4.7 The Basics of Filtering in the Frequency Domain

Additional Characteristics of the Frequency Domain

By observing the 2-D discrete Fourier transform (DFT)

$$F(u,v) = \sum_{x=0}^{M-1}\sum_{y=0}^{N-1} f(x,y)e^{-j2\pi(ux/M+vy/N)}, \qquad (4.5\text{-}15)$$

we find that each term of $F(u,v)$ contains all values of $f(x,y)$, and it usually is impossible to make direct associations between specific components of an image and its transform.

However, we can make some general statements about the relationship between the frequency components of the Fourier transform and spatial features of an image.

For instance, since frequency is directly related to spatial rates of changes, frequencies in the Fourier transform are associated with patterns of intensity variations in an image. One of examples is that the slowest varying frequency component, $u = v = 0$, is proportional to the average intensity of an image.

Filtering techniques in the frequency domain are based on modifying the Fourier transform to achieve a specific objective and then computing the IDFT to get us back to the image domain.

According to

$$F(u,v) = |F(u,v)|e^{j\phi(u,v)}, \qquad (4.6\text{-}15)$$

the two components of the transform to which we have access are the transform magnitude (spectrum) and the phase angle.

Although the visual analysis of the phase component generally is not very useful, however, the spectrum provides some useful guidelines as to gross characteristics of the image from which the spectrum was generated.

Figure 4.29 shows an example.

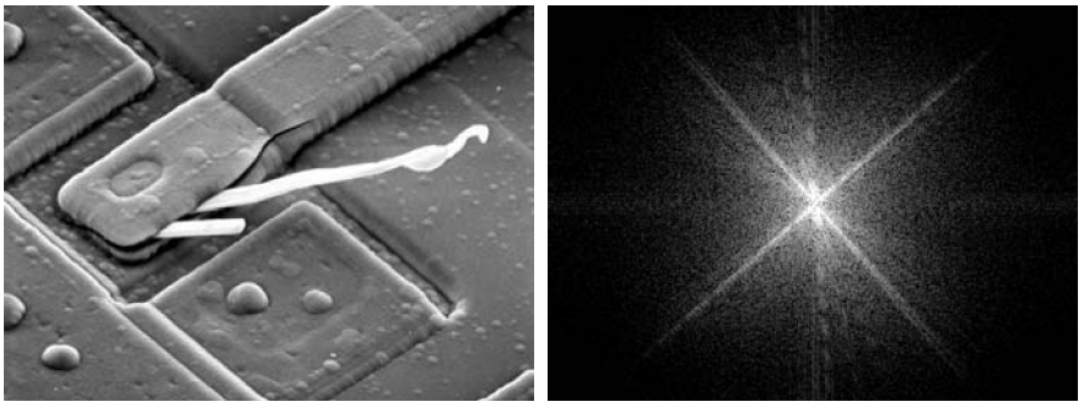

a b

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.)

Figure 4.29 (a) shows an electron microscope image of an integrated circuit. It has two notable features: strong edges that run approximately at $\pm 45^o$ and two white oxide protrusions.

The Fourier spectrum in Figure 4.29 (b) shows prominent components along the $\pm 45^o$ directions that correspond to the edges mentioned in Figure 4.29 (a).

There is a vertical component that is off-axis slightly to the left in Figure 4.29 (b). It was caused by the edges of the white oxide protrusions.

The abovementioned facts are typical of the types of associations that can be made in general between the frequency and spatial domains.

The relationships between frequency content and rate of change of intensity levels in an image can lead to some very useful results.

Frequency Domain Filtering Fundamentals

Filtering in the frequency domain consists of modifying the Fourier transform of an image and then computing the inverse transform to obtain the processed result.

Given a digital image $f(x,y)$, of size $M \times N$, the basic filtering equation has the form

$$g(x,y) = \mathcal{F}^{-1}[H(u,v)F(u,v)], \qquad (4.7\text{-}1)$$

where $\mathcal{F}^{-1}$ is the IDFT, $F(u,v)$ is the DFT of the input image, $f(x,y)$, $H(u,v)$ is a filter function, and $g(x,y)$ is the filtered image.

One of the simplest filters we can construct is a filter $H(u,v)$ that is 0 at the center of the transform and 1 elsewhere. This filter will reject the *dc* term and pass all other terms.

Figure 4.30 shows the result of applying this filter.

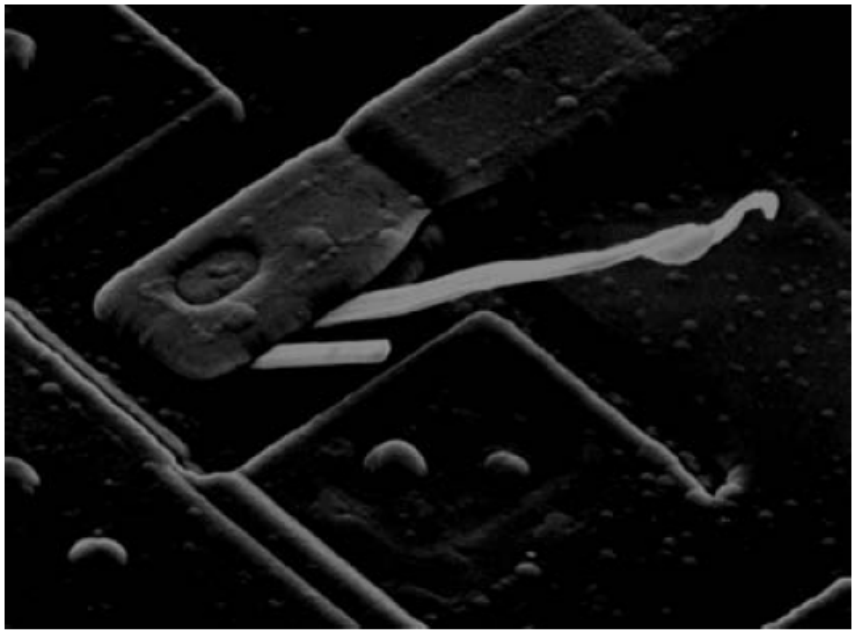

FIGURE 4.30 Result of filtering the image in Fig. 4.29(a) by setting to 0 the term $F(M/2, N/2)$ in the Fourier transform.

Figure 4.30 is not a true representation of the original, as all negative intensities were set to 0 for display purposes.

Since high frequencies are caused by sharp transitions in intensity, such as edges and noise, we would expect a filter $H(u,v)$ that attenuates high frequencies while passing low frequencies would blur an image. This is called a lowpass filter.

On the other hand, a highpass filter would enhance sharp details, but could cause a reduction in contrast in the image.

Figure 4.31 shows these effects.

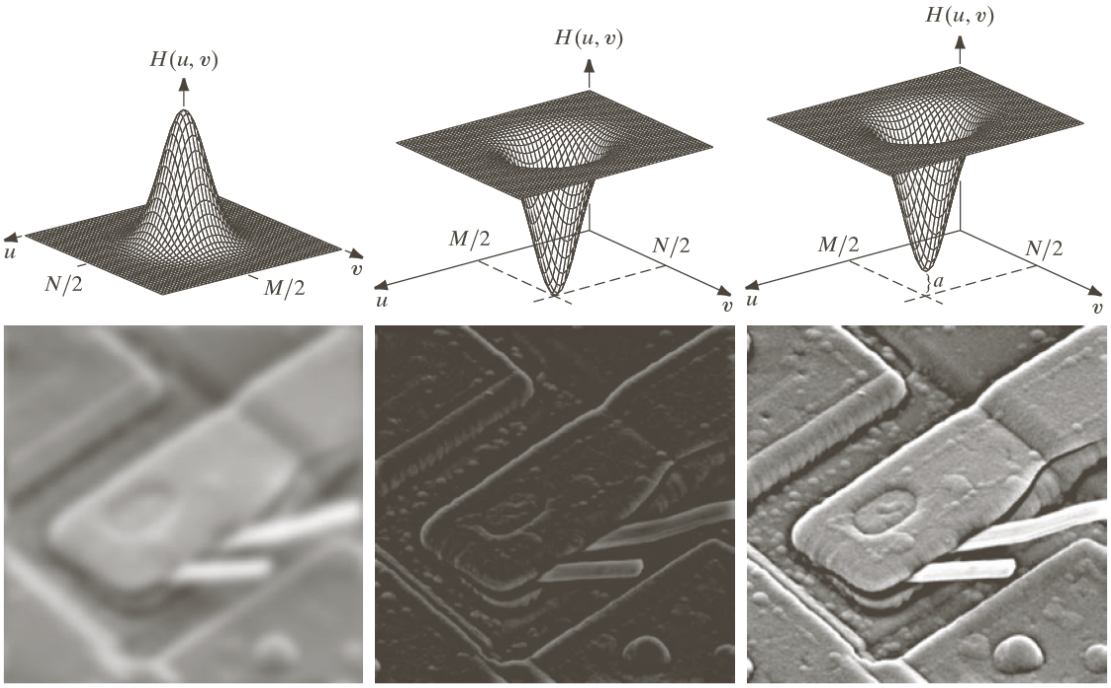

a b c
d e f

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).

We know that, while applying convolution, if the functions are not padded, the wraparound error is expected. Figure 4.32 shows an example.

a b c

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).

Figure 4.33 illustrates the reason for discrepancy between Figure 4.32 (b) and (c).

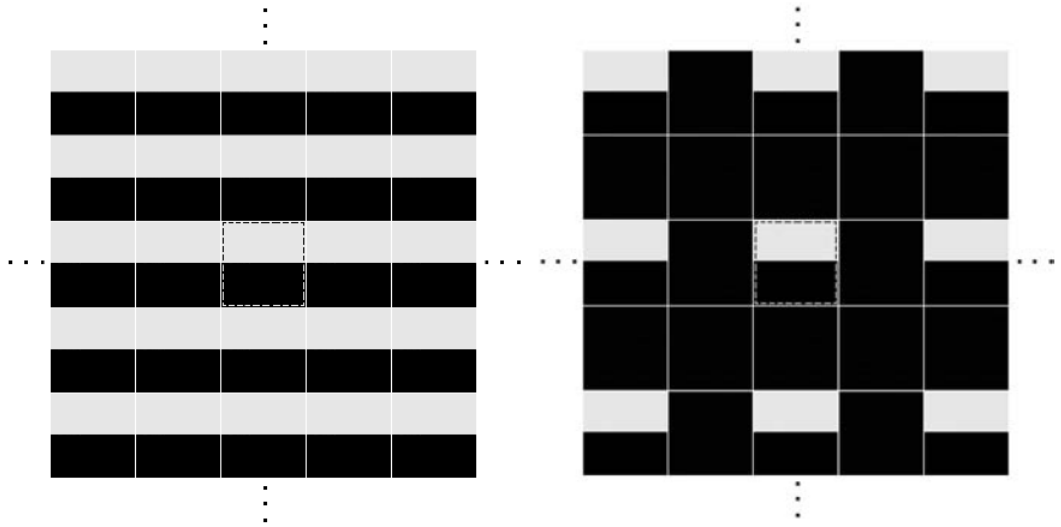

a b

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.)

Although

$$g(x,y) = \mathcal{F}^{-1}[H(u,v)F(u,v)] \qquad (4.7\text{-}1)$$

involves a filter that can be specified either in the spatial or the frequency domain, padding is done only in the spatial domain.

What is the relationship between spatial padding and filters specified directly in the frequency domain?

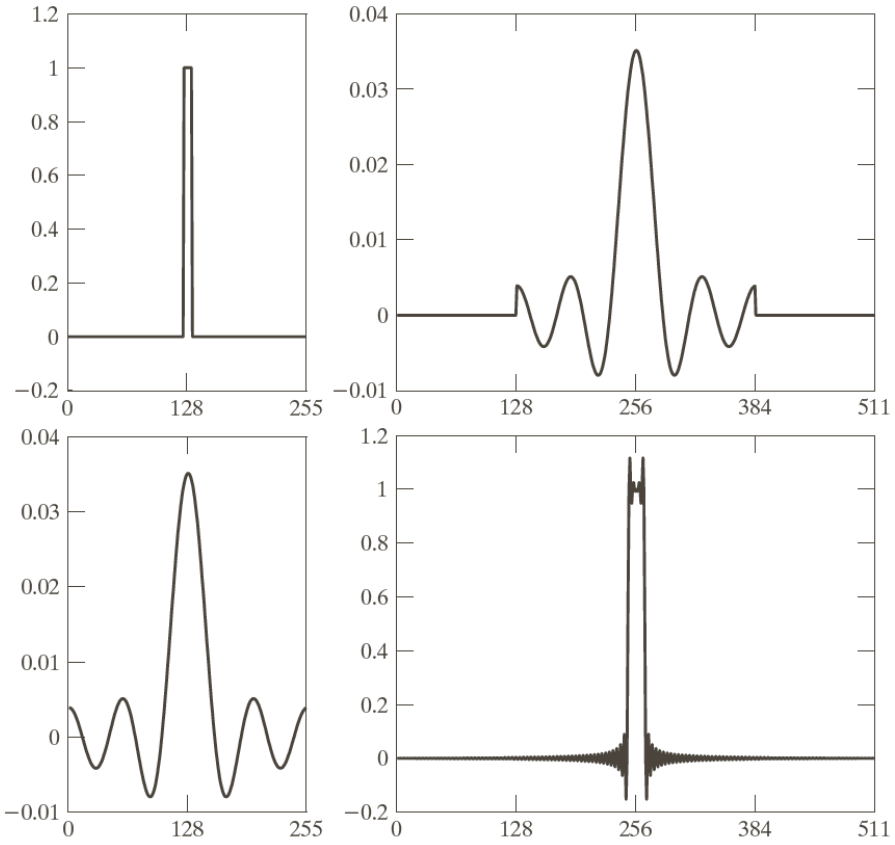

a c
b d

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.)

Figure 4.34 (a) shows a 1-D ideal lowpass filter in the frequency domain. The filter is real and has even symmetry. From Property 8 in Table 4.1

$$f(x,y) \text{ real and even} \quad \Leftrightarrow \quad F(u,v) \text{ real and even}$$

its IDFT will be real and even symmetry as well.

Figure 4.34 (b) shows the corresponding spatial filter.

Figure 4.34 (c) shows the two discontinuities created by zero-padding.

To get back to the frequency domain, we compute the DFT of the zero-padded spatial filter. Figure 4.34 (d) shows the result. The discontinuities in the spatial filter have created ringing in its frequency counterpart.

Any spatial truncation of the filter to implement zero-padding will introduce discontinuities, which will in general result in ringing in the frequency domain.

Now we analyze the phase angle of the filtered transform.

Since the DFT is a complex array, we can express it in terms of its real and imaginary parts:

$$F(u,v) = R(u,v) + jI(u,v)\,. \qquad (4.7\text{-}2)$$

Equation (4.7-1) then becomes

$$g(x,y) = \mathcal{F}^{-1}\left[H(u,v)R(u,v) + jH(u,v)I(u,v)\right], \qquad (4.7\text{-}3)$$

which will not alter the phase angle because $H(u,v)$ cancels out when the ratio of the imaginary and real parts is formed in

$$\phi(u,v) = \arctan\left[\frac{I(u,v)}{R(u,v)}\right]. \qquad (4.6\text{-}17)$$

Filters that affect the real and imaginary parts equally, and thus have no effect on the phase, are called zero-phase-shift filters.

Even small changes in the phase angle can have dramatic effects on the filtered output.

Figure 4.35 shows the effect of a scalar change.

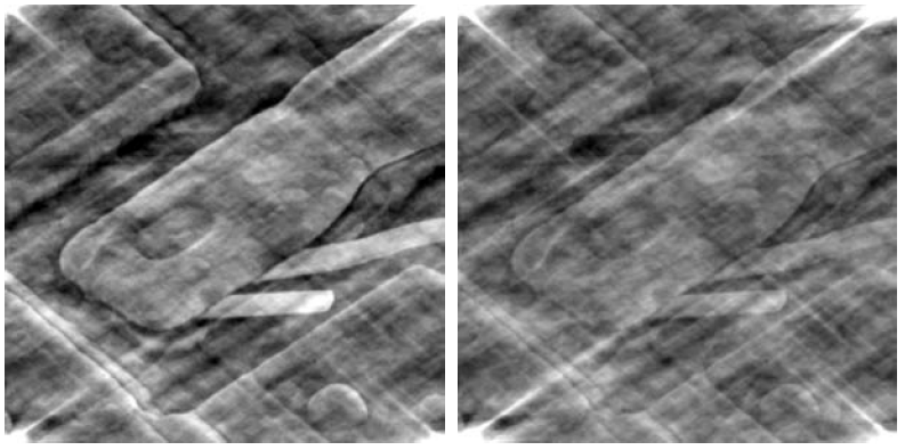

a b

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.

Figure 4.35 (a) shows an image resulting from multiplying the angle array in

$$F(u,v) = |F(u,v)|e^{j\phi(u,v)} \qquad (4.6\text{-}15)$$

by 0.5, without changing $|F(u,v)|$, and then computing the IDFT. Although the basic shapes remain unchanged, the intensity distribution is quite distorted.

Figure 4.35 (b) shows the result of multiplying the phase by 0.25. The image is close to unrecognizable.

Summary of Steps for Filtering in the Frequency Domain

1. Given an input image $f(x,y)$ of size $M \times N$, obtain the padding parameters P and Q from

   $$P \geq 2M - 1 \qquad \text{(4.6-31)}$$

   and

   $$Q \geq 2N - 1. \qquad \text{(4.6-32)}$$

   Typically, we select $P = 2M$ and $Q = 2N$.

2. Form a padded image, $f_p(x,y)$, of size $P \times Q$ by appending the necessary number of zeros to $f(x,y)$.

3. Multiply $f_p(x,y)$ by $(-1)^{x+y}$ to center its transform.

4. Compute the DFT, $F(u,v)$, of the image from Step 3.

5. Generate a real, symmetric filter function, $H(u,v)$, of size $P \times Q$ with center at $(P/2, Q/2)$. Form the product $G(u,v) = H(u,v)F(u,v)$ using array multiplication $G(i,k) = H(i,k)F(i,k)$.

6. Obtain the processed image:

   $$g_p(x,y) = \left\{\text{real}\left[\mathcal{F}^{-1}[G(u,v)]\right]\right\}(-1)^{x+y}$$

   where the real part is selected in order to ignore parasitic complex components resulting from computational in accuracies.

7. Obtain the final processed result, $g(x,y)$, by extracting the $M \times N$ region from the top, left quadrant of $g_p(x,y)$.

Figure 4.36 illustrates the preceding steps.

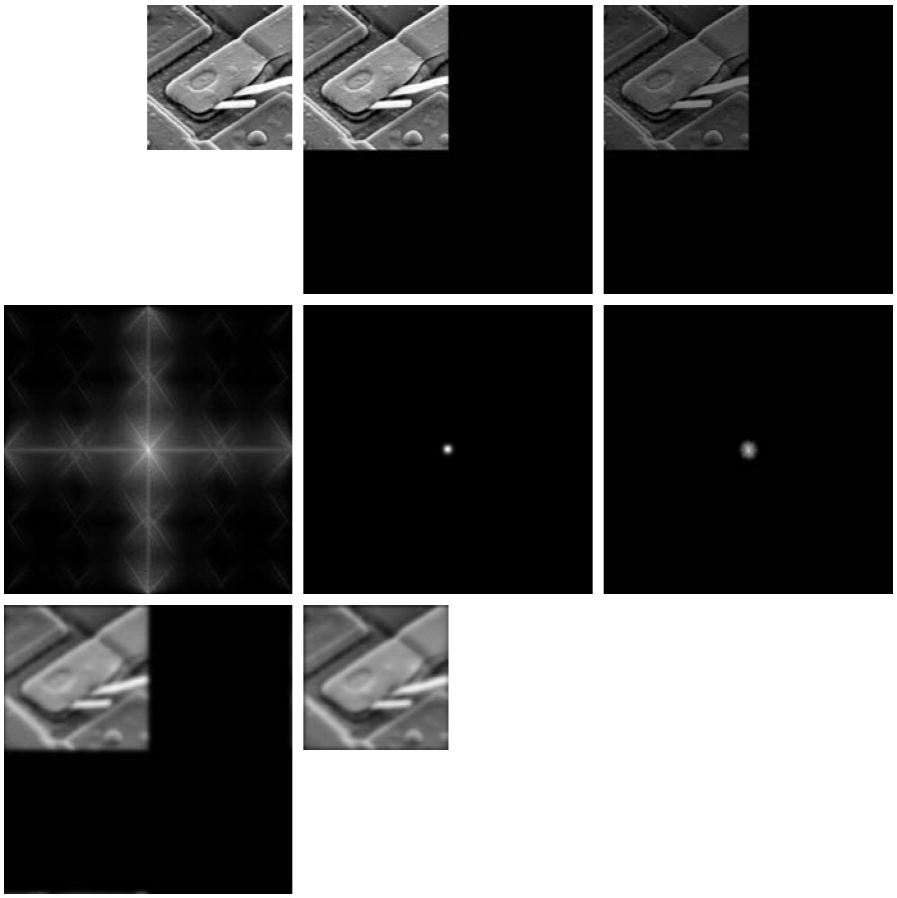

a b c
d e f
g h

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.

Correspondence between Filtering in the Spatial and Frequency Domains

The link between filtering in the spatial and frequency domains is the convolution theorem.

Suppose we want to find the equivalent representation of $H(u,v)$ in the spatial domain. If we let

$$f(x,y) = \delta(x,y)\,,$$

it follows from Table 4.3 that

$$F(u,v) = 1\,.$$

Then, from

$$\mathrm{g}(x,y) = \mathcal{F}^{-1}\left[H(u,v)F(u,v)\right], \qquad (4.7\text{-}1)$$

the filtered output is $\mathcal{F}^{-1}\{H(u,v)\}$. This is the inverse transform of the frequency domain filter, which is the corresponding filter in the spatial domain.

Conversely, we can obtain the frequency domain representation by taking the forward Fourier transform of the spatial filter. Therefore, the two filters form a Fourier transform pair

$$h(x,y) \Leftrightarrow H(u,v)\,. \qquad (4.7\text{-}4)$$

Since $h(x,y)$ can be obtained from the response of a frequency domain filter to an impulse, it is referred to the impulse response of $H(u,v)$.

In practice, we prefer to implement convolution filtering using

$$w(x,y)\star f(x,y) = \sum_{s=-a}^{a}\sum_{t=-b}^{b} w(s,t)f(x-s,y-t) \qquad (3.4\text{-}2)$$

with smaller filter masks because of easier computation and implementation.

Filtering concepts are more intuitive in the frequency domain.

One way to use the properties of both domains is to specify a filter in the frequency domain, compute its IDFT, and then use the resulting full-size spatial filter as a guide to construct smaller spatial filter masks.

Since both the forward and inverse Fourier transforms of a Gaussian function are real Gaussian functions, the Gaussian filters are used here to show how frequency domain filters can be a guide to construct smaller spatial filter masks.

Let $H(u)$ be the 1-D frequency domain Gaussian filter

$$H(u) = Ae^{-u^2/2\sigma^2} \qquad (4.7\text{-}5)$$

where σ is the standard deviation of the Gaussian curve.

The corresponding filter in the spatial domain is obtained by taking IDFT of $H(u)$:

$$h(x) = \sqrt{2\pi}\sigma A e^{-\pi^2\sigma^2 x^2}. \qquad (4.7\text{-}6)$$

Equations (4.7-5) and (4.7-6) are a Fourier transform pair, and both of their components are Gaussian and real. Also, these functions behave reciprocally.

When the value of σ increases, $h(x)$ has a narrow profile. In fact, as σ approaches infinity, $H(u)$ tends toward a constant function and $h(x)$ tends toward an impulse, which means that there is no filtering in both domains.

Example:

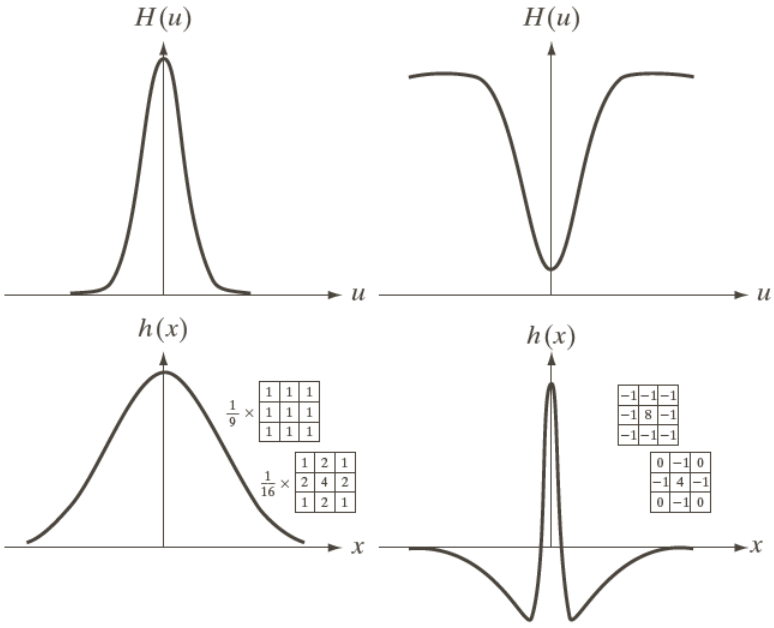

a c
b d

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.

More complex filters can be constructed using the basic Gaussian function of (4.7-5).

We can construct a highpass filter as the different of Gaussians:

$$H(u) = Ae^{-u^2/2\sigma_1^2} - Be^{-u^2/2\sigma_2^2} \qquad (4.7\text{-}7)$$

with $A \geq B$ and $\sigma_1 > \sigma_2$. The corresponding filter in the spatial domain is

$$h(x) = \sqrt{2\pi}\sigma_1 Ae^{-2\pi^2\sigma_1^2 x^2} - \sqrt{2\pi}\sigma_2 Be^{-2\pi^2\sigma_2^2 x^2} \qquad (4.7\text{-}8)$$

Figure 4.37 (c) and (d) show plots of these two functions.

The most important feature here is that $h(x)$ has a positive center term with negative terms on either side. The two masks shown in Figure 4.37 (d) were used in Chapter 3 as sharpening filters.

In practice, the frequency domain can be viewed as a "laboratory" in which we take advantage of the correspondence between frequency content and image appearance.

Example 4.15: Obtaining a frequency domain filter from a small spatial mask.

Figure 4.38 (a) shows a 600×600 image, $f(x,y)$. Figure 4.38 (b) shows its spectrum.

a b

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

Figure 4.39 (a) shows the Sobel mask, $h(x,y)$. To avoid wraparound error, we have padded f and h to size 602×602.

−1	0	1
−2	0	2
−1	0	1

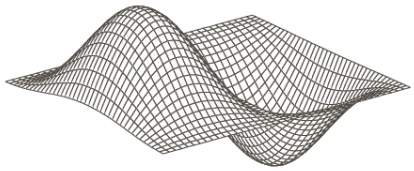

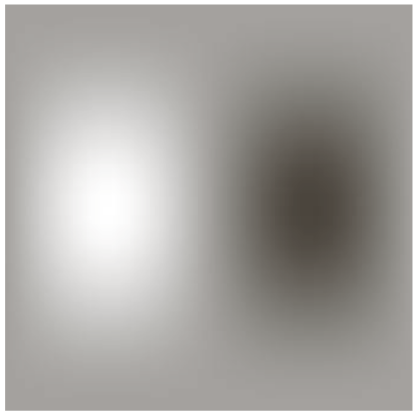

a b
c d

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.

The Sobel mask exhibits odd symmetry, provided that it is embedded in an array of zeros of even size. To maintain this symmetry, we place $h(x, y)$ so that its center is at the center of the 602×602 padded array.

If we preserve the odd symmetry with respect to the padded array in forming $h_p(x, y)$, according to the Property 9 in Table 4.1, $H(u, v)$ will be purely imaginary. This will yield results that are identical to filtering the image spatially using $h(x, y)$.

The procedure used to generate $H(u, v)$ is:

(1) multiply $h_p(x, y)$ by $(-1)^{x+y}$ to center the frequency domain filter;

(2) compute the forward DFT of the result in (1);

(3) set the real part of the resulting DFT to 0 to account for parasitic real part;

(4) and multiply the result by $(-1)^{u+v}$.

Figure 4.39 (a) shows a plot of $H(u, v)$, and Figure 4.39 (b) shows $H(u, v)$ as an image.

Figure 4.39 (c) is the result of using the filter just obtained in the procedure outlined in Page 208 (of the Lecture Notes) to filter the image in Figure 4.38 (a). As expected from a derivative filter, edges are enhanced and all the constant intensity areas are reduced to zero.

Figure 4.39 (d) is the result of filtering the same image in the spatial domain directly, using $h(x, y)$. The results are identical.

4.8 Image Smoothing Using Frequency Domain Filters

Smoothing (blurring) is achieved in the frequency domain by lowpass filtering.

Ideal Lowpass Filters

A 2-D lowpass filter that passes all frequencies within a circle of radius D_0 from the origin and "cuts off" all frequencies outside this circle is called an ideal lowpass filter (ILPF). It is specified by the function

$$H(u,v) = \begin{cases} 1 & \text{if } D(u,v) \le D_0 \\ 0 & \text{if } D(u,v) > D_0 \end{cases} \qquad (4.8\text{-}1)$$

where D_0 is a positive constant, and $D(u,v)$ is the distance between a point (u,v) in the frequency domain and the center of the frequency rectangle

$$D(u,v) = \sqrt{(u - P/2)^2 + (v - Q/2)^2}\,, \qquad (4.8\text{-}2)$$

where P and Q are the padded sized from

$$P \ge 2M - 1 \qquad (4.6\text{-}31)$$

and

$$Q \ge 2N - 1\,. \qquad (4.6\text{-}32)$$

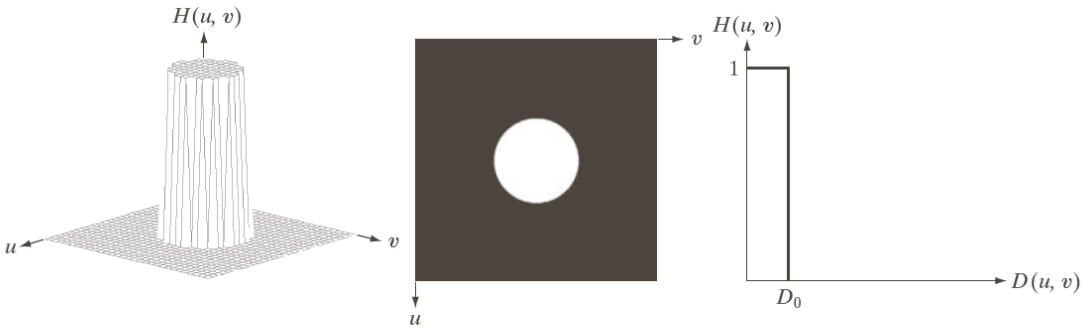

a b c

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

For an ILPF cross section, the point of transition between $H(u,v) = 1$ and $H(u,v) = 0$ is called the cutoff frequency. For example, the cutoff frequency is D_0 in Figure 4.40 (c).

One way to establish a set of standard cutoff frequency loci is to compute circles that enclose specified amounts of total image power P_T. It can be obtained by summing the components of the power spectrum of the padded images at each point (u,v),

$$P_T = \sum_{u=0}^{P-1}\sum_{v=0}^{Q-1} P(u,v) \quad , \qquad (4.8\text{-}3)$$

where $P(u,v)$ is given in

$$P(u,v) = |F(u,v)|^2 = R^2(u,v) + I^2(u,v) \, . \qquad (4.6\text{-}18)$$

If the DFT has been centered, a circle of radius D_0 within the frequency rectangle encloses an α percent of the power, where

$$\alpha = 100\left[\sum_u\sum_v P(u,v) \, / \, P_T\right] . \qquad (4.8\text{-}4)$$

Figure 4.41 shows a test pattern image and its spectrum.

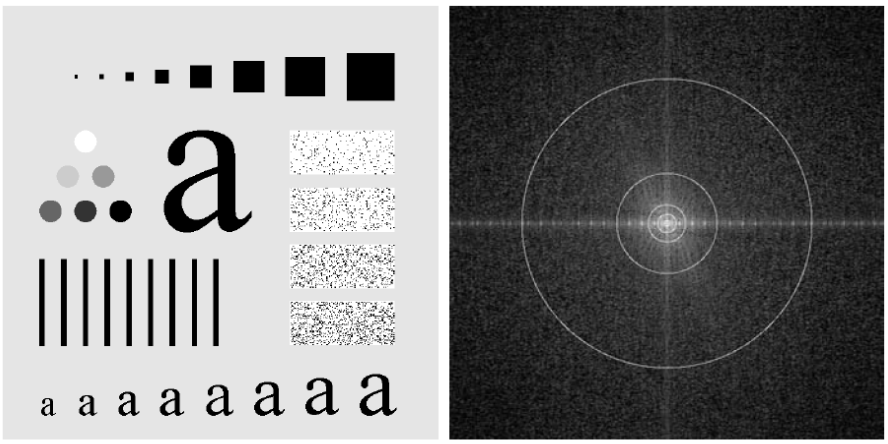

a b

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.

Example 4.16: Image smoothing using an ILPF

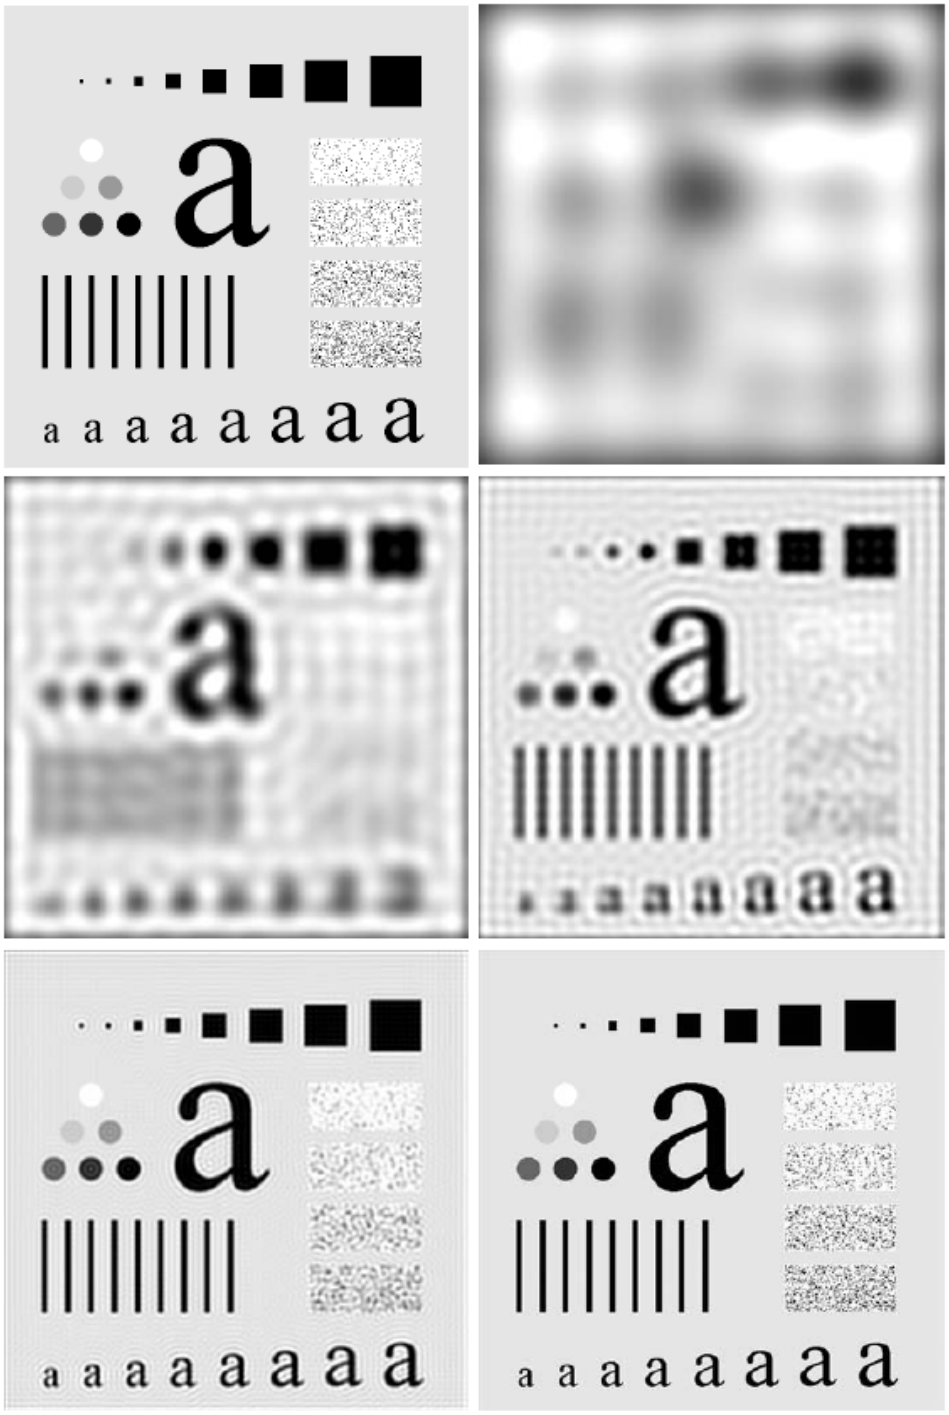

a b
c d
e f

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.

The blurring and ringing properties of ILPFs can be explained using the convolution theorem.

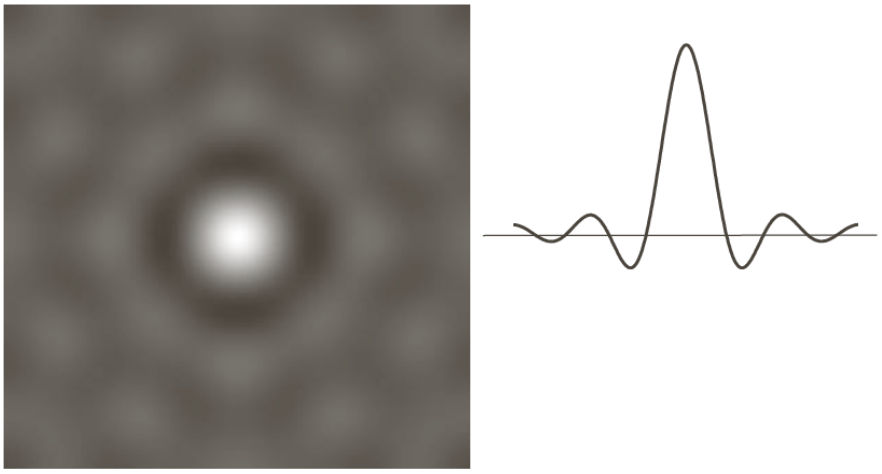

a b

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.

Since a cross section of the LIPF in the frequency domain looks like a box filter, the corresponding spatial filter has the shape of a sinc function.

The center lobe of the sinc is the principal cause of blurring, while the outer, smaller lobes are mainly responsible for ringing.

Butterworth Lowpass Filters

The transfer function of a Butterworth lowpass filter (BLPF) of order n, and with cutoff frequency at a distance D_0 from the origin, is defined as

$$H(u,v) = \frac{1}{1 + [\,D(u,v)\,/\,D_0\,]^{2n}}, \qquad (4.8\text{-}5)$$

where $D(u,v)$ is given by

$$D(u,v) = \sqrt{(u - P\,/\,2)^2 + (v - Q\,/\,2)^2}\,. \qquad (4.8\text{-}2)$$

Figure 4.44 shows a perspective plot, image display, and radial cross sections of the BLPF function.

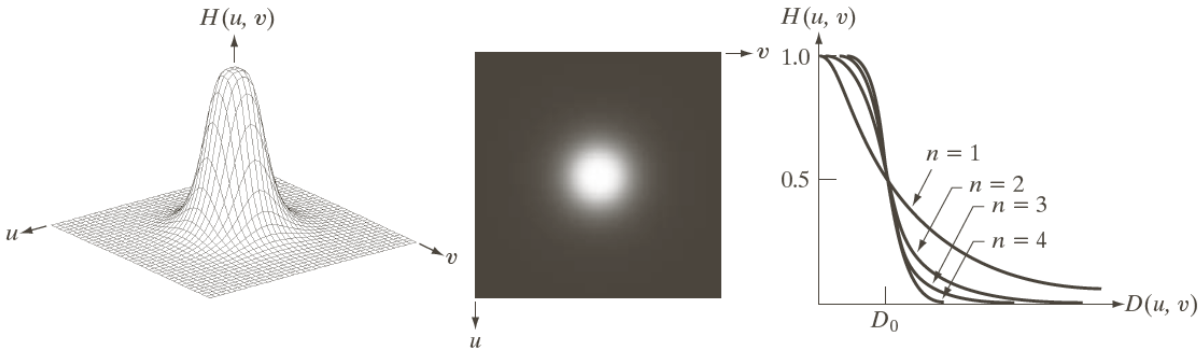

a b c

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.

Unlike the ILPF, the BLPF transfer function does not have a sharp discontinuity.

Example 4.17: Image smoothing with a Butterworth lowpass filter

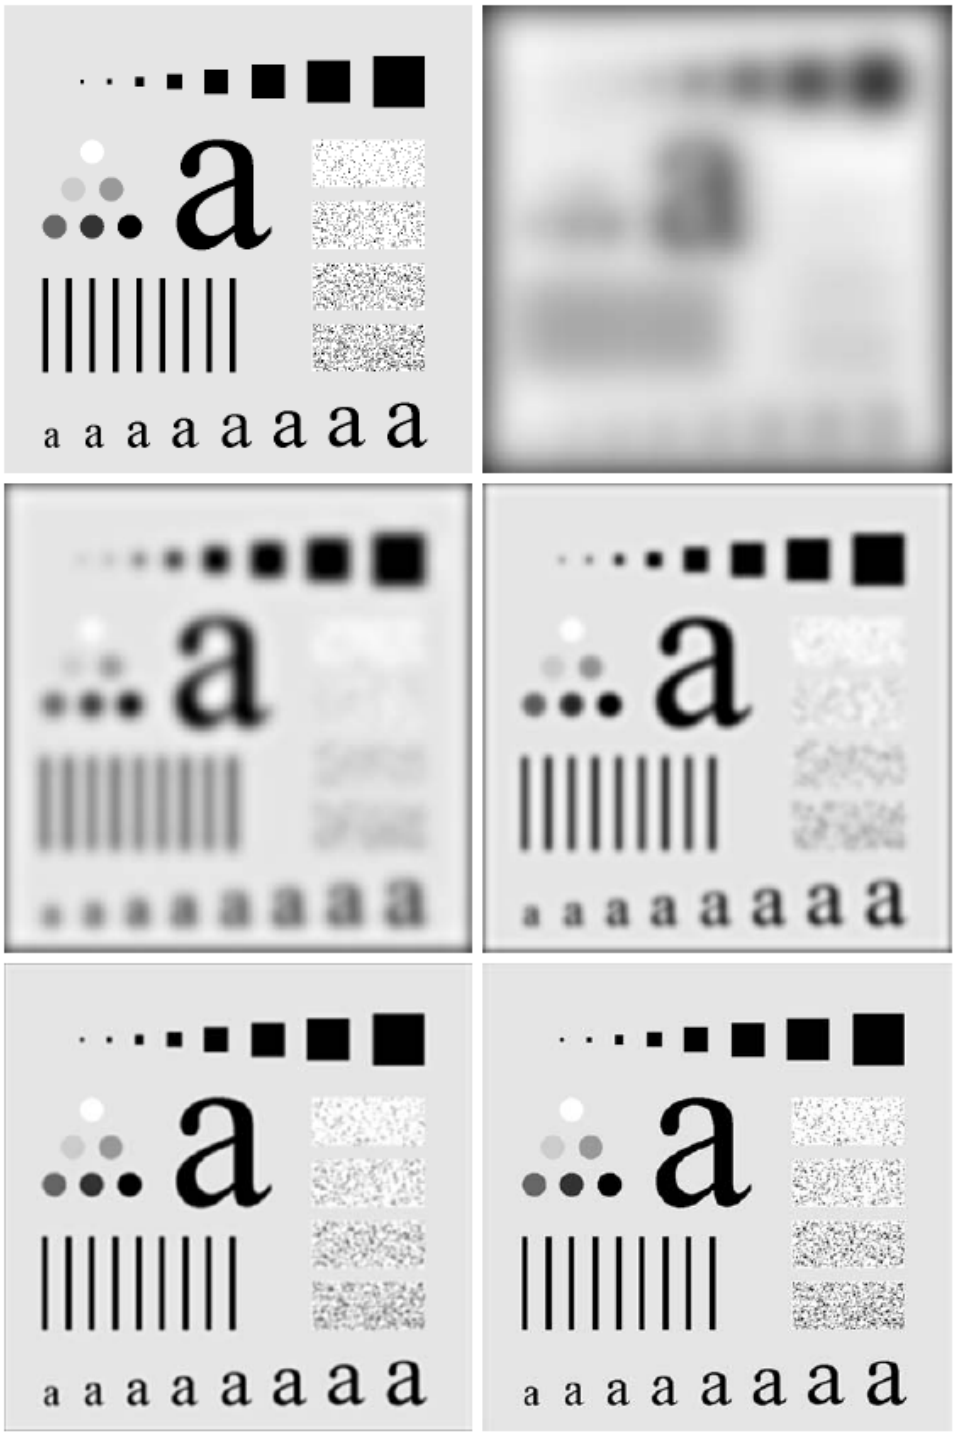

a b
c d
e f

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.

A BLPF of order 1 has no ringing in the spatial domain. Ringing can become significant in filters of higher order.

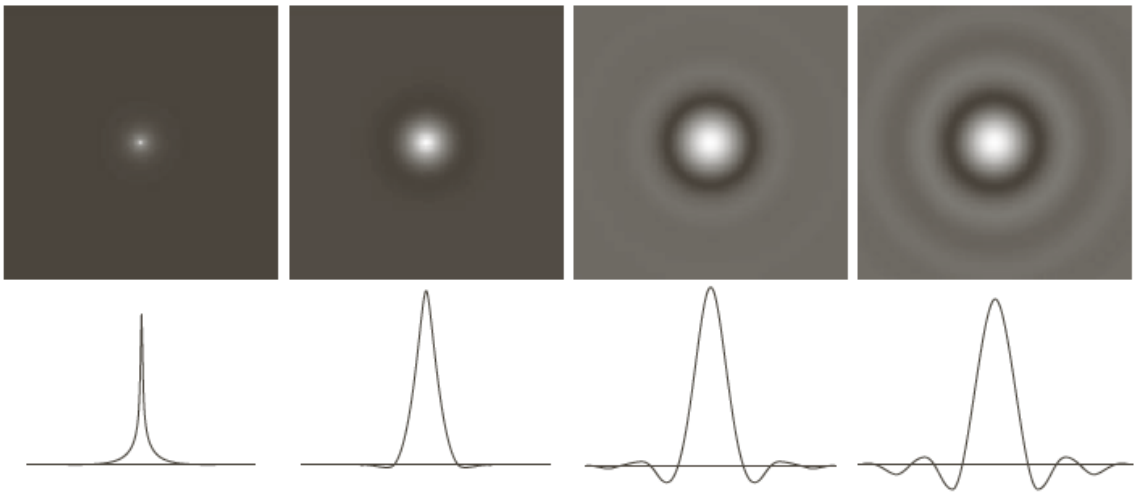

a b c d

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.