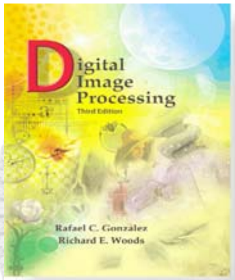

Lecture #10

- Wraparound and padding
- Image Correlation
- Image Processing in the frequency domain
- A simple frequency domain filter
- Frequency domain filters
  - High-pass, low-pass
  - Apodization
- Zero-phase filtering
- Frequency domain filters
  - ideal, Butterworth, Gaussian
  - Ringing

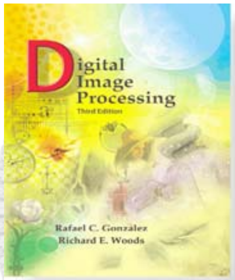

Wraparound and Padding

DFT makes f & h periodic

$$f(x)\star h(x) = \frac{1}{M}\sum_{m=0}^{\infty} f(m)h(x-m)$$

Mirror h to get h(-m)

Now shift h(-m) to get h(x-m)

If f and g are 400 points f★g goes out to 800 points

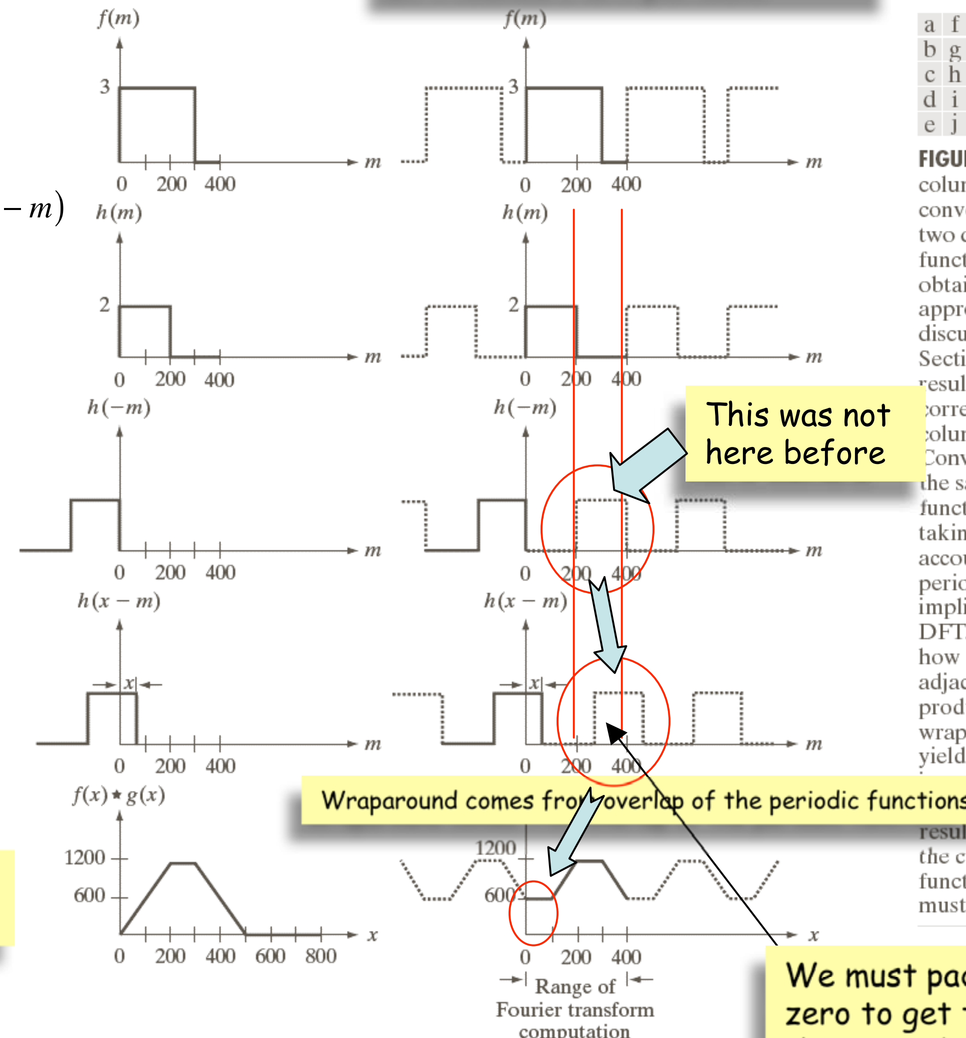

FIGURE 4.28 Left column: convolution of two discrete functions obtained using the approach discussed in Section 3.4.2. The result in (e) is correct. Right column: Convolution of the same functions, but taking into account the periodicity implied by the DFT. Note in (j) how data from adjacent periods produce wraparound error, yielding an … result. To obtain the correct result, function padding must be used.

We must pad h with 399 points of zero to get the correct result with this periodic function

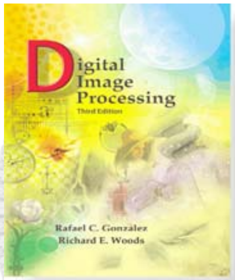

Wraparound and Padding

We also have to pad images but in two-dimensions to get the correct convolution results

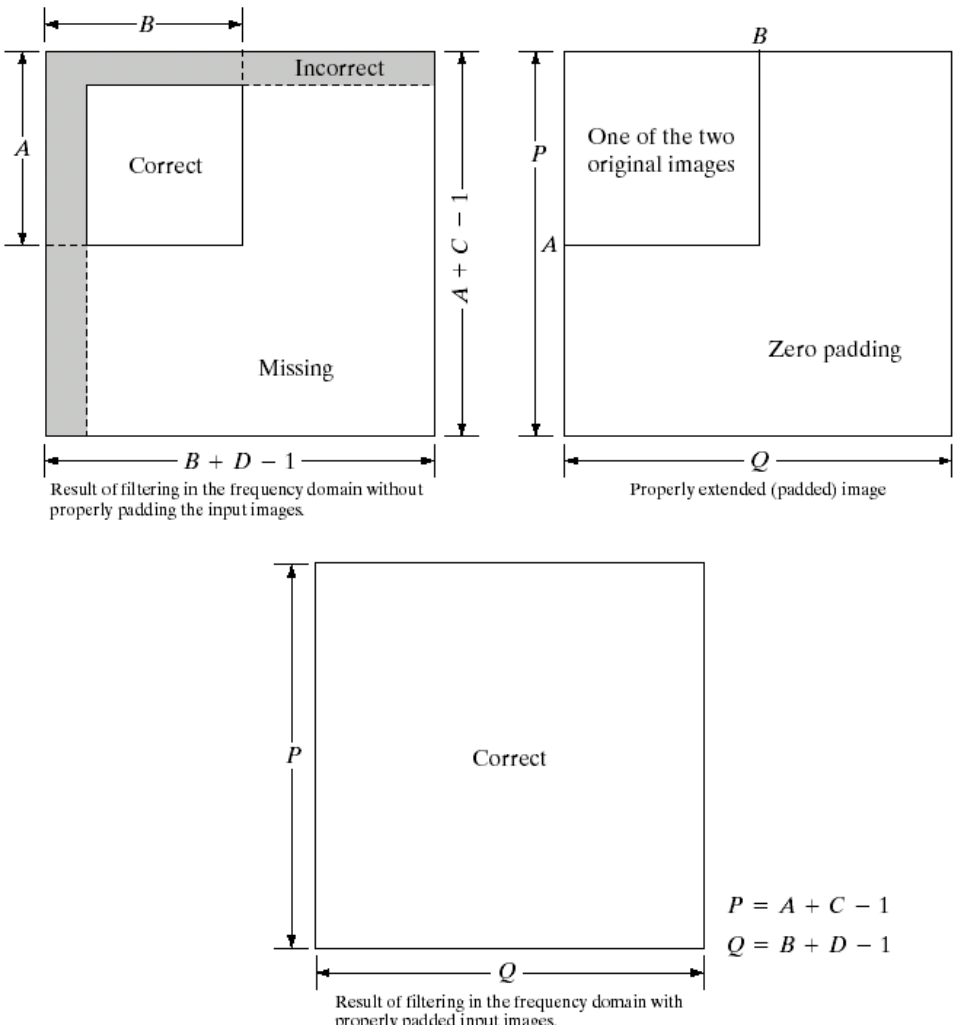

a b
c

FIGURE 4.38 Illustration of the need for function padding. (a) Result of performing 2-D convolution without padding. (b) Proper function padding. (c) Correct convolution result.

$$P = A + C - 1$$
$$Q = B + D - 1$$

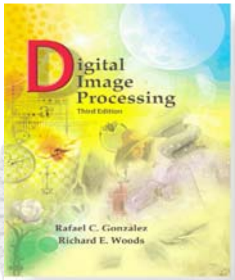

Wraparound and Padding

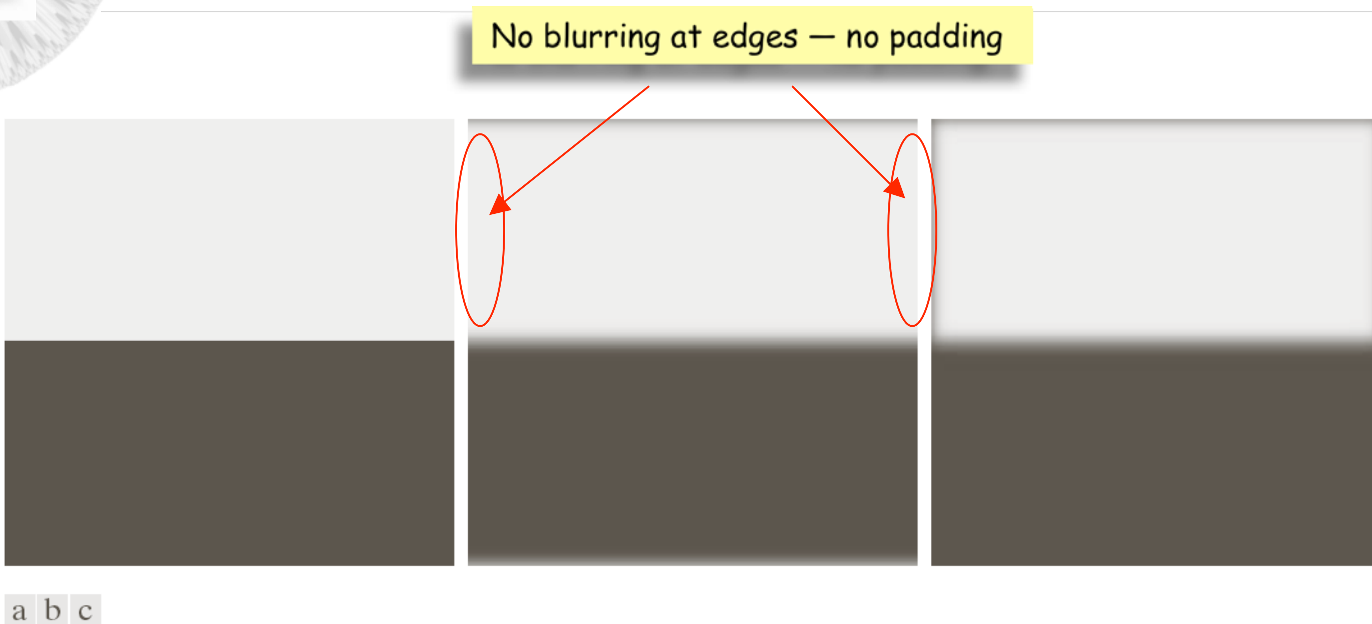

FIGURE 4.32 (a) A simple image. (b) Result of blurring with a Gaussian lowpass filter without padding. (c) Result of lowpass filtering with padding. Compare the light area of the vertical edges in (b) and (c).

$$\frac{1}{16}\begin{bmatrix} 1 & 2 & 1 \\ 2 & 4 & 2 \\ 1 & 2 & 1 \end{bmatrix}$$

Wraparound and Padding

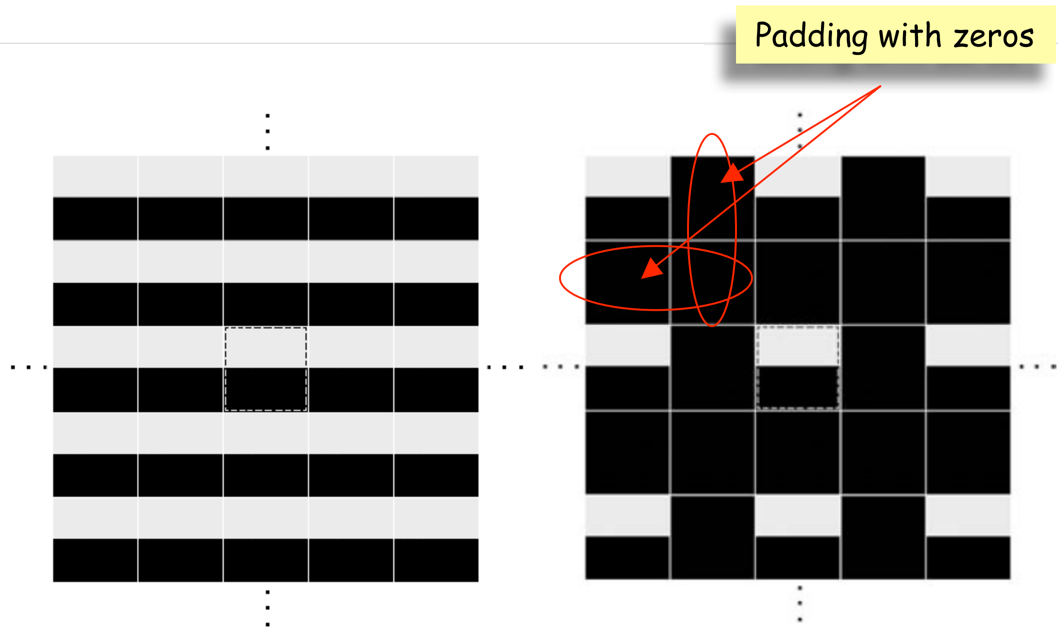

a b

FIGURE 4.33 2-D image periodicity inherent in using the DFT. (a) Periodicity without image padding. (b) Periodicity after padding with 0s (black). The dashed areas in the center correspond to the image in Fig. 4.32(a). (The thin white lines in both images are superimposed for clarity; they are not part of the data.)

Wraparound and Padding

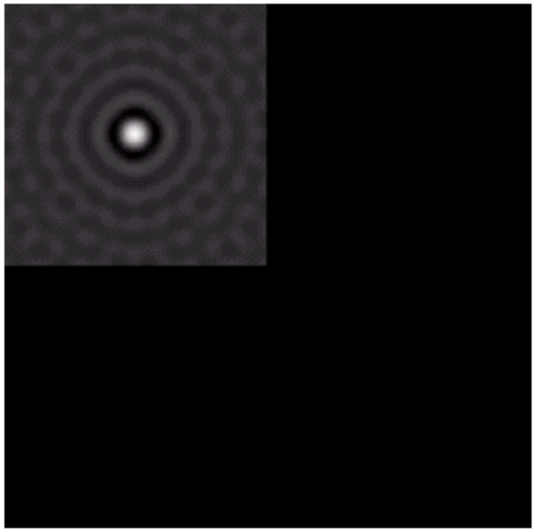

FIGURE 4.39 Padded lowpass filter is the spatial domain (only the real part is shown).

FIGURE 4.40 Result of filtering with padding. The image is usually cropped to its original size since there is little valuable information past the image boundaries.

Image Correlation

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

Original images

Padded images

Correlation output

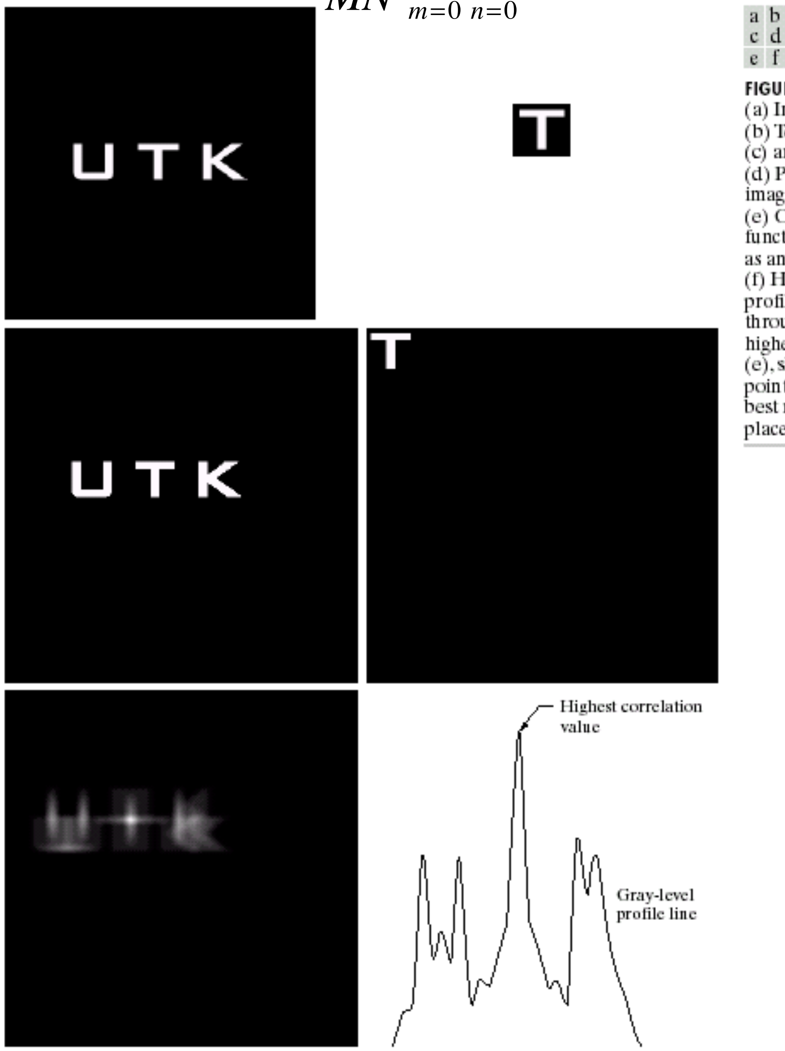

FIGURE 4.41 (a) Image. (b) Template. (c) and (d) Padded images. (e) Correlation function displayed as an image. (f) Horizontal profile line through the highest value in (e), showing the point at which the best match took place.

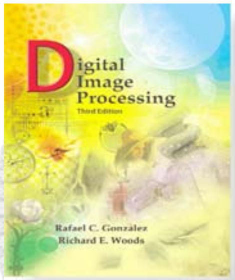

Image Processing in the Frequency Domain

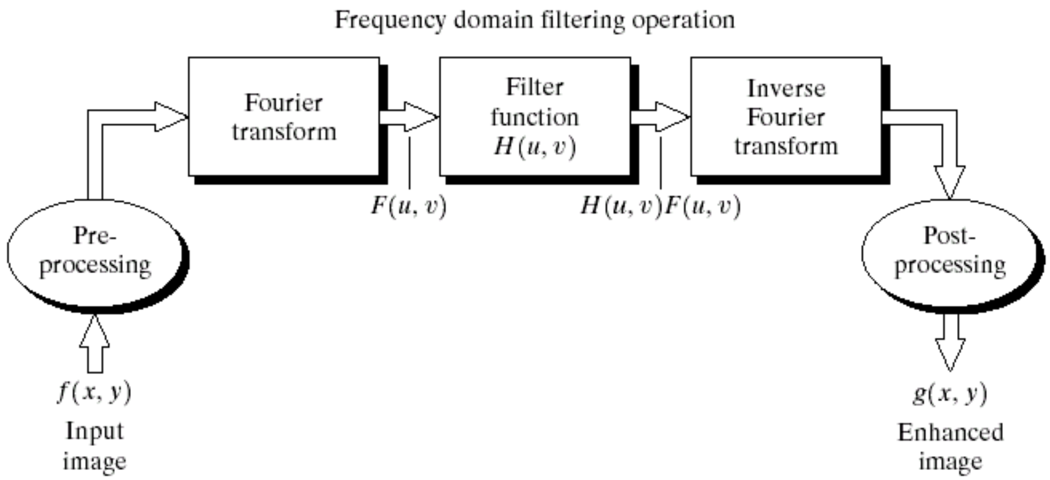

FIGURE 4.5 Basic steps for filtering in the frequency domain.

a. Pad original image to PxQ
b. Multiply padded image by $(-1)^{x+y}$
d. Fourier transform F(u,v)
e. Construct real, symmetric filter H(u,v) of size PxQ
f. Compute F(u,v)H(u,v)
g. Compute $\mathrm{Re}[\mathcal{F}^{-1}\{F(u,v)H(u,v)\}]$ and multiply by $(-1)^{x+y}$
h. Crop to remove padding in result

Fourier Transform

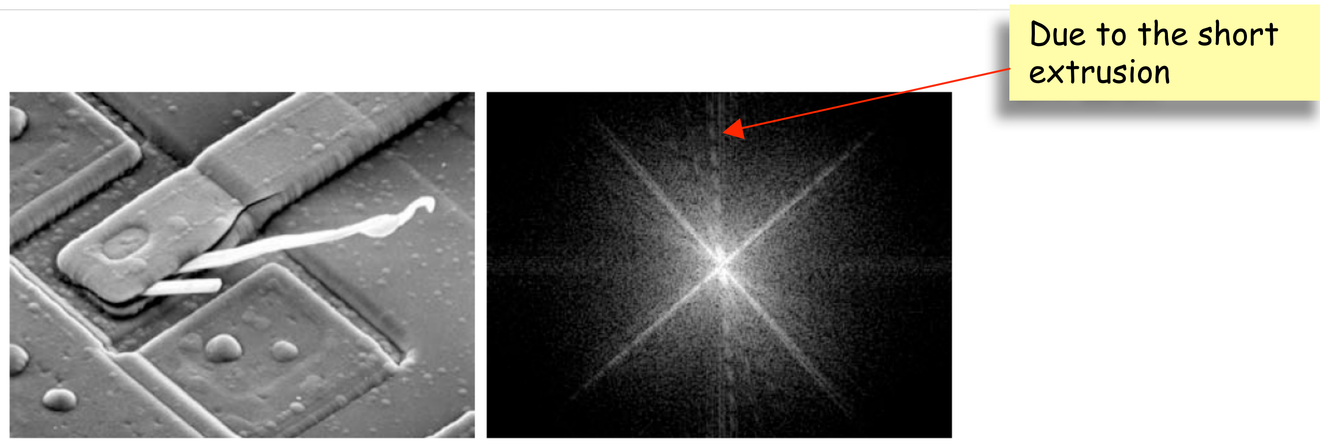

a b

FIGURE 4.29 (a) SEM image of a damaged integrated circuit. (b) Fourier spectrum of (a). (Original image courtesy of Dr. J. M. Hudak, Brockhouse Institute for Materials Research, McMaster University, Hamilton, Ontario, Canada.)

Simple Frequency Domain Filter

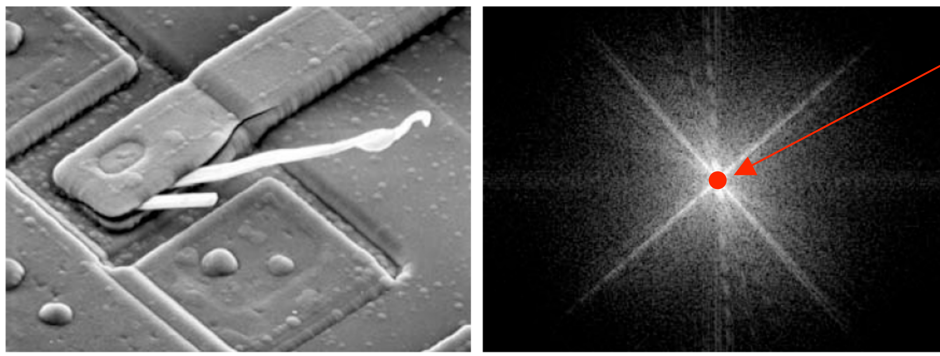

a b

FIGURE 4.29 (a) SEM image of a damaged integrated circuit. (b) Fourier spectrum of (a). (Original image courtesy of Dr. J. M. Hudak, Brockhouse Institute for Materials Research, McMaster University, Hamilton, Ontario, Canada.)

What if we simply multiply F(0,0) by zero?

This is the minimal high-pass filter

Simple Frequency Domain Filtering

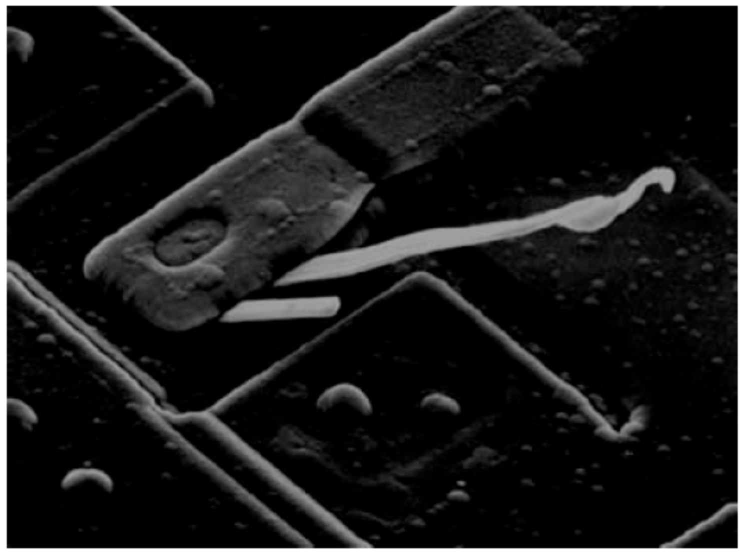

FIGURE 4.30 Result of filtering the image in Fig. 4.29(a) by setting to 0 the term $F(M/2, N/2)$ in the Fourier transform.

Multiplying F(0,0) by zero makes the average image value zero— these are negative image values which cannot be displayed, but get cropped to zero!

Gaussian Filters

Frequency domain filters

Corresponding spatial domain filters

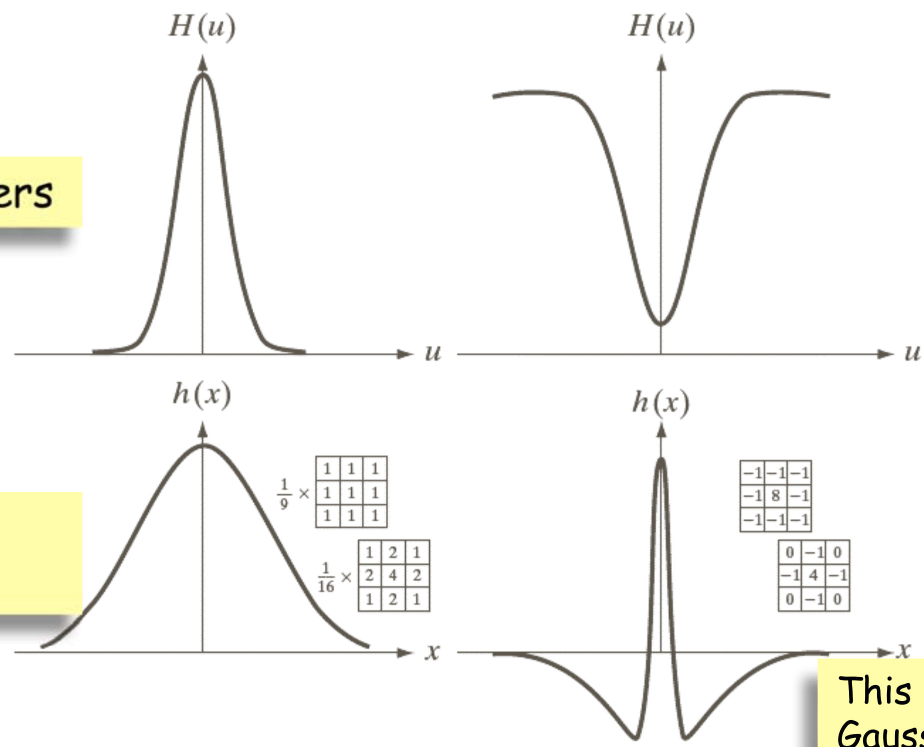

a c
b d

FIGURE 4.37 (a) A 1-D Gaussian lowpass filter in the frequency domain. (b) Spatial lowpass filter corresponding to (a). (c) Gaussian highpass filter in the frequency domain. (d) Spatial highpass filter corresponding to (c). The small 2-D masks shown are spatial filters we used in Chapter 3.

This is a Difference of Gaussians, called a DoG

Low-pass filters

High-pass filters

Gaussian Filters

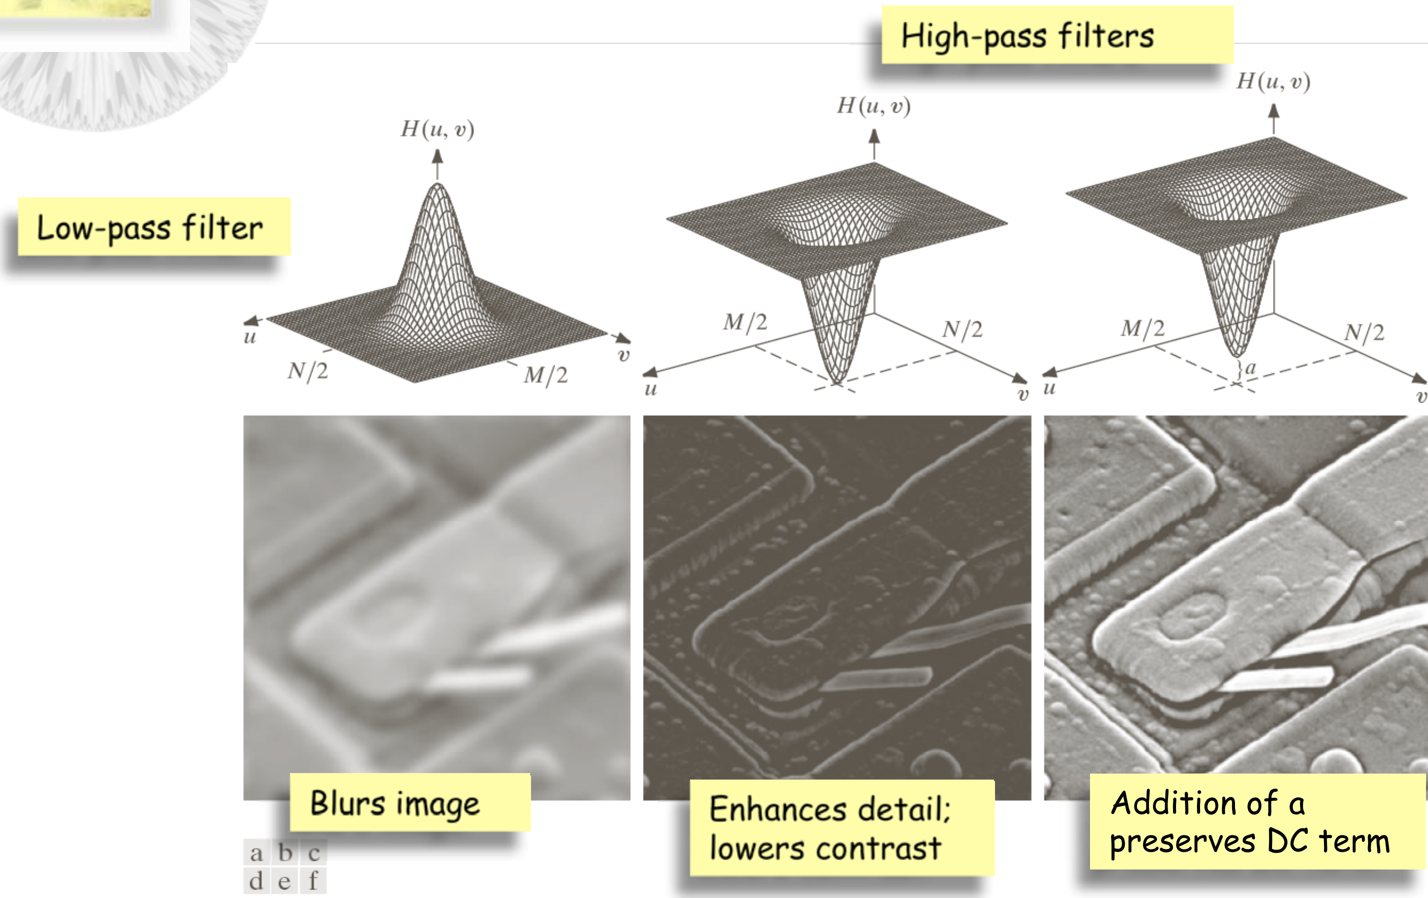

FIGURE 4.31 Top row: frequency domain filters. Bottom row: corresponding filtered images obtained using Eq. (4.7-1). We used $a = 0.85$ in (c) to obtain (f) (the height of the filter itself is 1). Compare (f) with Fig. 4.29(a).

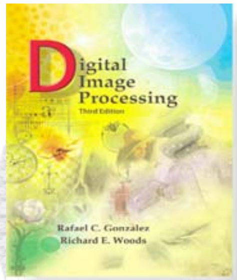

Frequency Domain Filters

1. Original frequency domain filter H(f)

2. Spatial domain filter h(x)= $\mathcal{F}^{-1}\{H(f)\}$

3. Padded spatial domain filter $h_p(x)$

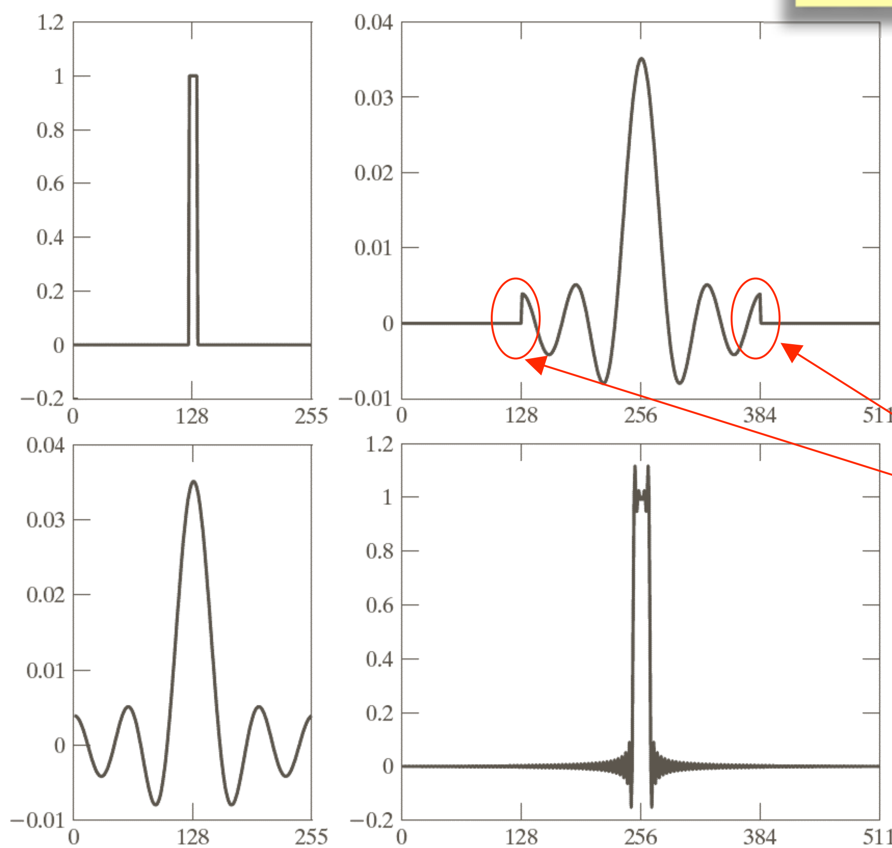

a c
b d

FIGURE 4.34 (a) Original filter specified in the (centered) frequency domain. (b) Spatial representation obtained by computing the IDFT of (a). (c) Result of padding (b) to twice its length (note the discontinuities). (d) Corresponding filter in the frequency domain obtained by computing the DFT of (c). Note the ringing caused by the discontinuities in (c). (The curves appear continuous because the points were joined to simplify visual analysis.)

Discontinuities caused by padding

4. Fourier transform of passed spatial domain filter $h_p(x)$

Phase Filters

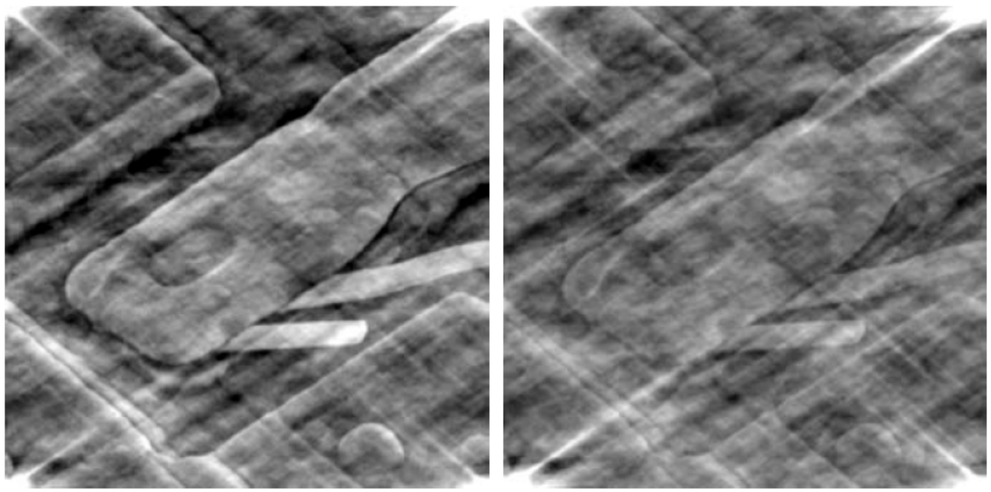

a b

FIGURE 4.35 (a) Image resulting from multiplying by 0.5 the phase angle in Eq. (4.6-15) and then computing the IDFT. (b) The result of multiplying the phase by 0.25. The spectrum was not changed in either of the two cases.

Results of multiplying ONLY the phase of an image by a constant

MOST image processing filters operate ONLY on the magnitude and are called zero-phase-shift filters

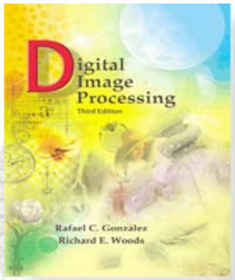

Frequency Domain Filtering Example

a. Original image
b. Padded image
c. Multiply by $(-1)^{x+y}$
d. Fourier transform F(u,v)
e. Centered Gaussian LPF H(u,v)
f. Compute F(u,v)H(u,v)
g. Compute $\mathcal{F}^{-1}\{F(u,v)H(u,v)\}$ and multiply by $(-1)^{x+y}$
h. Crop to remove padding

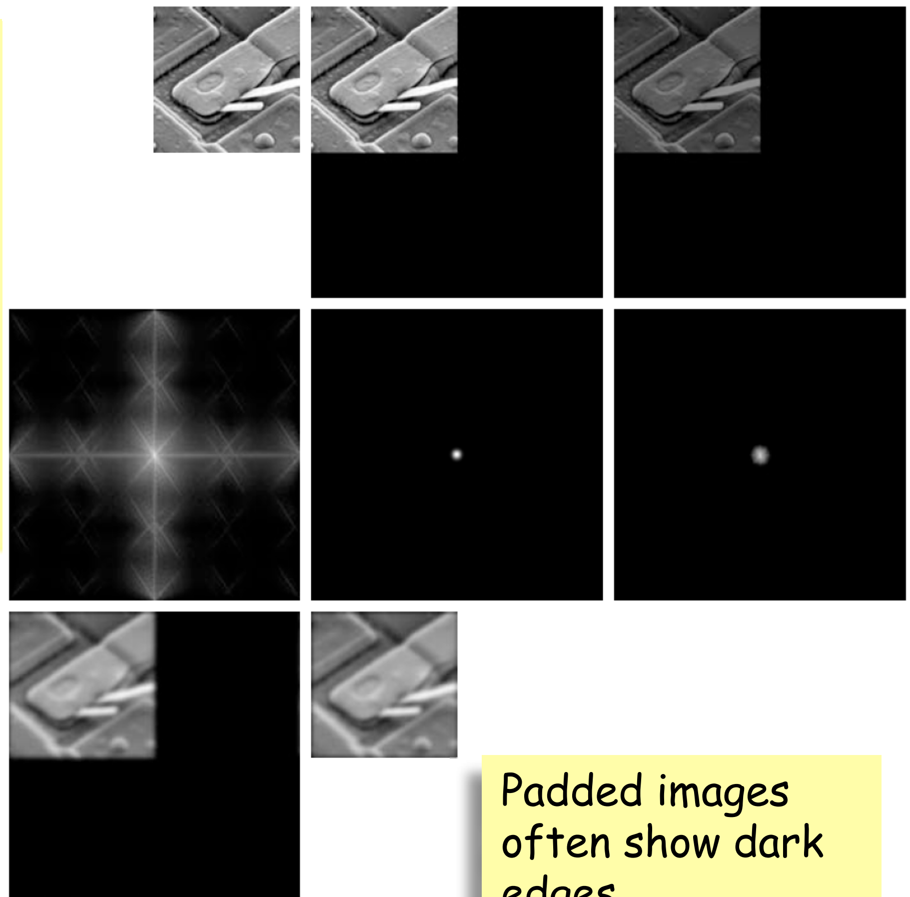

Padded images often show dark edges

a b c
d e f
g h

FIGURE 4.36 (a) An $M \times N$ image, f. (b) Padded image, f_p of size $P \times Q$. (c) Result of multiplying f_p by $(-1)^{x+y}$. (d) Spectrum of F_p. (e) Centered Gaussian lowpass filter, H, of size $P \times Q$. (f) Spectrum of the product HF_p. (g) g_p, the product of $(-1)^{x+y}$ and the real part of the IDFT of HF_p. (h) Final result, g, obtained by cropping the first M rows and N columns of g_p.

Fourier Transform of an Image

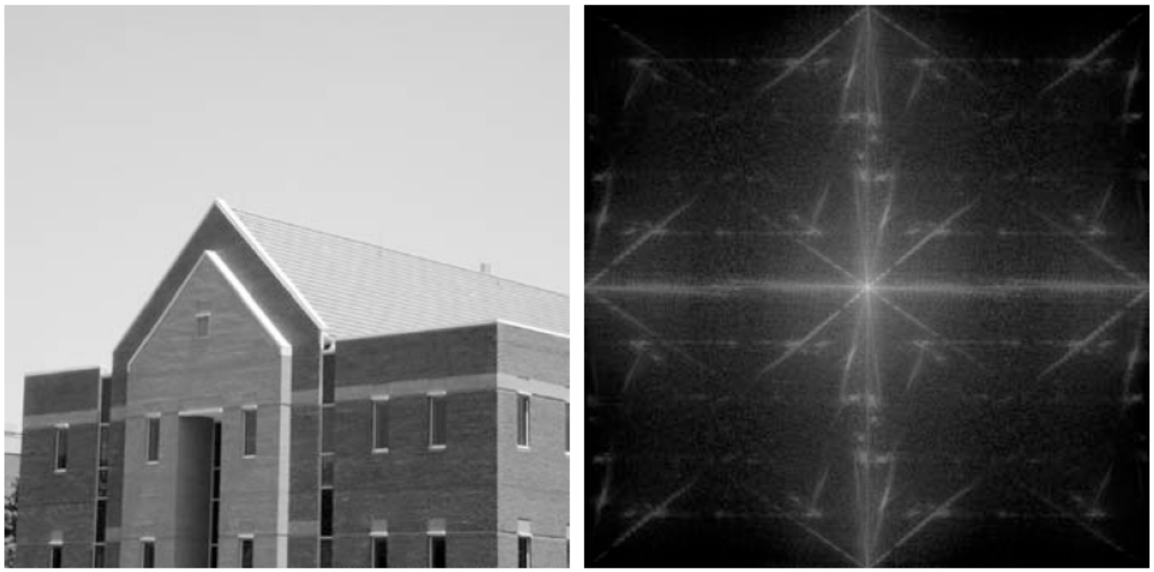

a b

FIGURE 4.38 (a) Image of a building, and (b) its spectrum.

600x600 pixels

F(u,v)

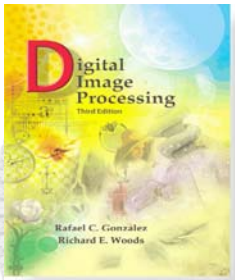

Sobel spatial/frequency domain filters

Spatial domain Sobel, 3x3

−1	0	1
−2	0	2
−1	0	1

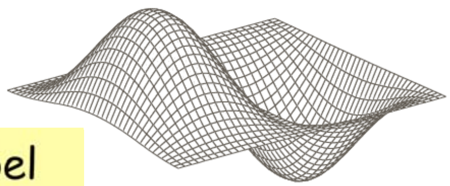

Frequency domain Sobel

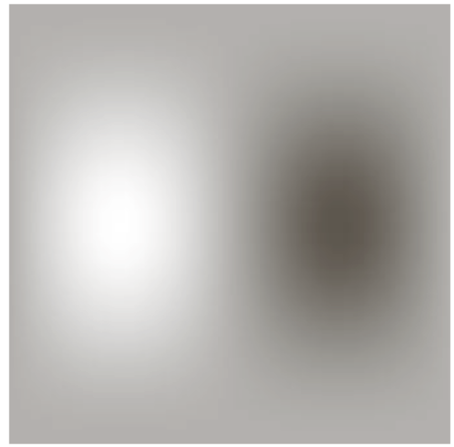

Frequency domain filtering where h(x,y) was padded to 602x602 and centered prior to multiplying by F(u,v) in the frequency domain

a b
c d

FIGURE 4.39 (a) A spatial mask and perspective plot of its corresponding frequency domain filter. (b) Filter shown as an image. (c) Result of filtering Fig. 4.38(a) in the frequency domain with the filter in (b). (d) Result of filtering the same image with the spatial filter in (a). The results are identical.

Spatial domain filtering

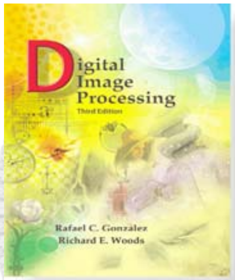

Ideal Low-pass Filter

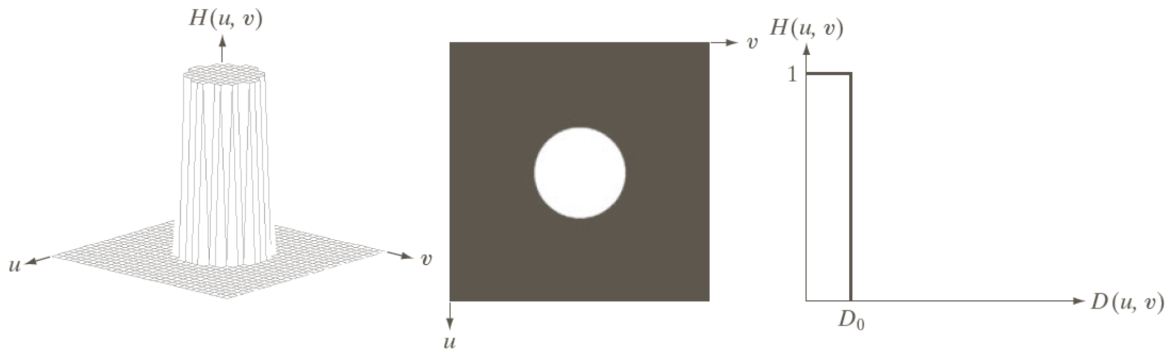

a b c

FIGURE 4.40 (a) Perspective plot of an ideal lowpass-filter transfer function. (b) Filter displayed as an image. (c) Filter radial cross section.

D_o is picked based upon the fraction of image power to be removed

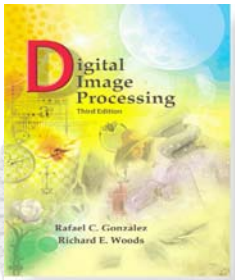

Test Pattern

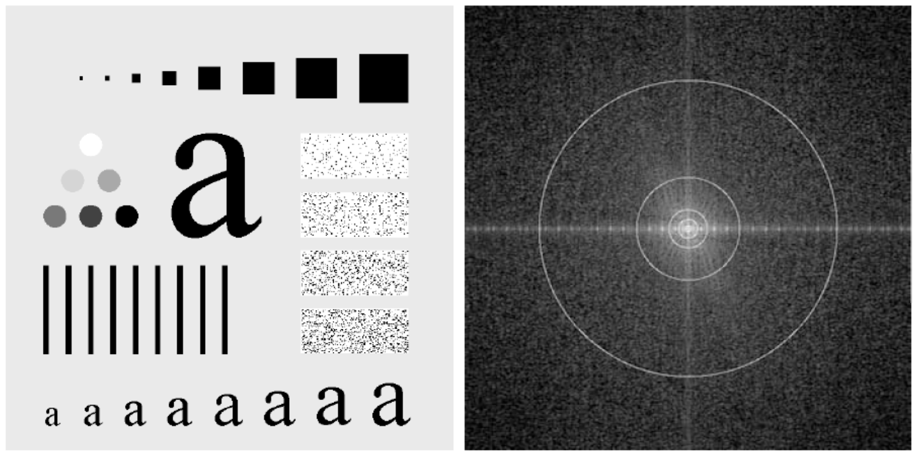

a b

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.

Magnitude transform of padded test image

$$\Delta u = \frac{1}{M\Delta x}$$

If M=688 and, say, there are 688 pixels/inch each pixel in the Fourier transforms corresponds to Δu=1/inch

Filtering with Ideal LPF

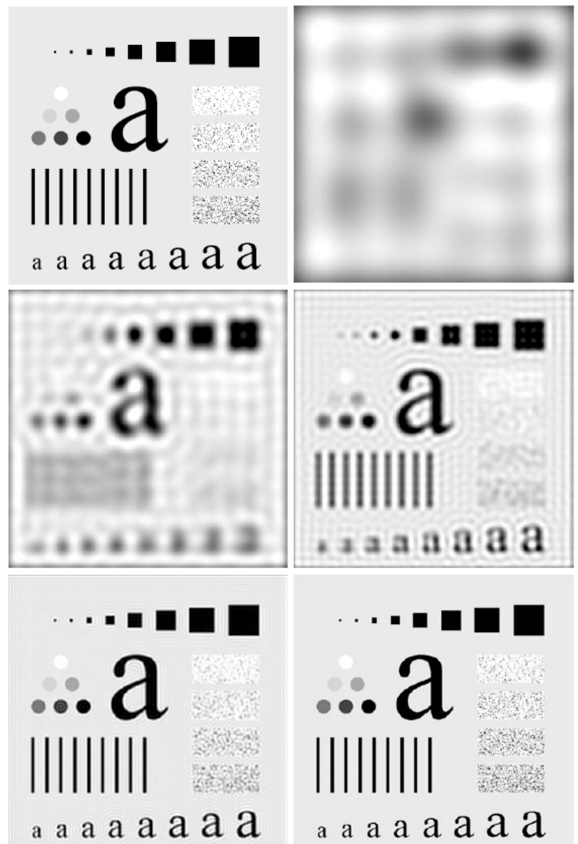

a b
c d
e f

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.

Filtering the padded test image with an ILPF in the frequency domain

Ringing

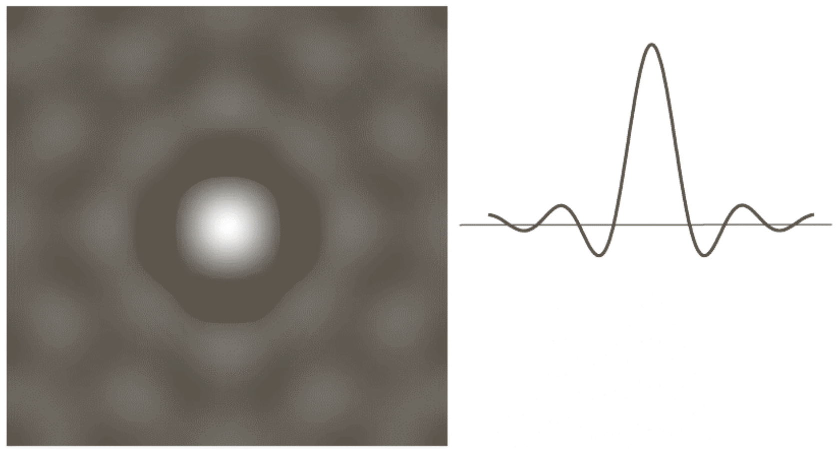

a b

FIGURE 4.43 (a) Representation in the spatial domain of an ILPF of radius 5 and size 1000×1000. (b) Intensity profile of a horizontal line passing through the center of the image.

Computing the inverse Fourier transform of the ILPF shows ringing in the spatial domain which is clearly shown in Figure 4.42

Butterworth LPF

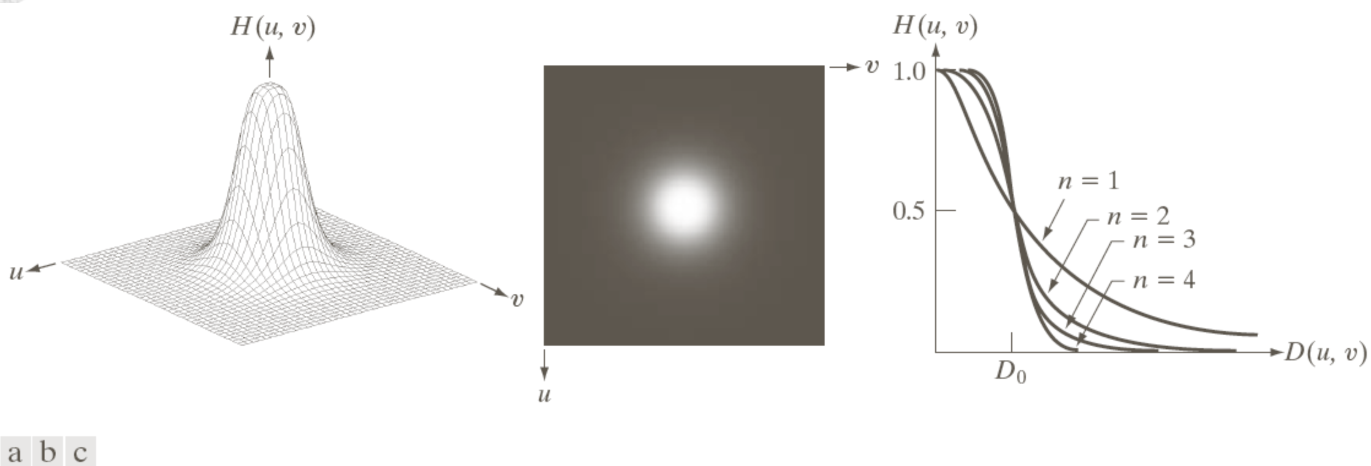

FIGURE 4.44 (a) Perspective plot of a Butterworth lowpass-filter transfer function. (b) Filter displayed as an image. (c) Filter radial cross sections of orders 1 through 4.

$$H(u,v)=\frac{1}{1+\left[\frac{D(u,v)}{D_0}\right]^{2n}}$$

For n=1 there is no ringing. The ringing will increase as n increases.

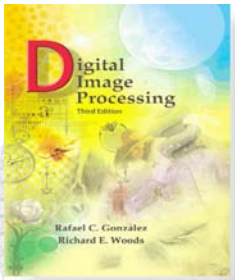

Filtering with Butterworth filter, n=2

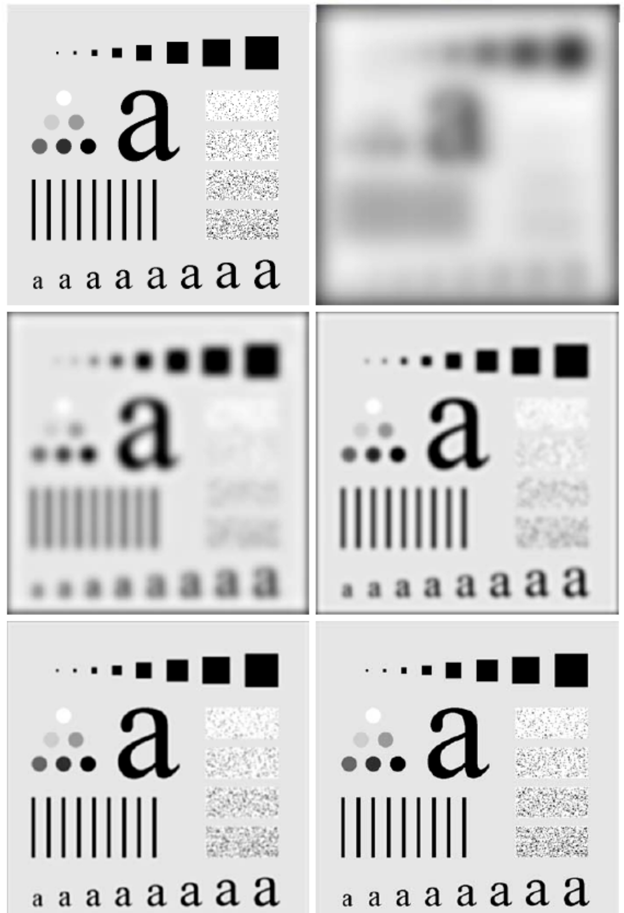

a b
c d
e f

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.

Spatial Butterworth filters

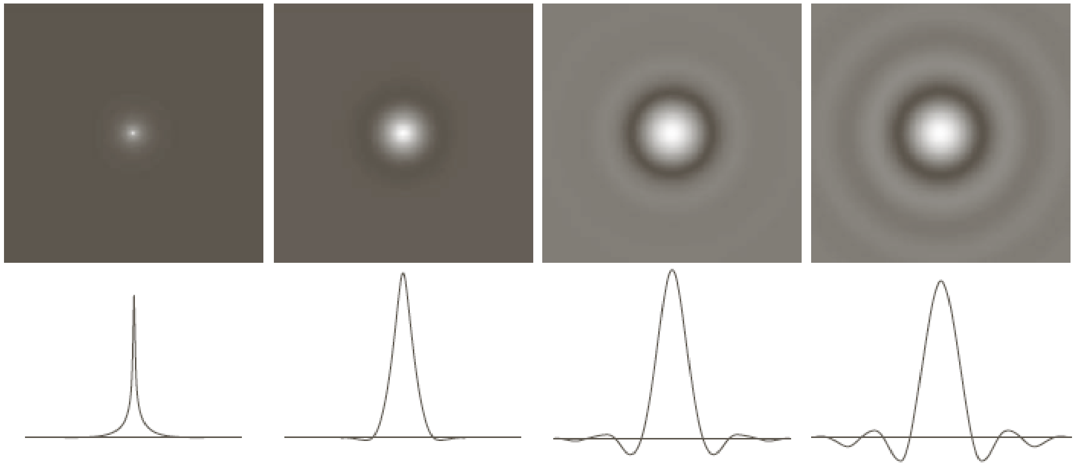

a b c d

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.

Gaussian LPF

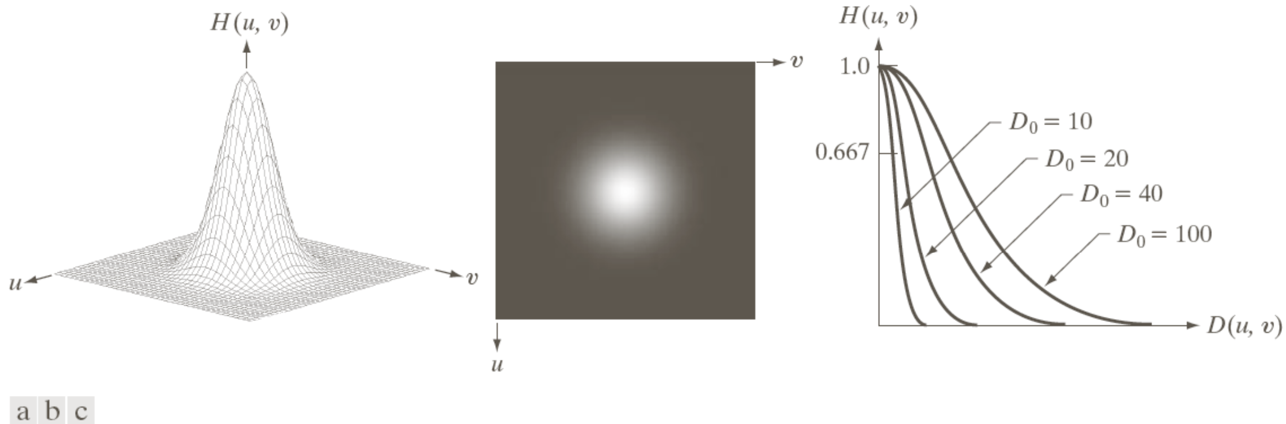

FIGURE 4.47 (a) Perspective plot of a GLPF transfer function. (b) Filter displayed as an image. (c) Filter radial cross sections for various values of D_0.

$$H(u,v) = e^{-\frac{D^2(u,v)}{2D_0^2}}$$

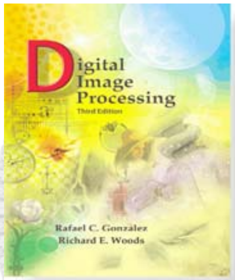

Low-pass Filters

TABLE 4.4
Lowpass filters. D_0 is the cutoff frequency and n is the order of the Butterworth filter.

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

Filtering with Gaussian LPF

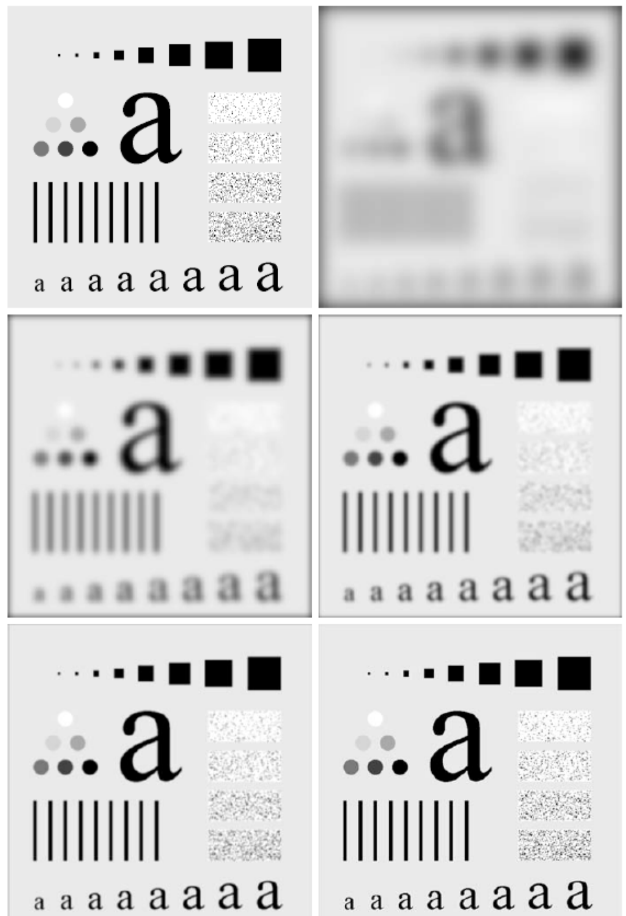

a b
c d
e f

FIGURE 4.48 (a) Original image. (b)–(f) Results of filtering using GLPFs with cutoff frequencies at the radii shown in Fig. 4.41. Compare with Figs. 4.42 and 4.45.

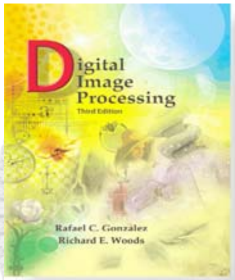

Filtering for OCR

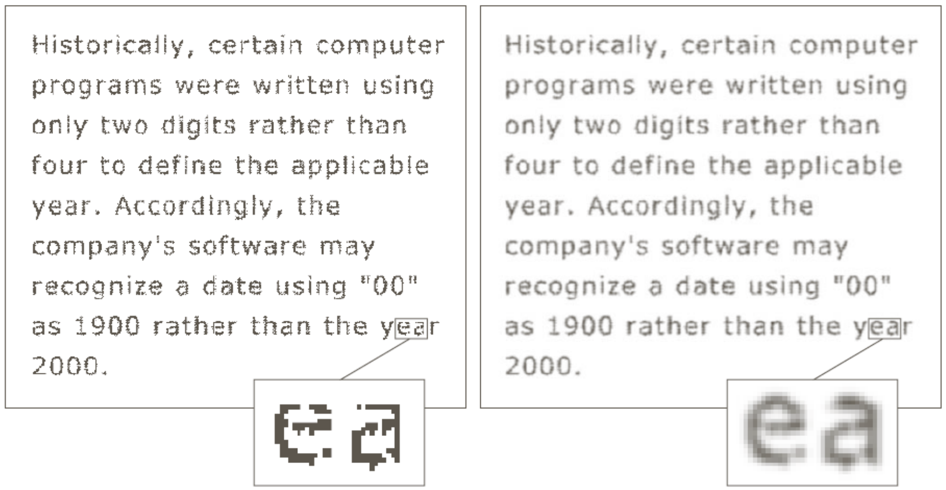

a b

FIGURE 4.49 (a) Sample text of low resolution (note broken characters in magnified view). (b) Result of filtering with a GLPF (broken character segments were joined).

No filtering. Broken characters are difficult to recognize.

Filtered with a GLPF with D_0=80. Characters are fuller and filled in.