^2 e^{-2\pi^2 D_o^2 (x^2+y^2)}$$

- Notice the reciprocal behavior in their widths

GLPF properties

- Since the spatial domain representation is also a Gaussian, there is no ringing artifact.
- The cutoff of this filter is not sharp at all. But results are comparable to BLPF with n=2. Good for cases where ringing is unacceptable.

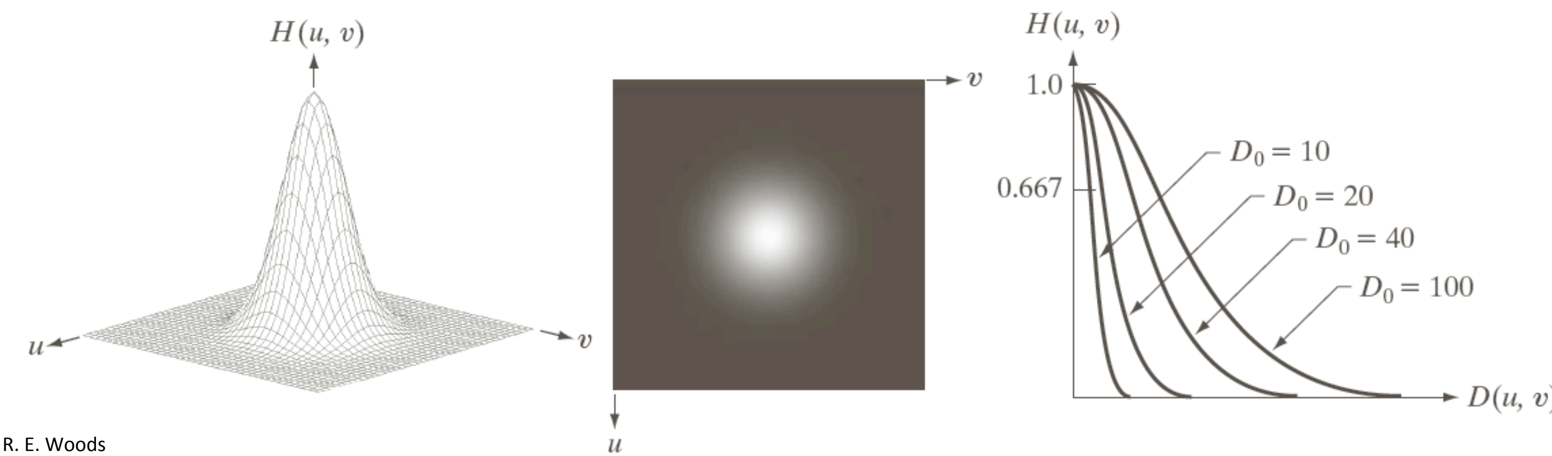

© 1992–2008 R. C. Gonzalez & R. E. Woods

High Pass Filtering

- HP = 1 - LP
  - All the same filters as HP apply
- Applications
  - Visualization of high-freq data (accentuate)
- High boost filtering
  - HB = (1- a) + a(1 - LP) = 1 - a*LP

High-Pass Filters

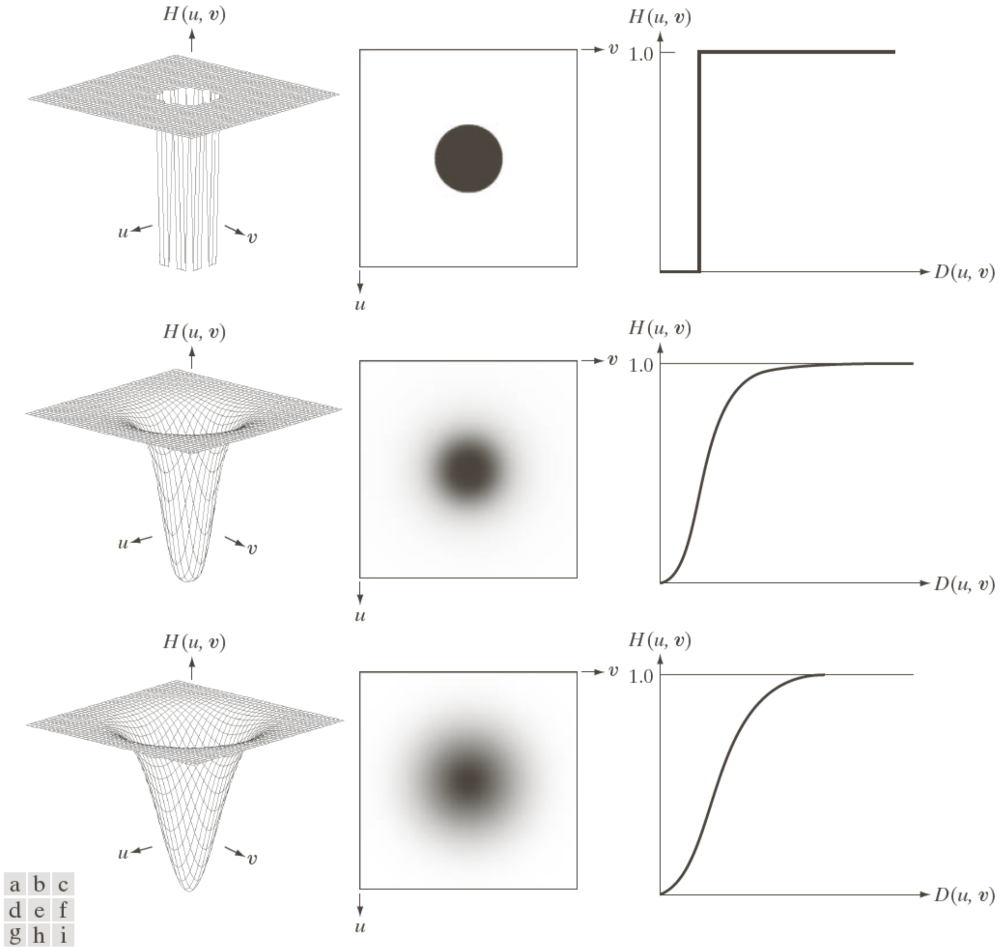

FIGURE 4.52 Top row: Perspective plot, image representation, and cross section of a typical ideal highpass filter. Middle and bottom rows: The same sequence for typical Butterworth and Gaussian highpass filters.

High-Pass Filters in Spatial Domain

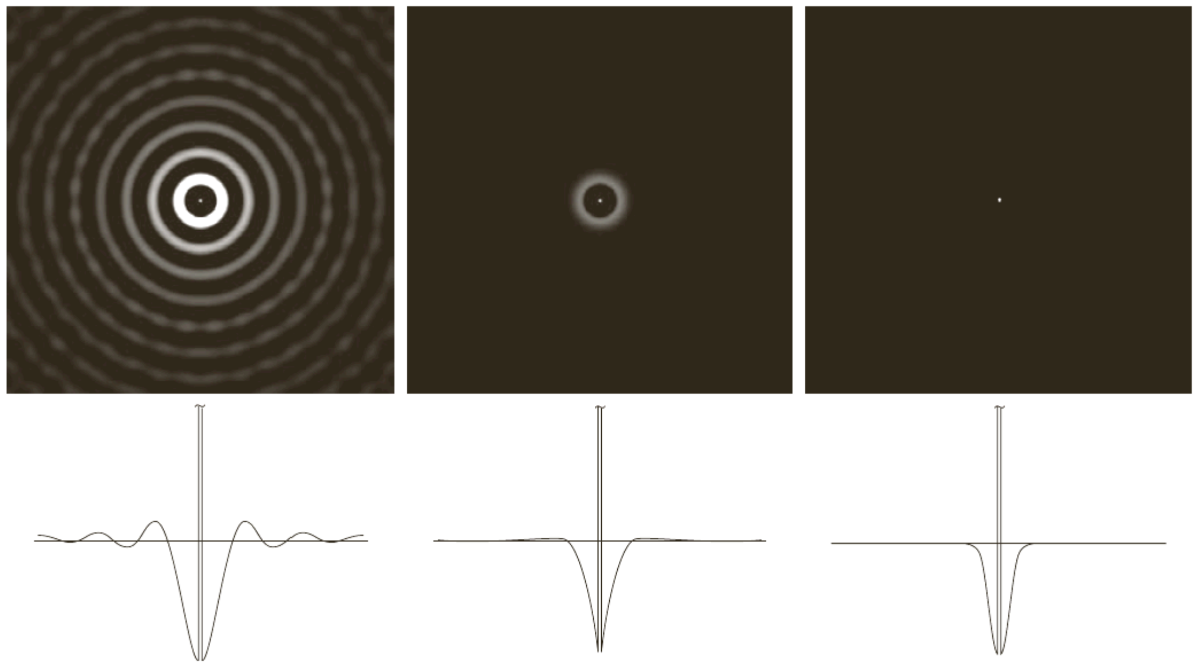

a b c

FIGURE 4.53 Spatial representation of typical (a) ideal, (b) Butterworth, and (c) Gaussian frequency domain highpass filters, and corresponding intensity profiles through their centers.

High-Pass Filtering with IHPF

a b c

FIGURE 4.54 Results of highpass filtering the image in Fig. 4.41(a) using an IHPF with $D_0 = 30, 60,$ and 160.

BHPF

a b c

FIGURE 4.55 Results of highpass filtering the image in Fig. 4.41(a) using a BHPF of order 2 with $D_0 = 30, 60,$ and 160, corresponding to the circles in Fig. 4.41(b). These results are much smoother than those obtained with an IHPF.

GHPF

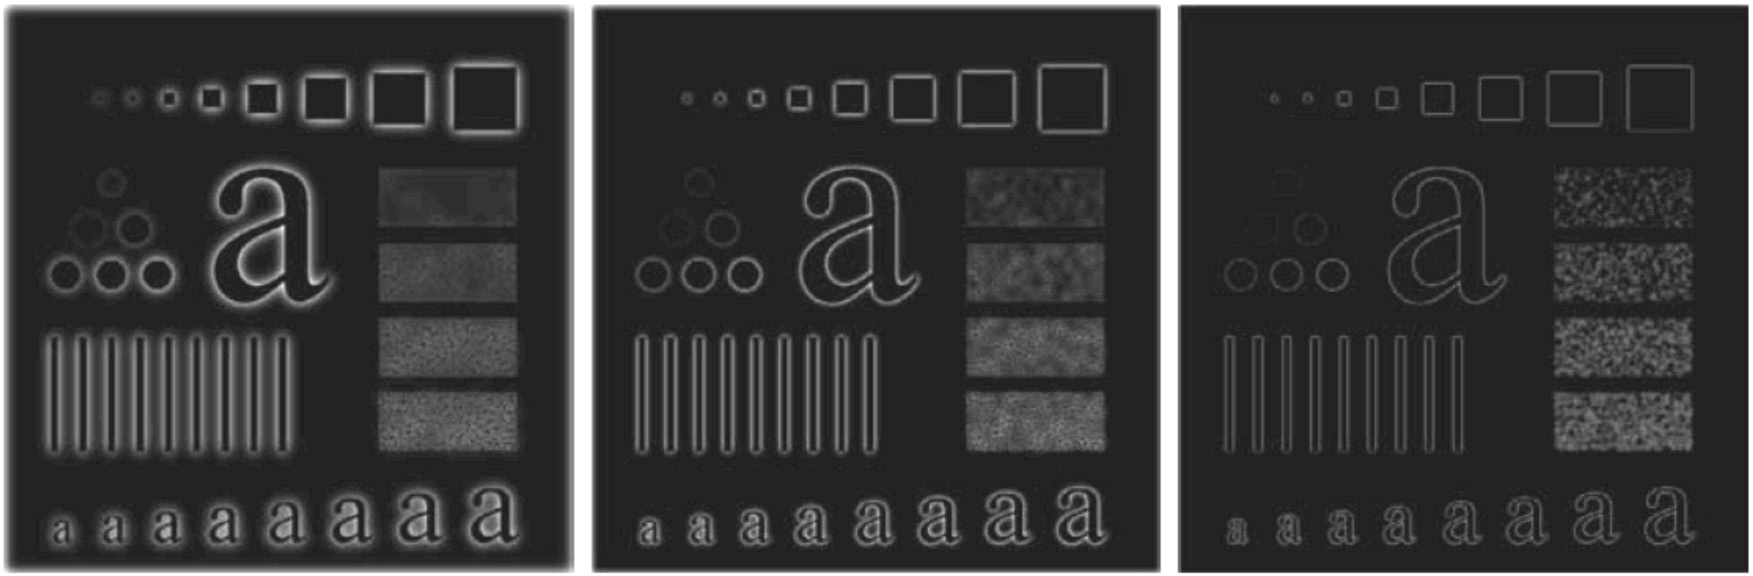

a b c

FIGURE 4.56 Results of highpass filtering the image in Fig. 4.41(a) using a GHPF with $D_0 = 30, 60,$ and 160, corresponding to the circles in Fig. 4.41(b). Compare with Figs. 4.54 and 4.55.

Highpass filter application

a b c

FIGURE 4.57 (a) Thumb print. (b) Result of highpass filtering (a). (c) Result of thresholding (b). (Original image courtesy of the U.S. National Institute of Standards and Technology.)

Butterworth highpass filter removes slow variations in intensity across the image. This in turn allows a single global threshold to separate the ridges and valleys of the fingerprint.

© 1992–2008 R. C. Gonzalez & R. E. Woods

HP, HB, HE

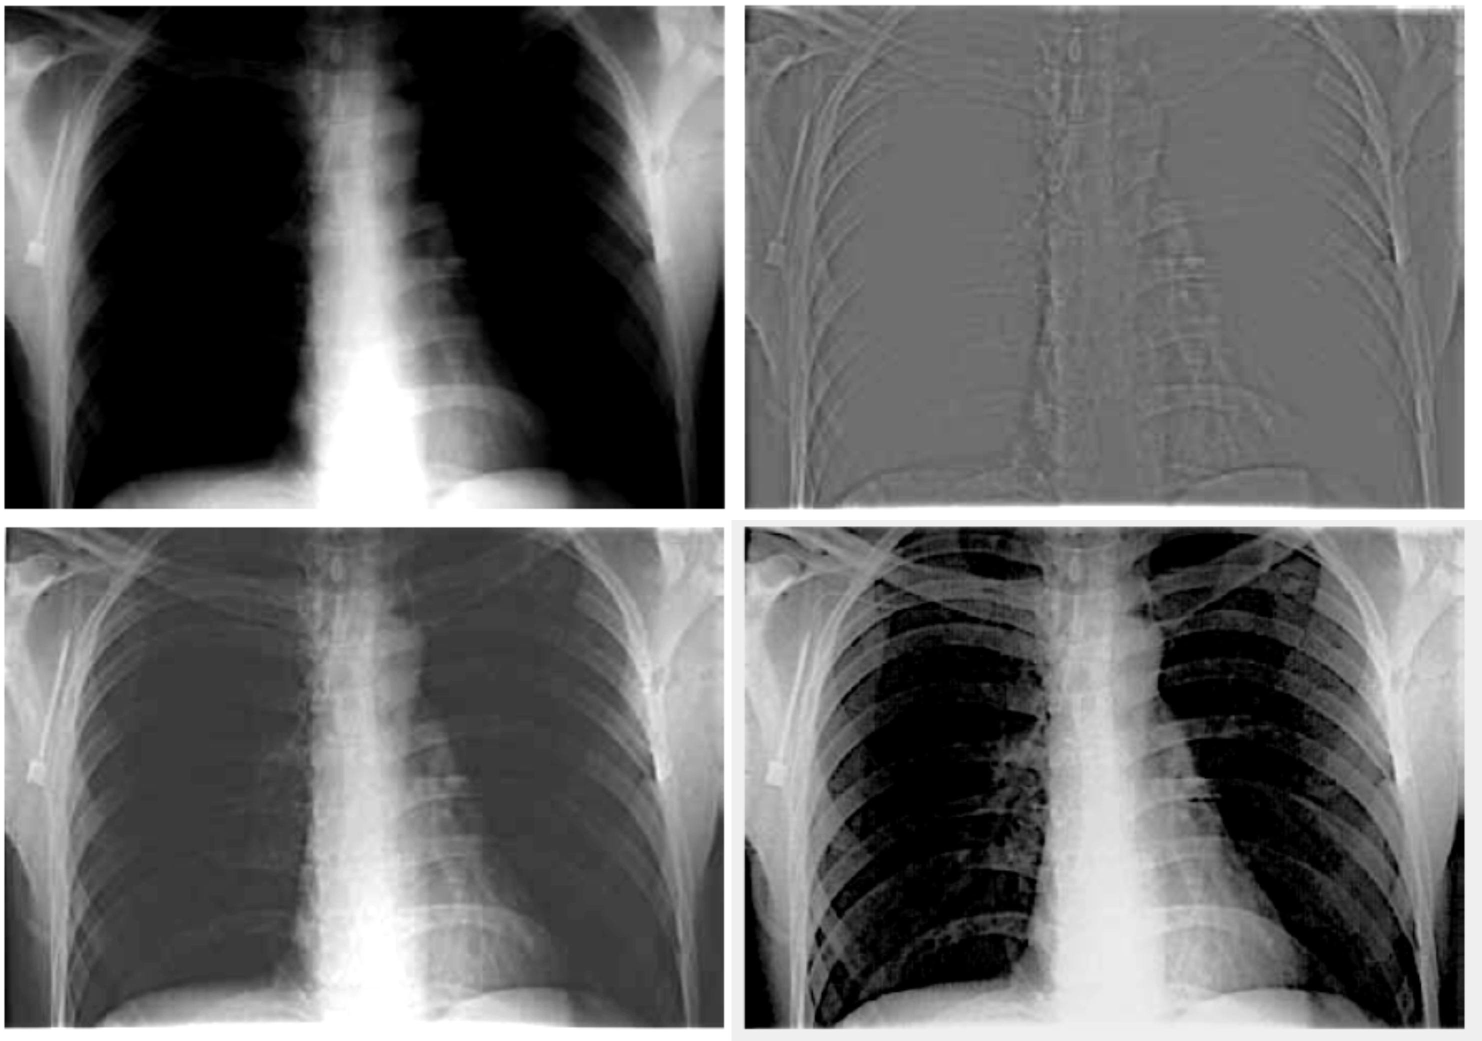

High-Boost Filtering

Band-Pass Filters

- Shift LP filter in Fourier domain by convolution with delta

$$\delta(s - s_0) + \delta(s + s_0)$$

LP

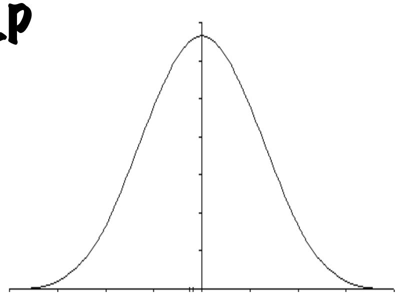

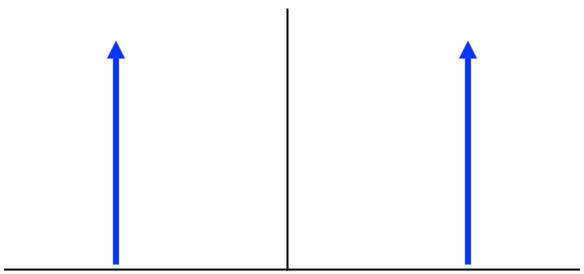

BP

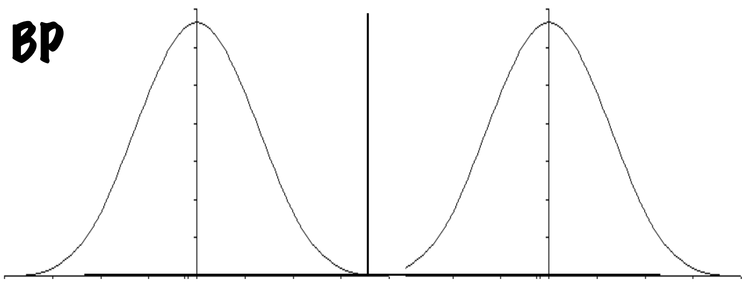

Typically 2-3 parameters
-Width
-Slope
-Band value

Band Pass - Two Dimensions

- Two strategies
  – Rotate
    - Radially symmetric
  – Translate in 2D
    - Oriented filters

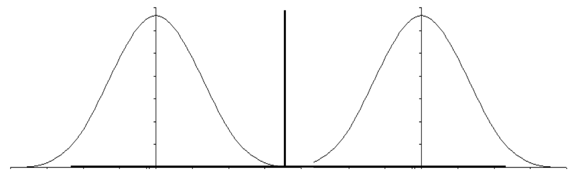

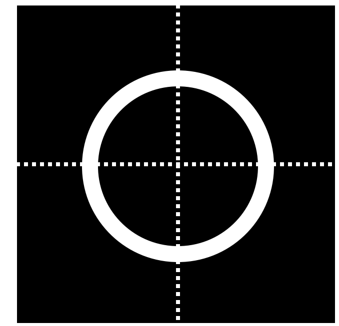

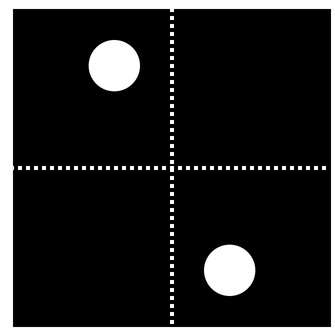

- Note:
  – Convolution with delta-pair in FD is multiplication with cosine in spatial domain

Band Bass Filtering

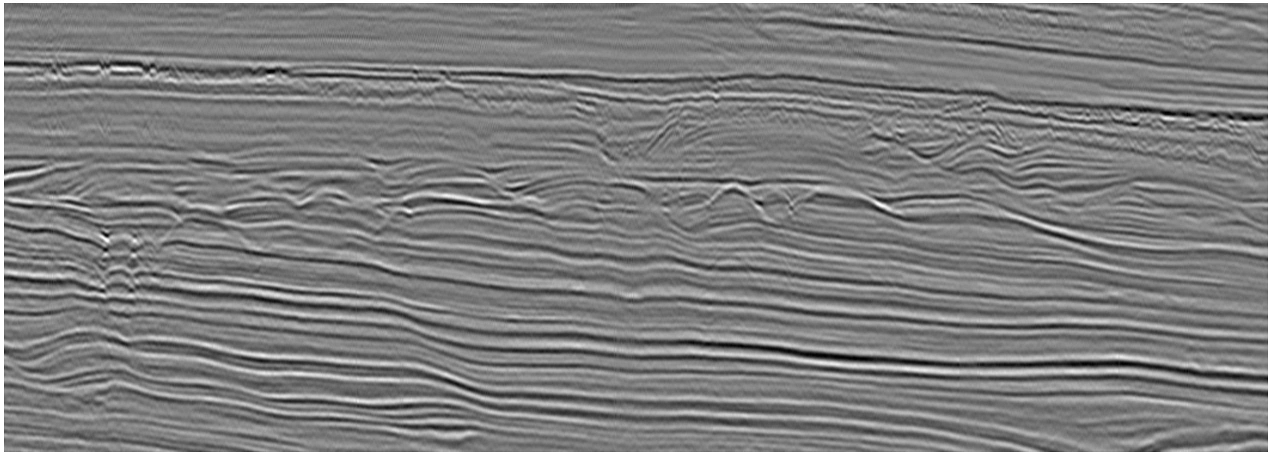

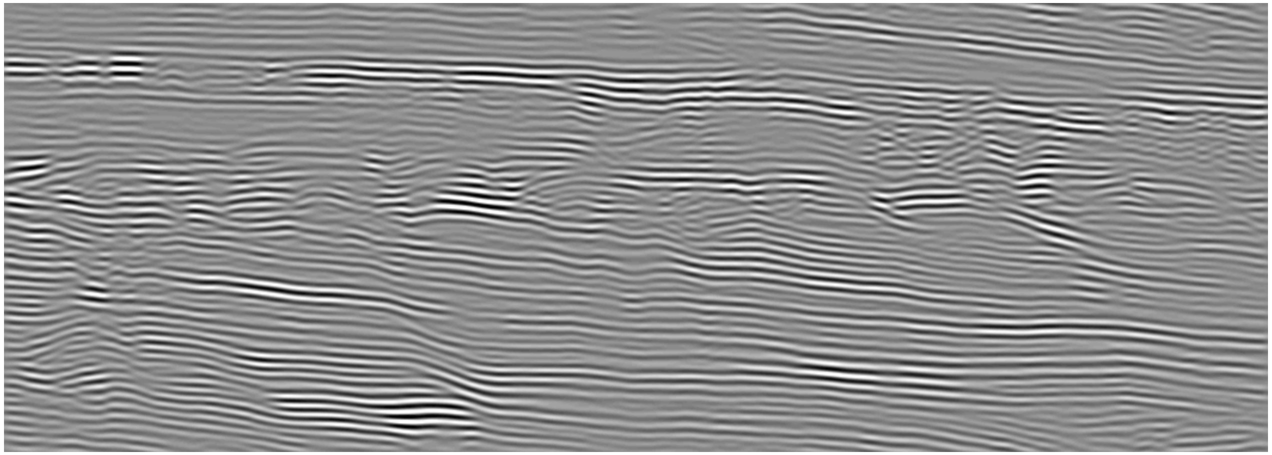

Bandreject filters

TABLE 4.6
Bandreject filters. W is the width of the band, D is the distance $D(u, v)$ from the center of the filter, D_0 is the cutoff frequency, and n is the order of the Butterworth filter. We show D instead of $D(u, v)$ to simplify the notation in the table.

Ideal	Butterworth	Gaussian
$H(u,v) = \begin{cases} 0 & \text{if } D_0 - \frac{W}{2} \leq D \leq D_0 + \frac{W}{2} \\ 1 & \text{otherwise} \end{cases}$	$H(u,v) = \dfrac{1}{1 + \left[\dfrac{DW}{D^2 - D_0^2}\right]^{2n}}$	$H(u,v) = 1 - e^{-\left[\frac{D^2 - D_0^2}{DW}\right]^2}$

- D is the distance from the center of frequency domain
- D_o is the radial center of the band to be rejected
- W is the width of the band to be rejected

© 1992–2008 R. C. Gonzalez & R. E. Woods

Bandpass – bandreject relationship

- We can obtain a bandpass filter from a bandreject filter: $H_{BP}(u,v) = 1 - H_{BR}(u,v)$

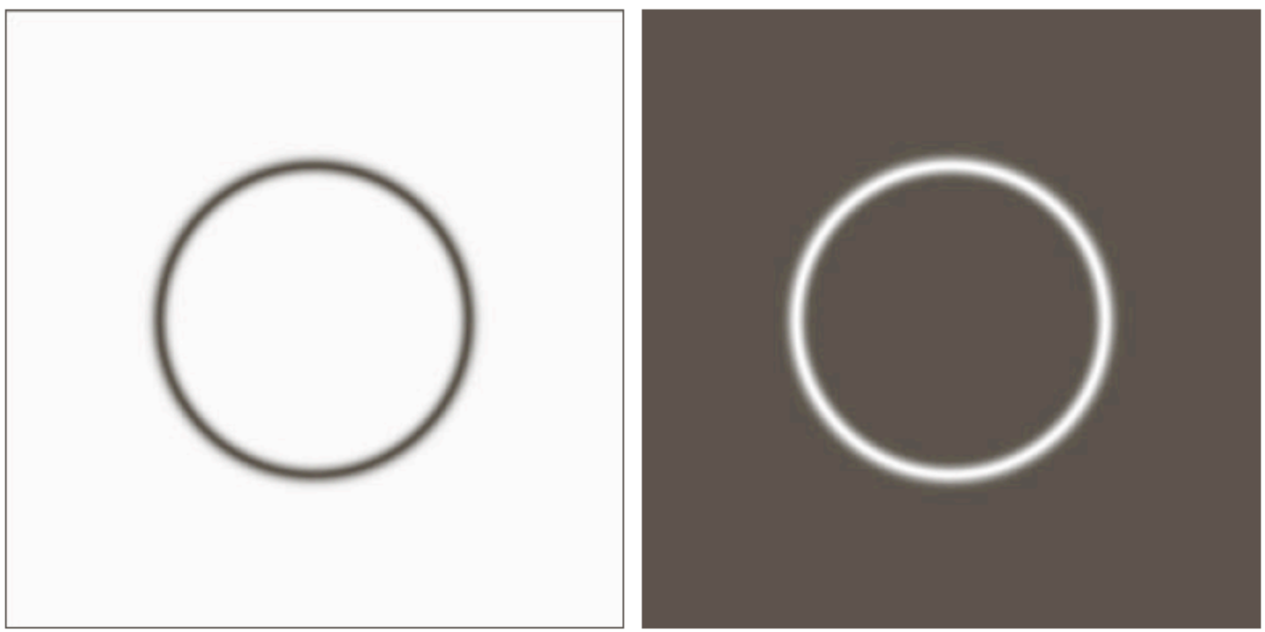

© 1992–2008 R. C. Gonzalez & R. E. Woods

Notch reject filters

- We can reject frequencies more selectively than bandreject filters

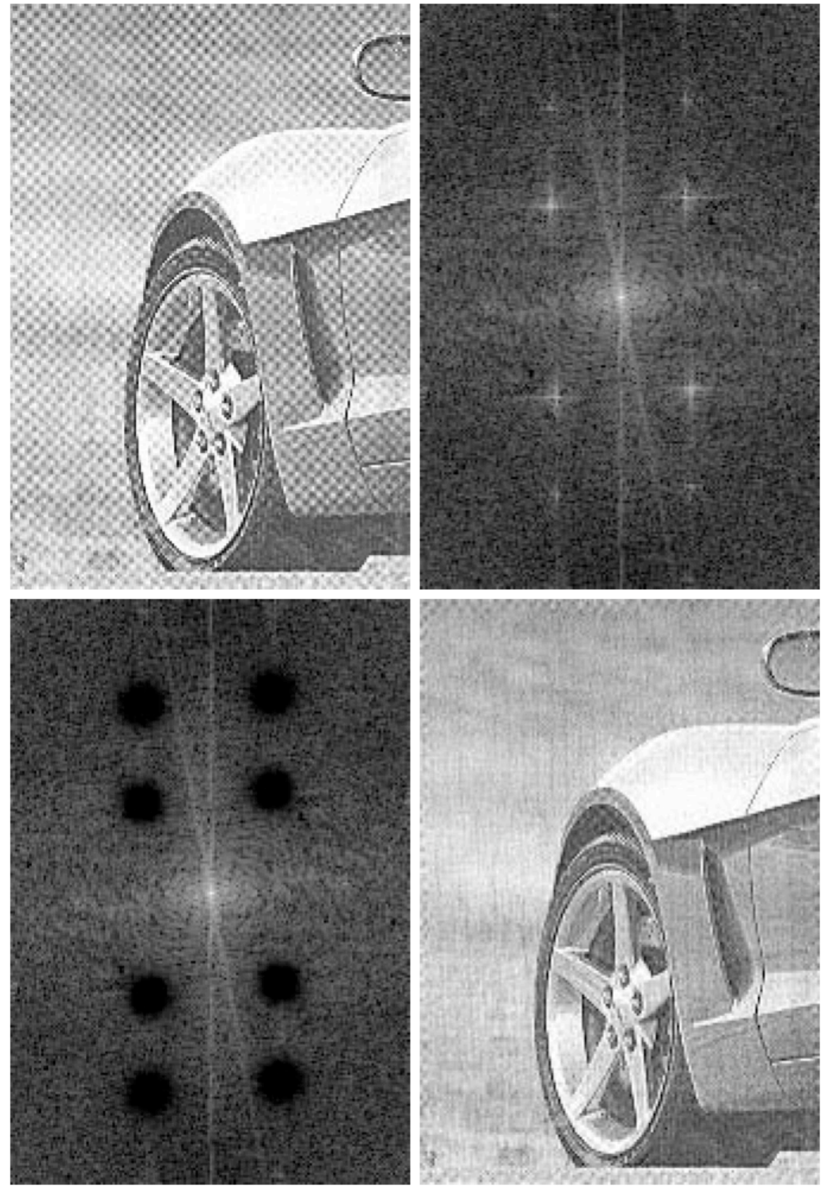

a b
c d

FIGURE 4.64 (a) Sampled newspaper image showing a moiré pattern. (b) Spectrum. (c) Butterworth notch reject filter multiplied by the Fourier transform. (d) Filtered image.

© 1992–2008 R. C. Gonzalez & R. E. Woods

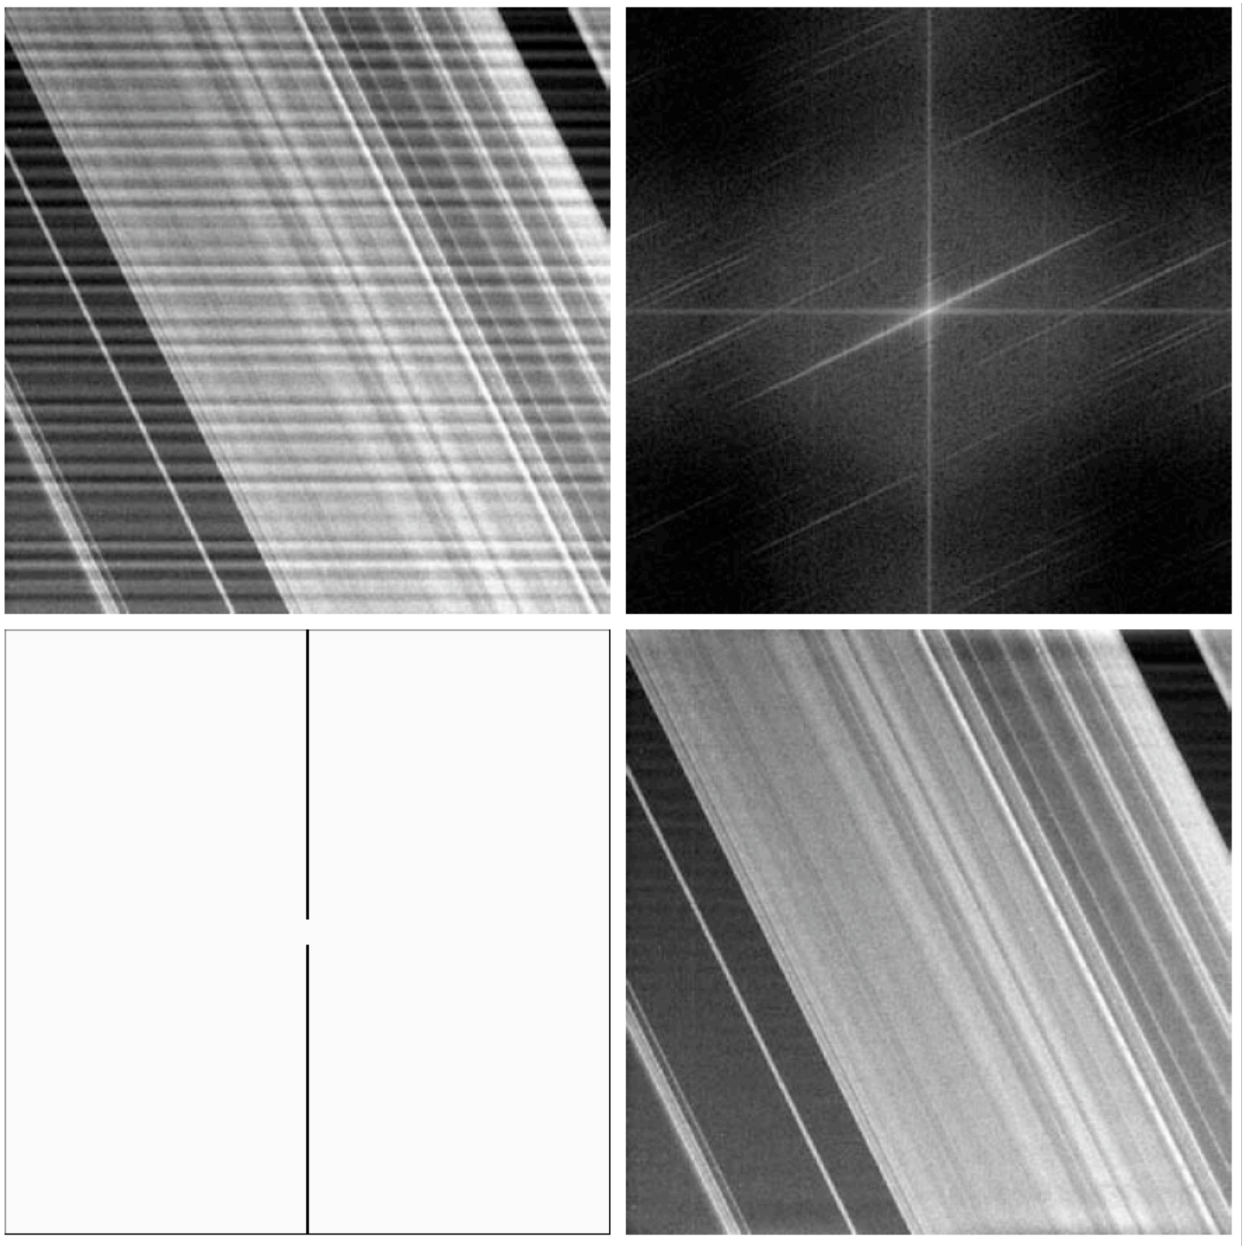

© 1992–2008 R. C. Gonzalez & R. E. Woods

What was rejected?

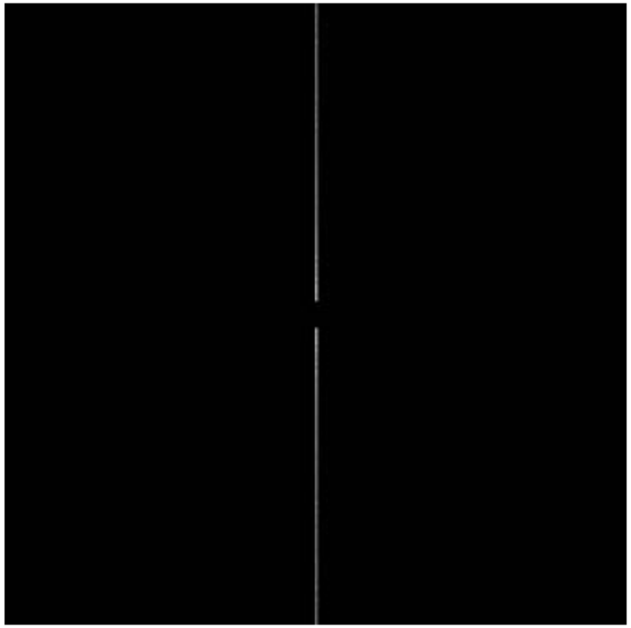

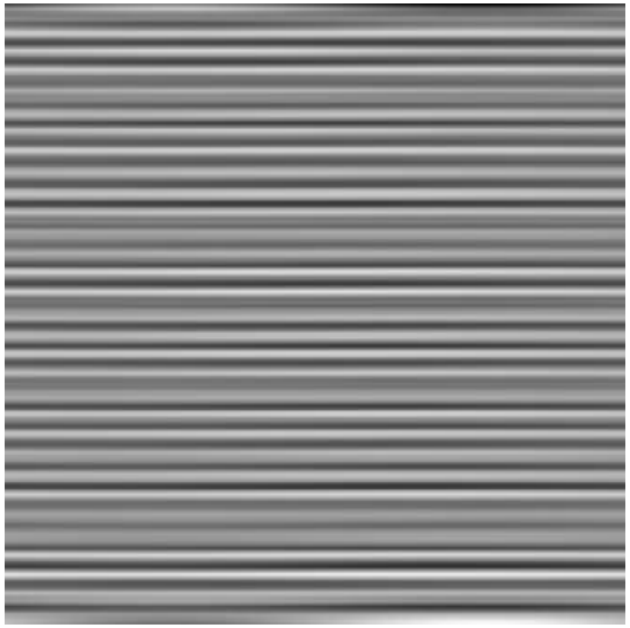

© 1992–2008 R. C. Gonzalez & R. E. Woods

Image Degradation/Restoration

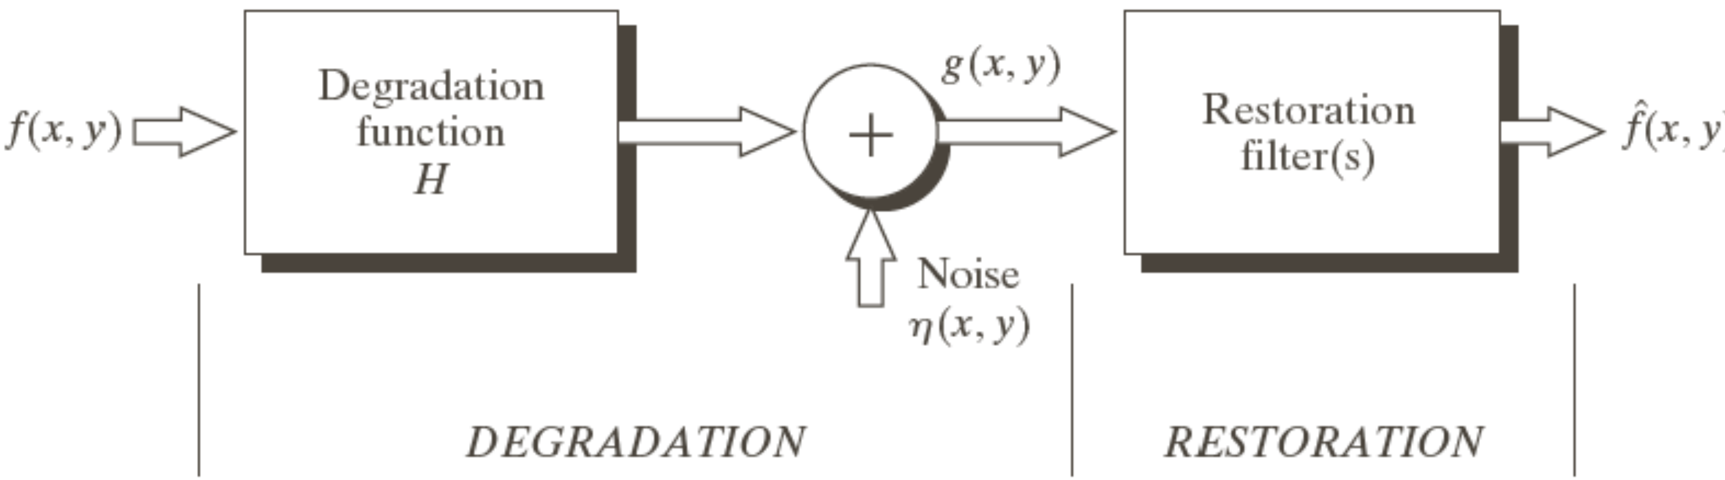

© 1992–2008 R. C. Gonzalez & R. E. Woods

Linear, space invariant degradations

- General model
  - g(x,y)=H[f(x,y)] + n(x,y)
- Linearity:
  - H[a f_1(x,y) + b f_2(x,y)] = a H[f_1(x,y)] + b H[f_2(x,y)]
- Position invariance:
  - H [f (x-a,y-b)] = g (x-a,y-b)
- Linear, position invariant model

$$g(x,y) = h(x,y) * f(x,y) + \eta(x,y)$$

$$G(u,v) = H(u,v)F(u,v) + N(u,v)$$

Point spread function

- A large class of degradations can be approximated as linear, space invariant processes
  - Example: most physical optical systems blur (spread) a point of light to some degree
  - Impulse response $h(x,y) = H[\ \delta(x,y)\ ]$ is sometimes called the point spread function (PSF)

a b

FIGURE 5.24 Degradation estimation by impulse characterization. (a) An impulse of light (shown magnified). (b) Imaged (degraded) impulse.

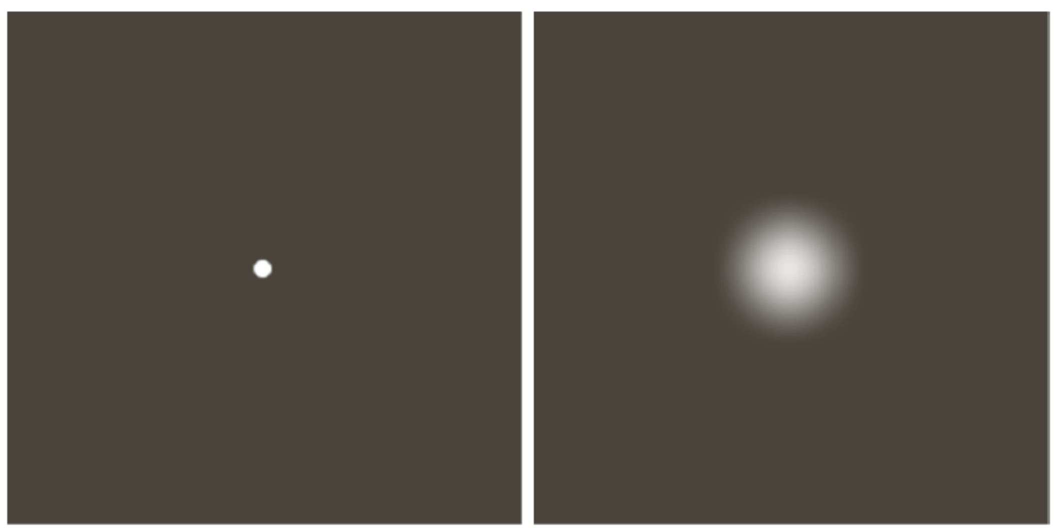

© 1992–2008 R. C. Gonzalez & R. E. Woods

Inverse filtering

$$G(u,v) = H(u,v)F(u,v) + N(u,v)$$

- The main idea is

$$\hat{F}(u,v) = \frac{G(u,v)}{H(u,v)} = F(u,v) + \frac{N(u,v)}{H(u,v)}$$

- This assumes we know H(u,v)…

Estimation by observation (for blurring H)

- Choose a small area of the image with strong signal content: $g_s(x,y)$
  - High contrast area
  - Edge
- Create an unblurred version of the area
  - Sharpening
  - Processing by hand

$$\hat{f}_s(x,y)$$

- Then, we can estimate (assuming negligible noise)

$$H_s(u,v) = \frac{G_s(u,v)}{\hat{F}_s(u,v)}$$

Estimation by Experimentation

- If you have access to the instrument used to acquire the image to be restored
- Capture the image of a small dot of light
  - Call the capture image g(x,y)
  - A dot of light is our approximation to the *Dirac delta* function: $f(x,y) = A\delta(x,y)$
  - $F(u,v) = F[\,A\delta(x,y)\,] = A$
  - A is the amplitude of the dot of light

$$H(u,v) = \frac{G(u,v)}{A}$$

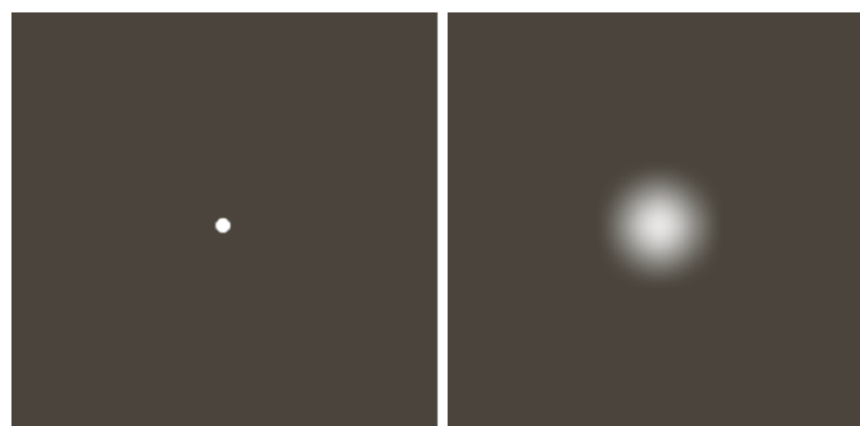

© 1992–2008 R. C. Gonzalez & R. E. Woods

Estimation by Modeling

- Sometimes it might be possible to mathematically derive a model for the degradation
- Example motion blur. Due to
  - Camera motion. Camera on a moving platform.
  - Objects moving

$$g(x,y) = \int_0^T f\left(x - x_o(t), y - y_o(t)\right) dt$$

$$g(x,y)=\int_0^T f\left(x-x_o(t),y-y_o(t)\right)dt$$

$$G(u,v)=\iint\left(\int_0^T f\left(x-x_o(t),y-y_o(t)\right)dt\right)e^{-j2\pi(ux+vy)}dxdy$$

$$G(u,v)=\int_0^T\left(\iint f\left(x-x_o(t),y-y_o(t)\right)e^{-j2\pi(ux+vy)}dxdy\right)dt$$

$$G(u,v)=\int_0^T F(u,v)e^{-j2\pi\left(ux_o(t)+vy_o(t)\right)}dt$$

$$G(u,v)=F(u,v)\boxed{\int_0^T e^{-j2\pi\left(ux_o(t)+vy_o(t)\right)}dt}$$

Motion blur
H(u,v)

Motion blur example

- If the motion variables are known then we can find H(u,v)
- Example: Linear motion
  - $x_o(t) = at/T,\ y_o(t)=bt/T$

$$H(u,v) = \frac{T}{\pi(ua+vb)} \sin\left[\pi(ua+vb)\right] e^{-j\pi(ua+vb)}$$

a b

FIGURE 5.26 (a) Original image. (b) Result of blurring using the function in Eq. (5.6-11) with $a = b = 0.1$ and $T = 1$.

© 1992–2008 R. C. Gonzalez & R. E. Woods

Inverse filtering

$$G(u,v) = H(u,v)F(u,v) + N(u,v)$$

- The main idea is

$$\hat{F}(u,v) = \frac{G(u,v)}{H(u,v)} = F(u,v) + \frac{N(u,v)}{H(u,v)}$$

- We now have an estimate of H(u,v)
- But there is still a problem

The problem with inverse filtering

$$\hat{F}(u,v) = \frac{G(u,v)}{H(u,v)} = F(u,v) + \frac{N(u,v)}{H(u,v)}$$

- Even if we know H(u,v) we can’t recover f(x,y) because N(u,v) is not known
- Also if H(u,v) has very small values for some (u,v), then the N(u,v)/H(u,v) part can easily dominate for those (u,v)

Inverse filtering example

Motion blur + additive noise

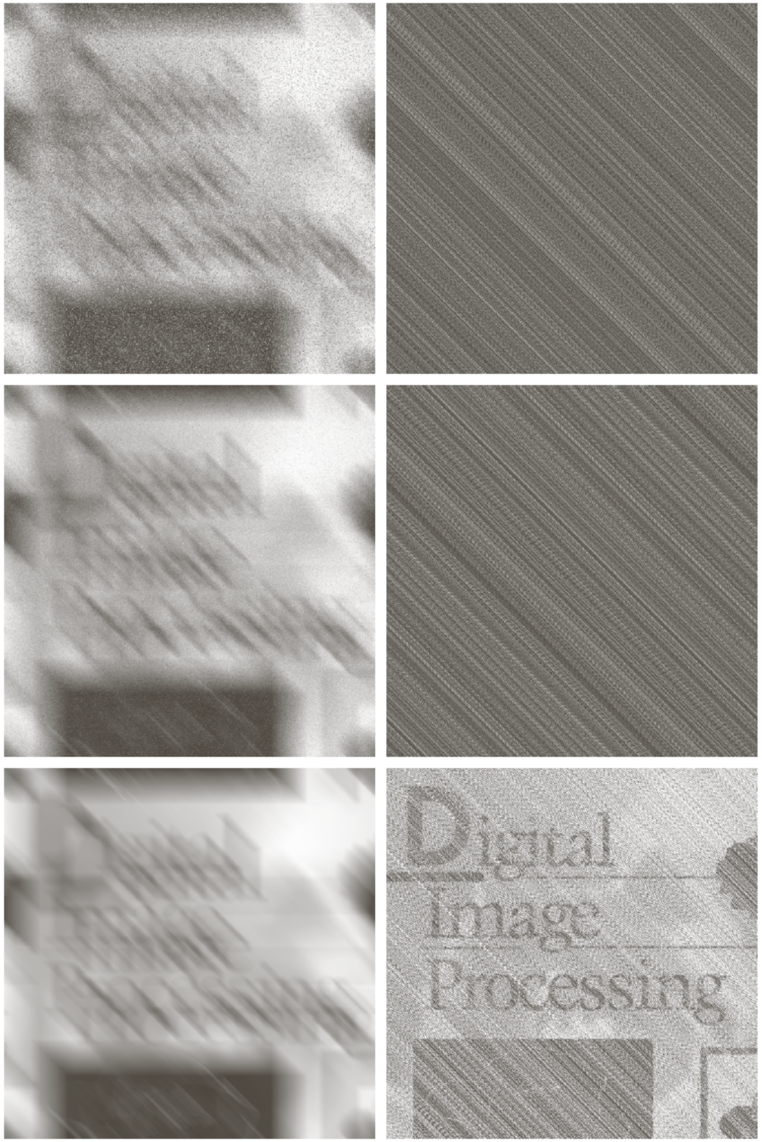

Smaller noise variance

Degraded | Inverse filtering

© 1992–2008 R. C. Gonzalez & R. E. Woods

Optimal/Weiner Filter

- Power spectrum of signal, noise are known
- H(u) is known
- The expected squared error of restoration with filter w

$$E\left[\left\|f - \hat{f}\right\|^2\right]$$

$$\hat{f} = w * g$$

$$g(x,y) = h(x,y) * f(x,y) + \eta(x,y)$$

Optimal/Weiner Filter

- Power spectrum of signal, noise are known
- H(u) is known
- The expected squared error of restoration with filter w, formulated in Fourier domain

$$E\left[\left\|F-\hat{F}\right\|^{2}\right]$$

$$\hat{F}(u,v)=W(u,v)G(u,v)$$

$$G(u,v)=H(u,v)F(u,v)+N(u,v)$$

Optimal/Weiner Filter

- Power spectrum of signal, noise are known
- H(u) is known
- Filter that minimizes the expected squared error of reconstruction is:

$$G(u) = \frac{H^*(u)S(u)}{|H(u)|^2 S(u) + N(u)}$$

Optimal/Weiner Filter

$$G(u) = \frac{H^*(u)S(u)}{|H(u)|^2 S(u) + N(u)}$$

$$G(u) = \frac{1}{H(u)} \frac{|H(u)|^2}{|H(u)|^2 + \frac{N(u)}{S(u)}}$$

$$G(u) = \frac{1}{H(u)} \frac{|H(u)|^2}{|H(u)|^2 + \frac{1}{SNR(u)}}$$

Wiener filtering example

Motion blur + additive noise

Smaller noise variance

Motion blur | Inverse filtering | Wiener filtering

© 1992–2008 R. C. Gonzalez & R. E. Woods

Another use of the frequency domain: Image Registration

- Find dx and dy that best matches two images
- Cross correlation can give the best translation between two images

Normalized Cross Correlation

- Subtract the mean of the image and divide by the S.D.
  - This maps the image to the unit sphere
  - A single integral is the dot product of these to vectors
    - angles between the two normalized images
  - Helps alleviate intensity differences

Phase Correlation

$$\mathbf{G}_a = \mathcal{F}\{g_a\},\ \mathbf{G}_b = \mathcal{F}\{g_b\}$$

$$R = \frac{\mathbf{G}_a \mathbf{G}_b^*}{|\mathbf{G}_a \mathbf{G}_b^*|}$$

$$r = \mathcal{F}^{-1}\{R\}$$

$$(\Delta x, \Delta y) = \arg\max_{(x,y)}\{r\}$$

$$g_b(x,y) \stackrel{\text{def}}{=} g_a((x-\Delta x) \bmod M, (y-\Delta y) \bmod N)$$

$$\mathbf{G}_b(u,v) = \mathbf{G}_a(u,v) e^{-2\pi i(\frac{u\Delta x}{M} + \frac{v\Delta y}{N})}$$

$$\begin{aligned} R(u,v) &= \frac{\mathbf{G}_a \mathbf{G}_b^*}{|\mathbf{G}_a \mathbf{G}_b^*|} \\ &= \frac{\mathbf{G}_a \mathbf{G}_a^* e^{2\pi i(\frac{u\Delta x}{M} + \frac{v\Delta y}{N})}}{|\mathbf{G}_a \mathbf{G}_a^* e^{2\pi i(\frac{u\Delta x}{M} + \frac{v\Delta y}{N})}|} \\ &= \frac{\mathbf{G}_a \mathbf{G}_a^* e^{2\pi i(\frac{u\Delta x}{M} + \frac{v\Delta y}{N})}}{|\mathbf{G}_a \mathbf{G}_a^*|} \\ &= e^{2\pi i(\frac{u\Delta x}{M} + \frac{v\Delta y}{N})} \end{aligned}$$

Phase Correlation

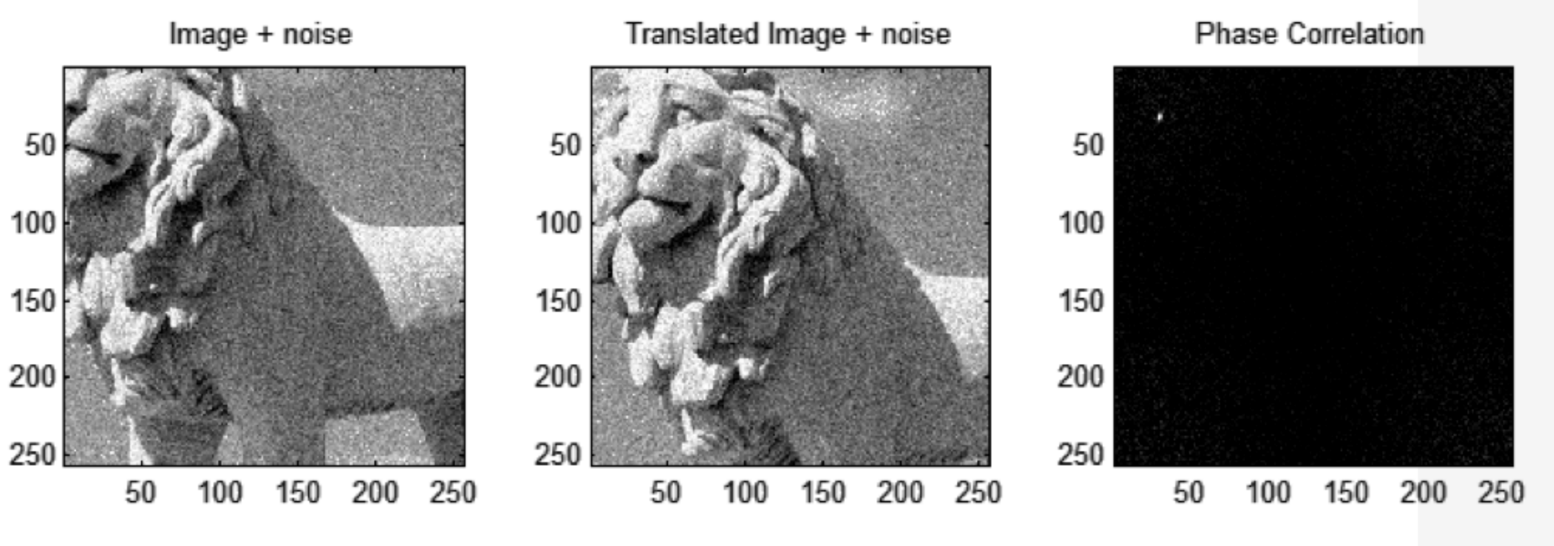

Connectomics

- Mapping connectivity patterns of neurons
  - Sparse/Dense neural circuit reconstruction
- *C. Elegans* complete neural circuit [1975]
  - Electron micrographs
  - 302 neurons, 10 years manual effort
- Imaging is no longer the bottleneck…
  - Serial section TEM (ssTEM)
    - Image alignment problem
  - Serial block face SEM (SBFSEM)
    - Stable block, no image alignment problem
    - In-plane resolution limited
  - Serial section SEM (ssSEM)
    - Fast, can generate very large datasets
    - Image alignment problem
  - Focused ion beam milling SEM (FIBSEM)
    - No image alignment problem, isotropic resolution
    - Smaller field of view
- Analysis is the bottleneck for reconstruction of larger neural circuits
  - Fly brain
  - Primate cortical column

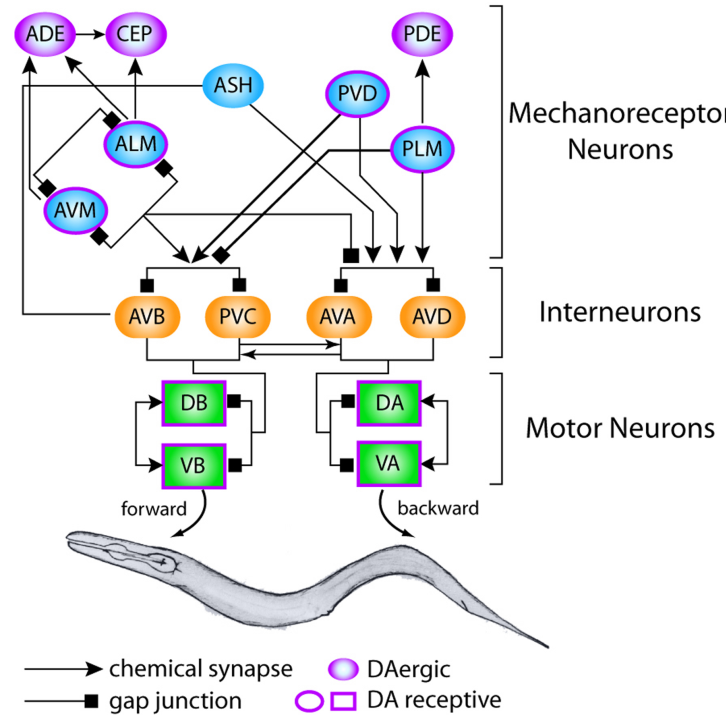

Wormbook, www.wormbook.org

Volume assembly

- Data:
  - 0.2mm x 0.2mm area
  - Each section: ~1000 tiles
  - 341 sections at 90nm thickness
  - Each image 4K x 4K pixels
- Assembling 2D sections
- Section-to-section registration

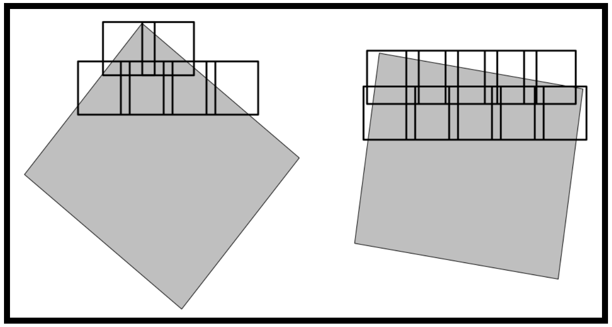

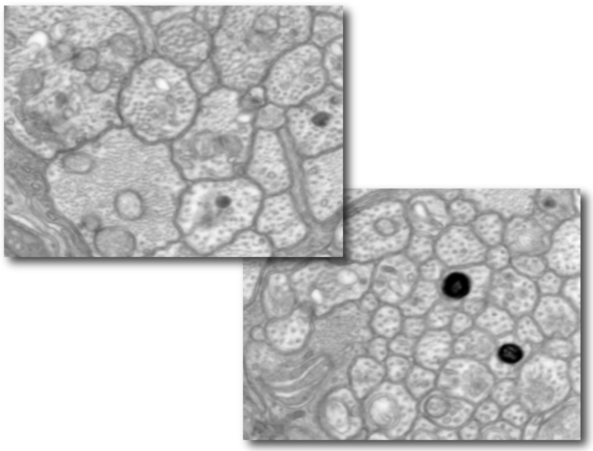

Assembling 2D mosaics

- Find the displacement between tiles
  - Closed-form solution using phase correlation, Girod and Kuo' 89

Fourier shift theorem

$$F\left[I_1\left(x-a\right)\right] = e^{-jwa} F\left[I_1\left(x\right)\right]$$

Normalized cross power spectrum

$$P\left(w\right) = \frac{F_1(w)F_2^*(w)}{\left|F_1(w)F_2^*(w)\right|} = e^{jwa}$$

Inverse Fourier transform of P interpreted as a displacement probability distribution

  - Need at least 10% overlap for reliable displacement estimate

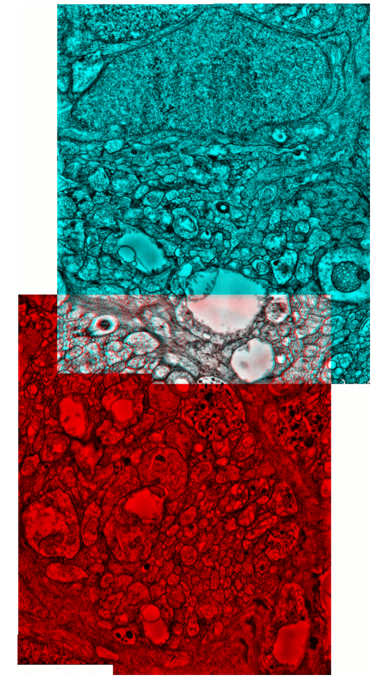

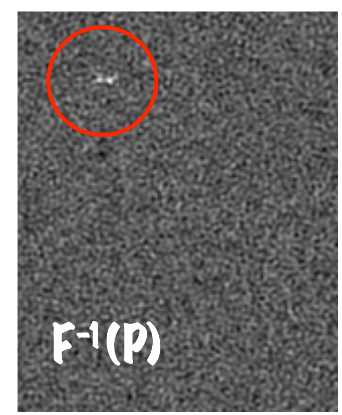

132 tile mosaic