UNIVERSITY OF SOUTH CAROLINA

CSCE 590 INTRODUCTION TO IMAGE PROCESSING

Convolution

Ioannis Rekleitis

Spatial Convolution: 1D Signal

Discrete impulse

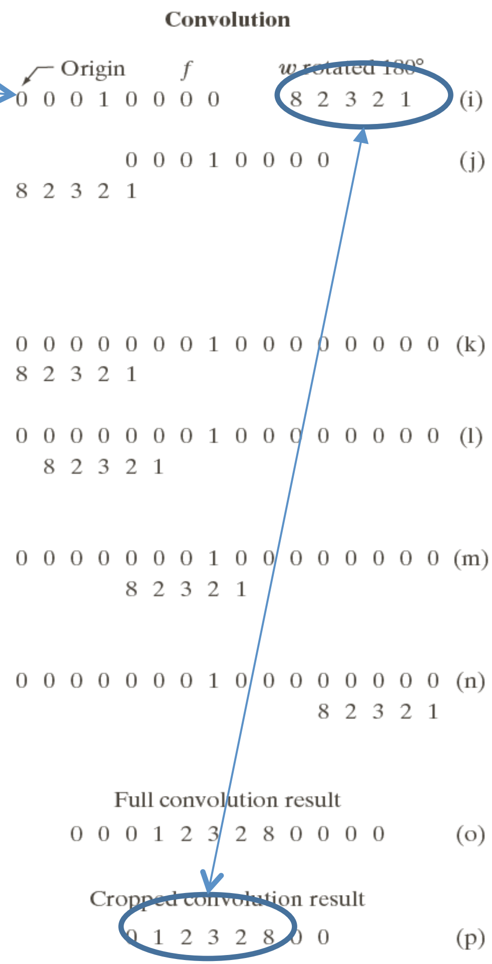

Flipped filter

1D convolution $$\sum_{s=-a}^{a} w(s) f(x-s)$$

The impulse response is the same as the filter

Extend to 2D Image: 2D Image Correlation

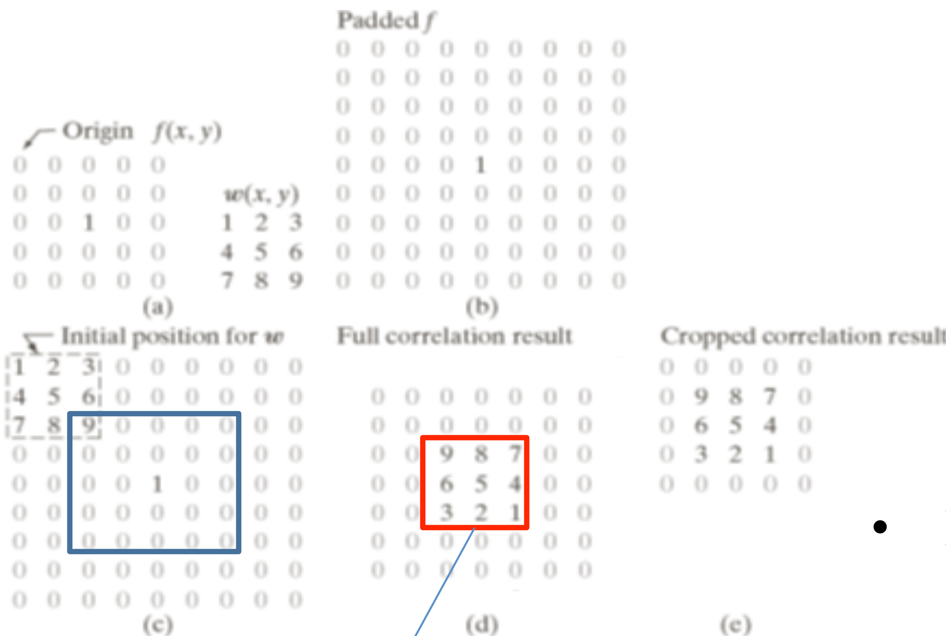

$$\sum_{s=-a}^{a}\sum_{t=-b}^{b} w(s,t)\, f(x+s, y+t)$$

- Full correlation result has the size of $(M+2a, N+2b)$
- Cropped result has the size of (M,N) – the size of the original image

The 2D impulse response of image correlation is a rotation of the filter by 180 degree

Extend to 2D Image: 2D Image Convolution

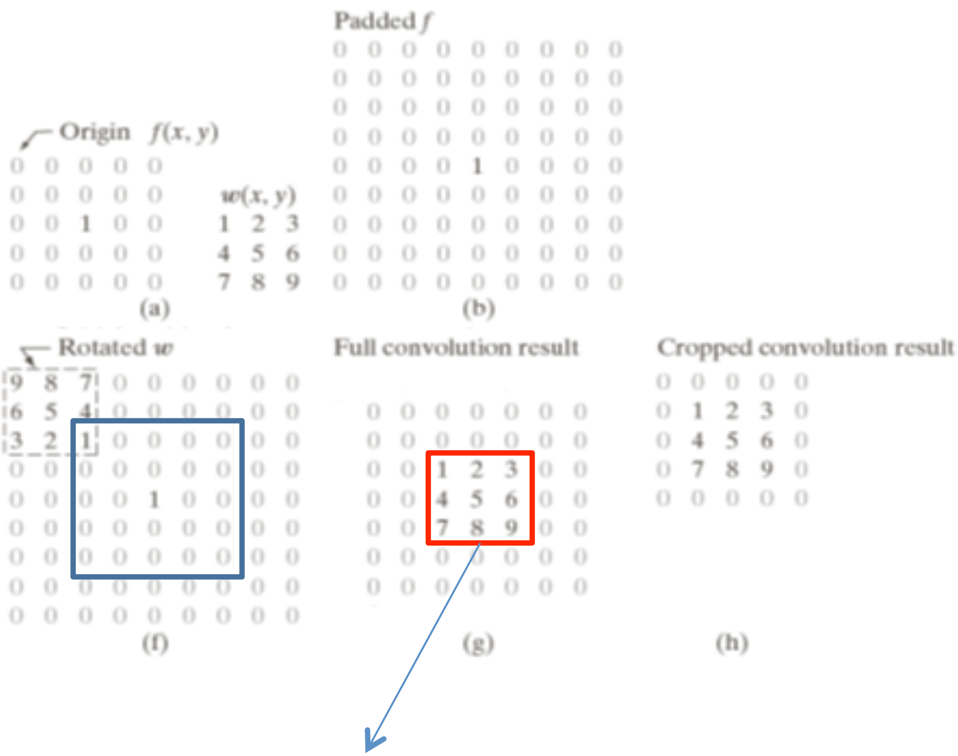

The 2D impulse response of image convolution is the same as the filter

$$\sum_{s=-a}^{a}\sum_{t=-b}^{b} w(s,t) f(x-s, y-t)$$

- Flip in both horizontal and vertical directions (rotate 180 degree) -> same if the filter is symmetric
- Convolution filter/mask/kernel
- Full convolution result has the size of $(M+2a, N+2b)$
- Cropped result has the size of (M,N) – the size of the original image

Slides courtesy of Prof. Yan Tong

Linear Filters

- General process:
  - Form new image whose pixels are a weighted sum of original pixel values, using the same set of weights at each point.
- Properties
  - Output is a linear function of the input
  - Output is a shift-invariant function of the input (i.e. shift the input image two pixels to the left, the output is shifted two pixels to the left)
- Example: smoothing by averaging
  - form the average of pixels in a neighborhood
- Example: smoothing with a Gaussian
  - form a weighted average of pixels in a neighborhood
- Example: finding an edge

Smoothing Spatial Filter – Low Pass Filters

Weighted average

$\frac{1}{9} \times$

1	1	1
1	1	1
1	1	1

$\frac{1}{16} \times$

1	2	1
2	4	2
1	2	1

a b

FIGURE 3.32 Two 3×3 smoothing (averaging) filter masks. The constant multiplier in front of each mask is equal to 1 divided by the sum of the values of its coefficients, as is required to compute an average.

$$g(x,y) = \frac{\sum_{s=-a}^{a}\sum_{t=-b}^{b} w(s,t) f(x+s, y+t)}{\sum_{s=-a}^{a}\sum_{t=-b}^{b} w(s,t)}$$

- Noise reduction
- reduction of "irrelevant details"
- edge blurred

Normalization factor

Slides courtesy of Prof. Yan Tong

Smoothing Spatial Filter

Image averaging $$R = \frac{1}{9}\sum_{i=1}^{9} z_i$$

1/9	1/9	1/9
1/9	1/9	1/9
1/9	1/9	1/9

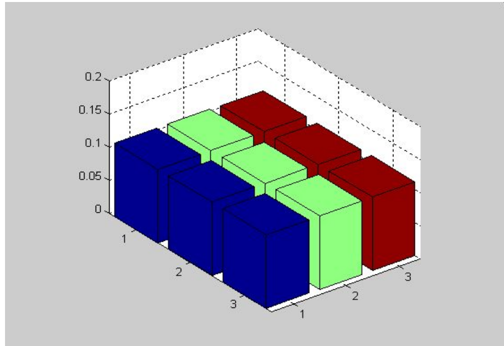

$$* \frac{1}{9}$$

1	1	1
1	1	1
1	1	1

=

Smoothing Spatial Filter

2D Gaussian filter

$$h(x, y) = \frac{1}{2\pi\sigma^2} e^{-\frac{x^2+y^2}{2\sigma^2}}$$

$*$

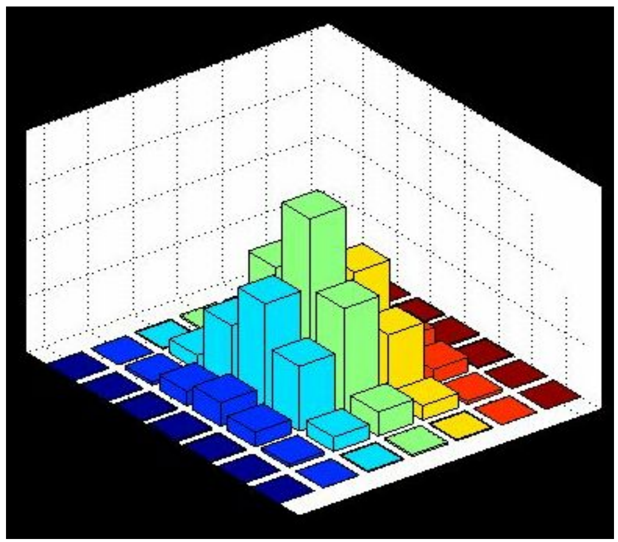

$=$

Slides courtesy of Prof. Yan Tong

Comparison using Different Smoothing Filters – Different Kernels

Average

Gaussian

Slides courtesy of Prof. Yan Tong

Comparison using Different Smoothing Filters: Different Size

FIGURE 3.33 (a) Original image, of size 500 × 500 pixels. (b)–(f) Results of smoothing with square averaging filter masks of sizes m = 3, 5, 9, 15, and 35, respectively. The black squares at the top are of sizes 3, 5, 9, 15, 25, 35, 45, and 55 pixels, respectively; their borders are 25 pixels apart. The letters at the bottom range in size from 10 to 24 points, in increments of 2 points; the large letter at the top is 60 points. The vertical bars are 5 pixels wide and 100 pixels high; their separation is 20 pixels. The diameter of the circles is 25 pixels, and their borders are 15 pixels apart; their intensity levels range from 0% to 100% black in increments of 20%. The background of the image is 10% black. The noisy rectangles are of size 50 × 120 pixels.

a b
c d
e f

Image Smoothing and Thresholding

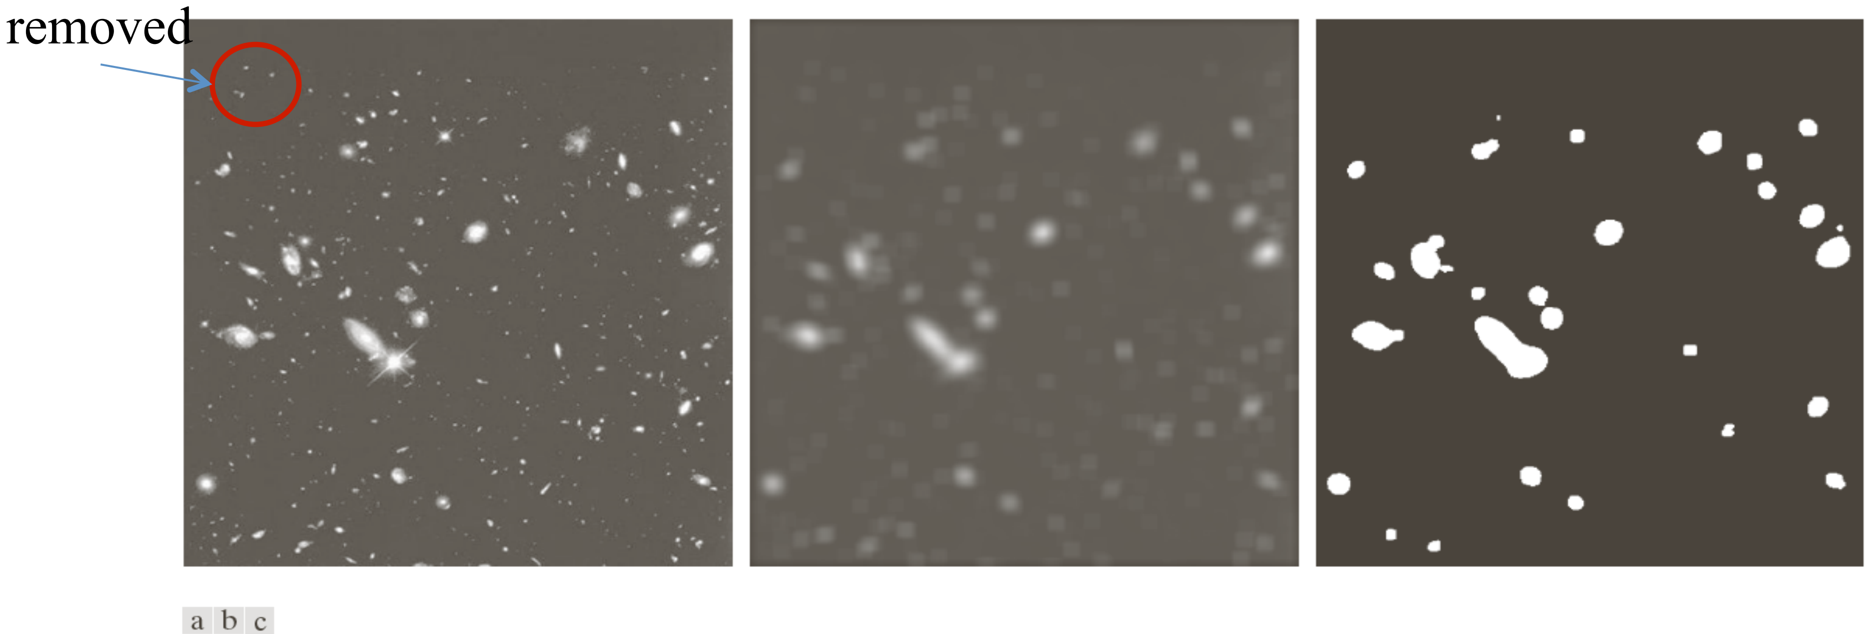

FIGURE 3.34 (a) Image of size 528 × 485 pixels from the Hubble Space Telescope. (b) Image filtered with a 15 × 15 averaging mask. (c) Result of thresholding (b). (Original image courtesy of NASA.)

Sharpening Spatial Filters

smooth → sharpen

http://www.bythom.com/sharpening.htm

Sharpening – highlight the transitions in intensity by differentiation

⇕

Smoothing – blur the transitions by summation

Sharpening Spatial Filters

Sharpening – highlight the transitions in intensity by differentiation

- Electric printing
- Medical imaging
- Industrial inspection

Compared to smoothing – blur the transitions by summation

Slides courtesy of Prof. Yan Tong

Perceived Intensity is Not a Simple Function of the Actual Intensity (1)

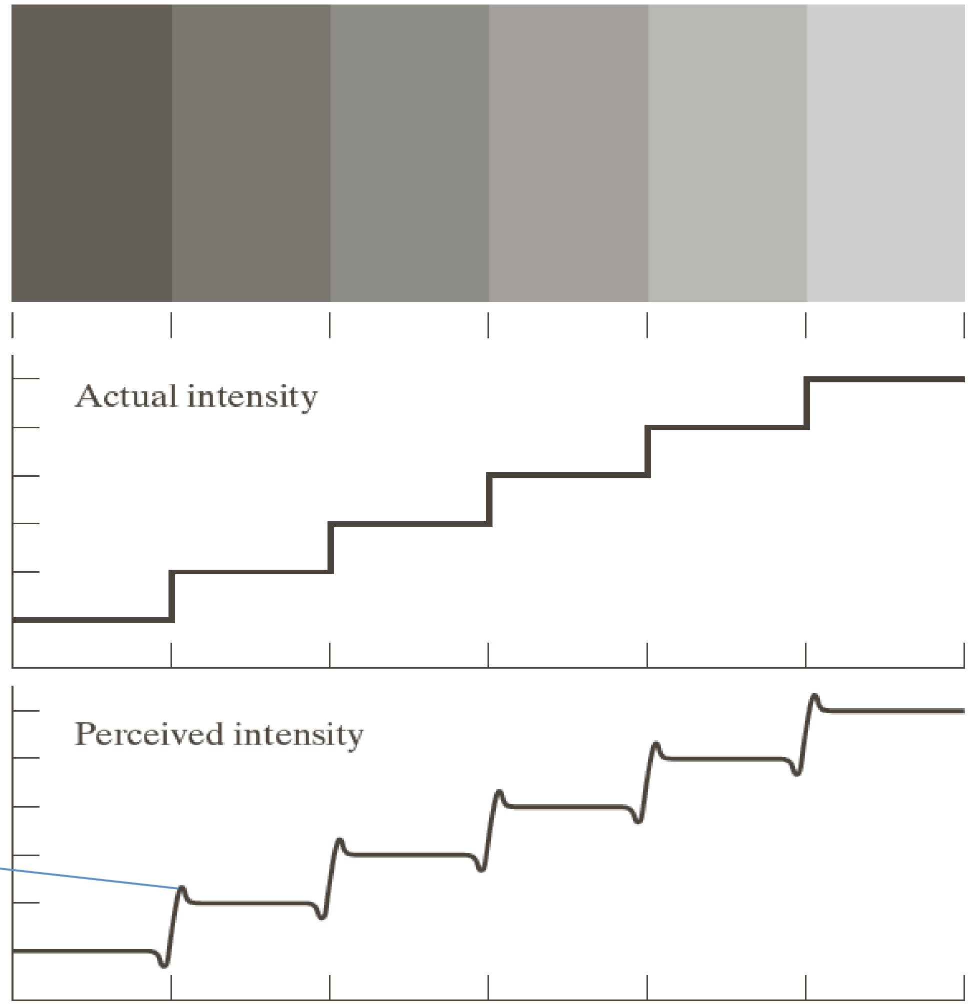

a
b
c

FIGURE 2.7 Illustration of the Mach band effect. Perceived intensity is not a simple function of actual intensity.

Enhance/amplify difference by image sharpening

Sharpening Spatial Filters

Original image

$$g(x,y) = f(x,y) + c * e(x,y)$$

Sharpened image — Magnifying factor — Edge map

Slides courtesy of Prof. Yan Tong

Spatial Filters for Edge Detection

$$\frac{\partial f}{\partial x} = f(x+1) - f(x)$$

Nonzero at

- onset of ramp and step
- along ramp or step

$$\frac{\partial^2 f}{\partial x^2} = f(x+1) + f(x-1) - 2f(x)$$

Nonzero at

- onset of ramp and step
- end of ramp and step

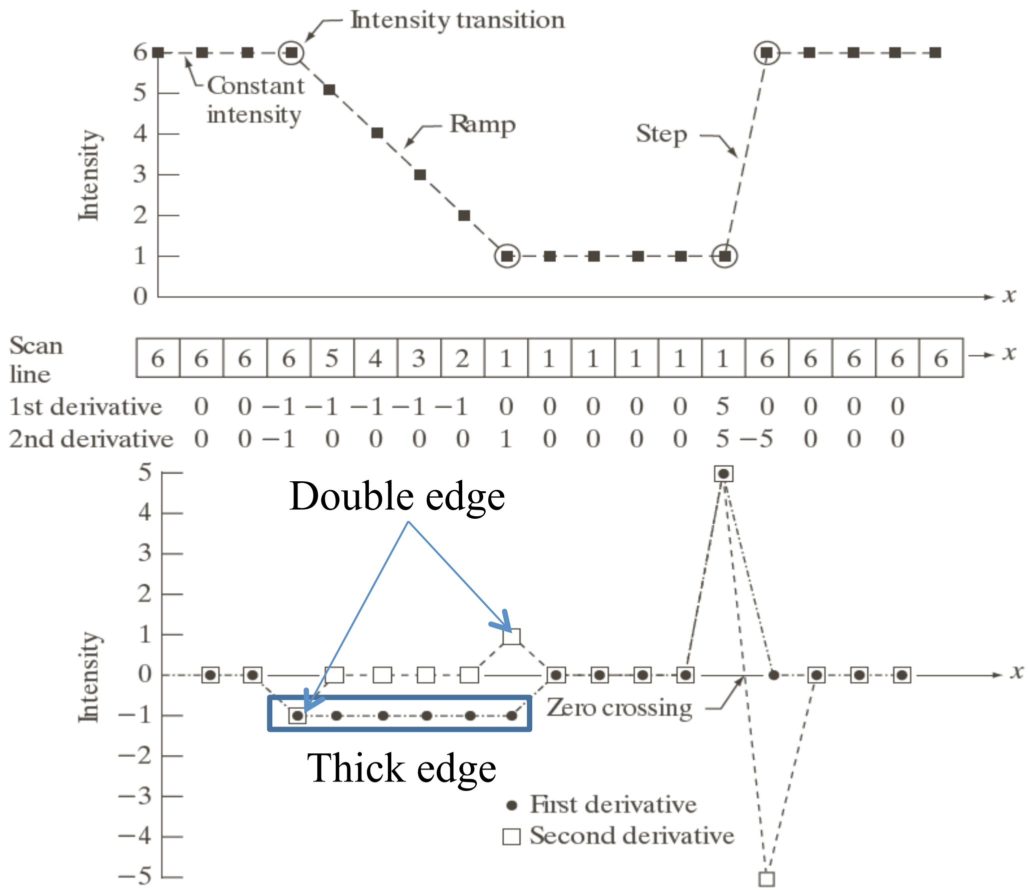

a
b
c

FIGURE 3.36 Illustration of the first and second derivatives of a 1-D digital function representing a section of a horizontal intensity profile from an image. In (a) and (c) data points are joined by dashed lines as a visualization aid.

Slides courtesy of Prof. Yan Tong

derivative

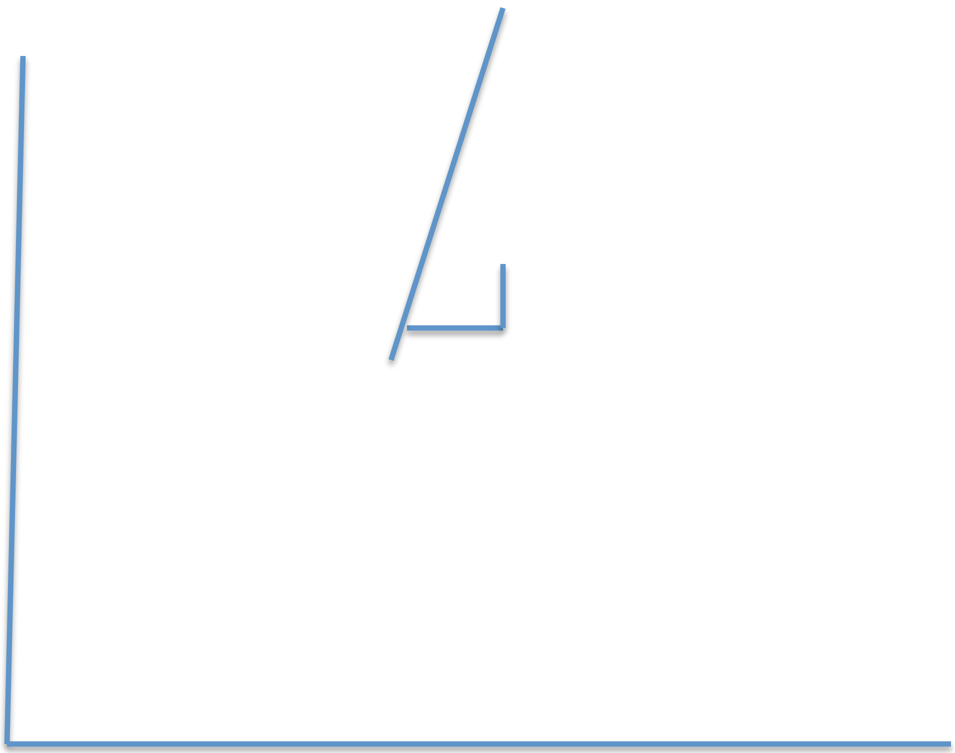

$$f(x) -$$

$$f(x1) - f(x2)$$

$$\frac{f(x1) - f(x2)}{x1 - x2}$$

First-order VS Second-order Derivative for Edge Detection

- First-order derivative produces thick edge along the direction of transition
- Second-order derivative produces thinner edges

Gradient for Image Sharpening

Direction of change

$$\nabla f = \text{grad}(f) = \begin{bmatrix} g_x \\ g_y \end{bmatrix} = \begin{bmatrix} \dfrac{\partial f}{\partial x} \\ \\ \dfrac{\partial f}{\partial x} \end{bmatrix}$$

Magnitude of change (gradient image)

$$M(x, y) = mag(\nabla f) = \sqrt{g_x^2 + g_y^2}$$

$$M(x, y) \approx |g_x| + |g_y|$$

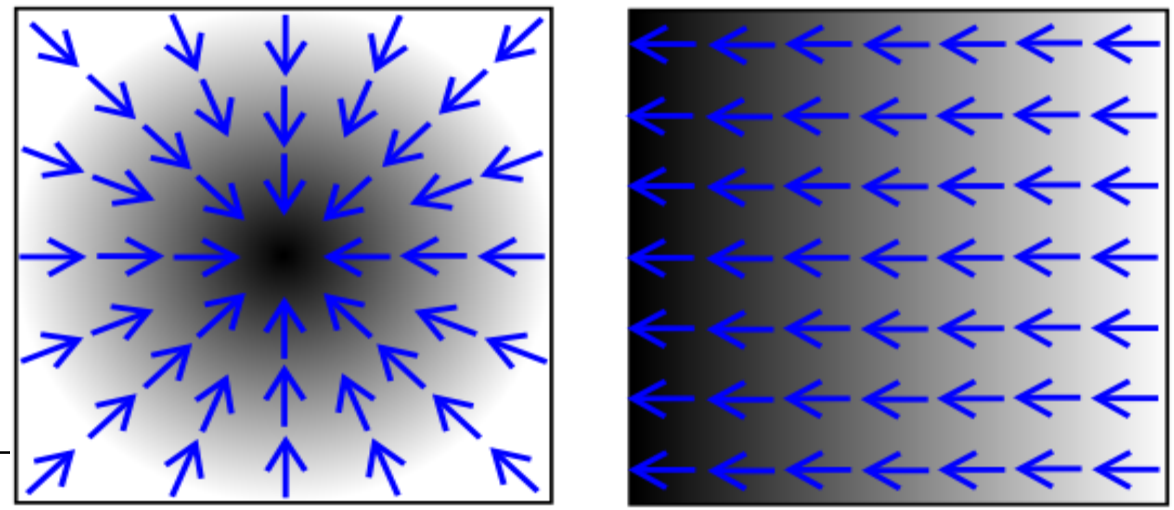

http://en.wikipedia.org/wiki/Image_gradient

Slides courtesy of Prof. Yan Tong

Gradient for Image Sharpening

Sum of the coefficients is 0 – the response of a constant region is 0

Edge detectors:

- Roberts cross – fast while sensitive to noise
- Sobel - smooth

z_1	z_2	z_3
z_4	z_5	z_6
z_7	z_8	z_9

−1	0
0	1

0	−1
1	0

−1	−2	−1
0	0	0
1	2	1

−1	0	1
−2	0	2
−1	0	1

a
b c
d e

FIGURE 3.41 A 3 × 3 region of an image (the zs are intensity values). (b)–(c) Roberts cross gradient operators. (d)–(e) Sobel operators. All the mask coefficients sum to zero, as expected of a derivative operator.

Laplacian for Image Sharpening

2D Isotropic filters – rotation invariant

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

Laplacian operator

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

$$g(x,y) = f(x,y) + c\left[\nabla^2 f(x,y)\right]$$

Image sharpening with Laplacian

(a)

0	1	0
1	−4	1
0	1	0

(b)

1	1	1
1	−8	1
1	1	1

(c)

0	−1	0
−1	4	−1
0	−1	0

(d)

−1	−1	−1
−1	8	−1
−1	−1	−1

a b
c d

FIGURE 3.37 (a) Filter mask used to implement Eq. (3.6-6). (b) Mask used to implement an extension of this equation that includes the diagonal terms. (c) and (d) Two other implementations of the Laplacian found frequently in practice.

Image Sharpening

Scale the Laplacian by shifting the intensity range to [0, L-1]

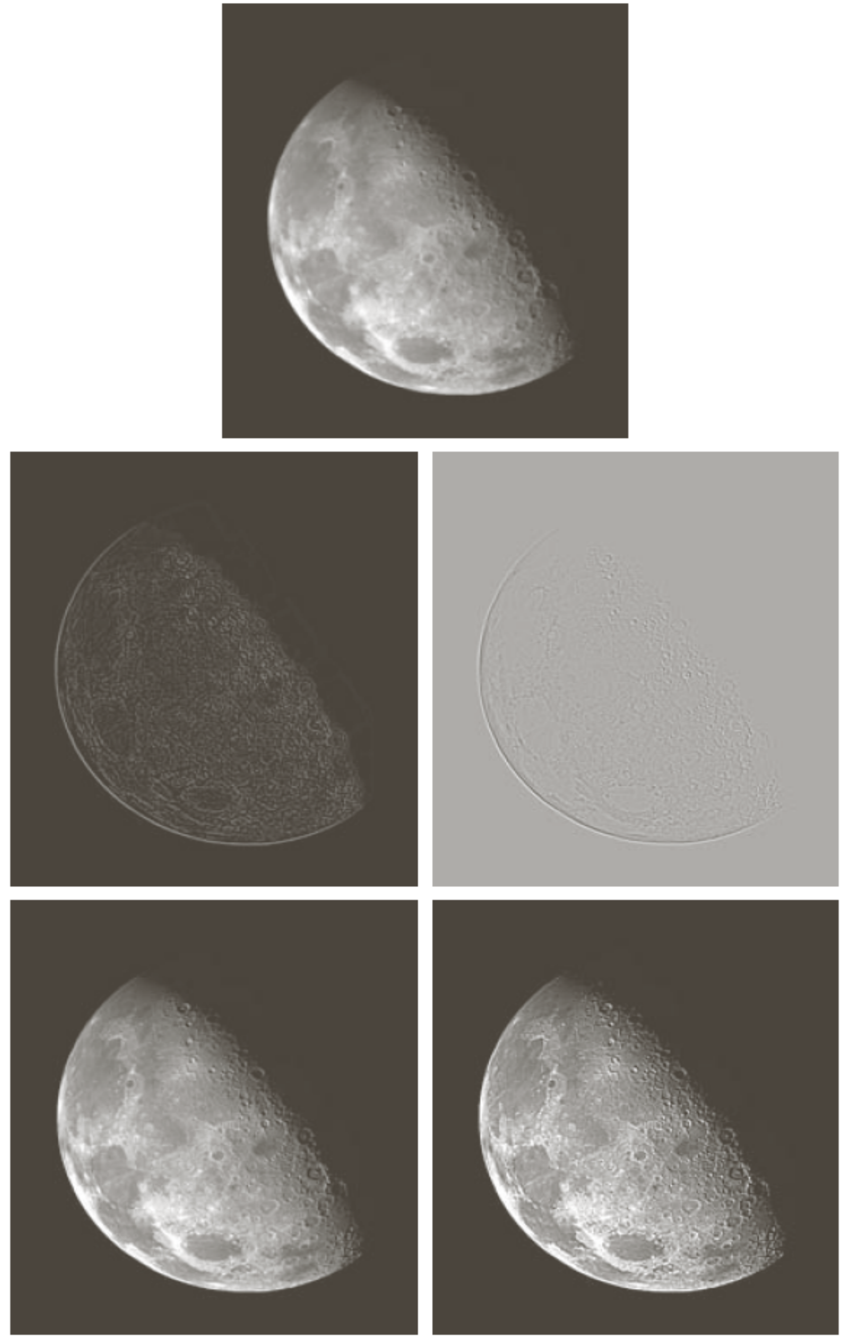

a
b c
d e

FIGURE 3.38 (a) Blurred image of the North Pole of the moon. (b) Laplacian without scaling. (c) Laplacian with scaling. (d) Image sharpened using the mask in Fig. 3.37(a). (e) Result of using the mask in Fig. 3.37(b). (Original image courtesy of NASA.)

Image Sharpening by Unsharp Masking and Highboost Filtering

1. **Blur the original image**
2. **Subtract the blurred image from the original to get the mask**
3. **Add the mask to the original**

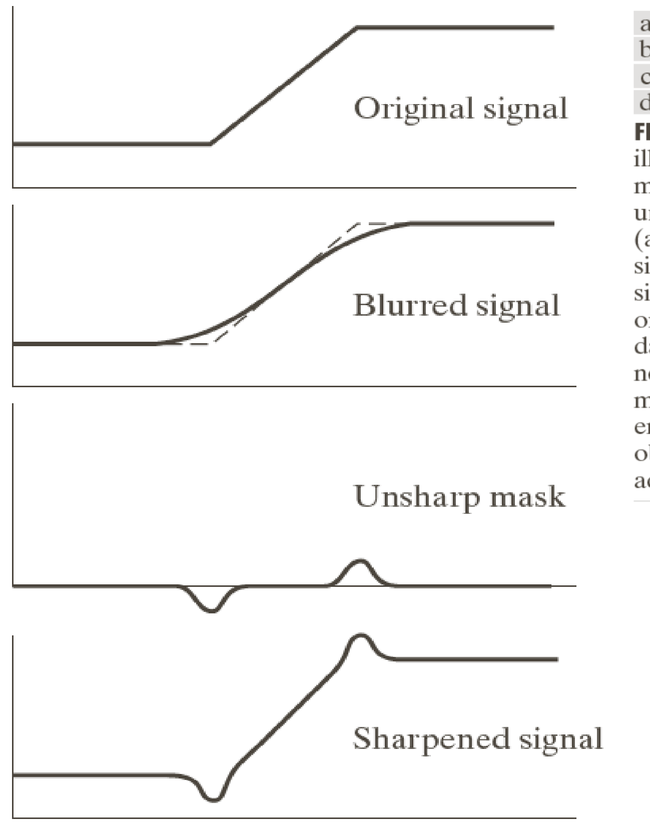

FIGURE 3.39 1-D illustration of the mechanics of unsharp masking. (a) Original signal. (b) Blurred signal with original shown dashed for reference. (c) Unsharp mask. (d) Sharpened signal, obtained by adding (c) to (a).

$$g(x,y) = f(x,y) + k * \left(f(x,y) - \bar{f}(x,y)\right)$$

$$k \geq 0$$

When k>1, it becomes a highboost filtering.

Slides courtesy of Prof. Yan Tong

An Example

a
b
c
d
e

FIGURE 3.40 (a) Original image. (b) Result of blurring with a Gaussian filter. (c) Unsharp mask. (d) Result of using unsharp masking. (e) Result of using highboost filtering.

Gradient for Image Sharpening -- Example

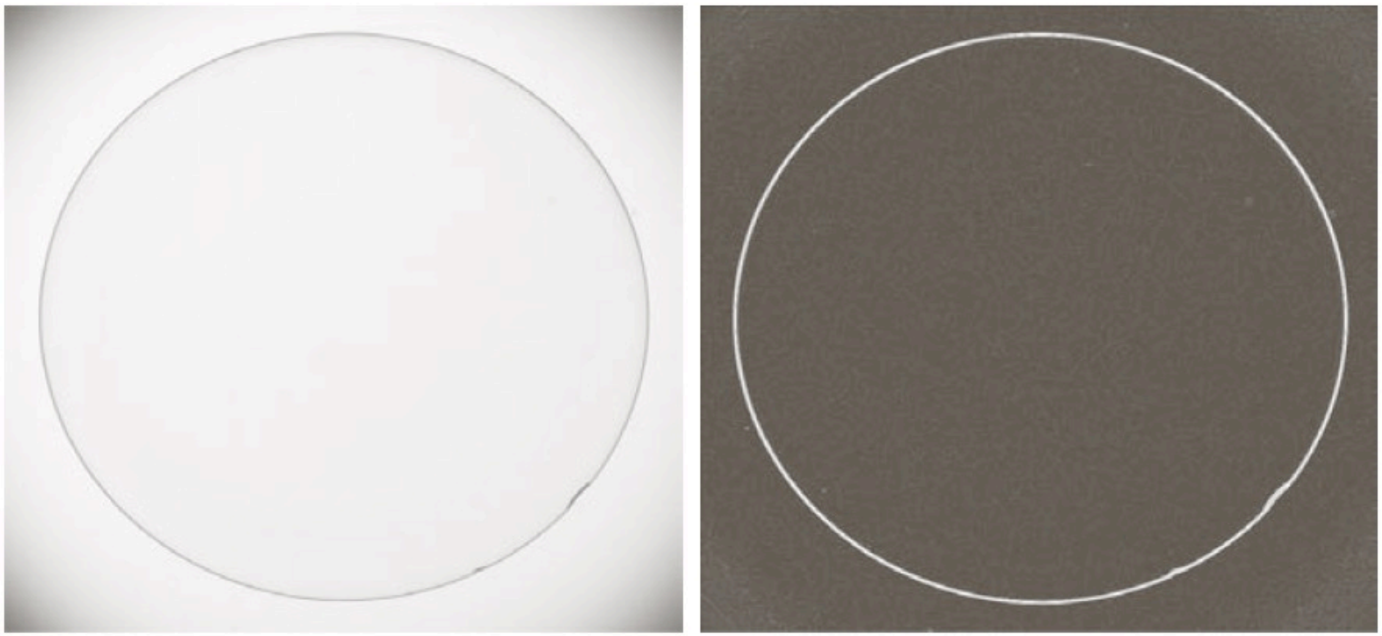

An application in industrial defect detection.

Combining Spatial Enhancement Methods

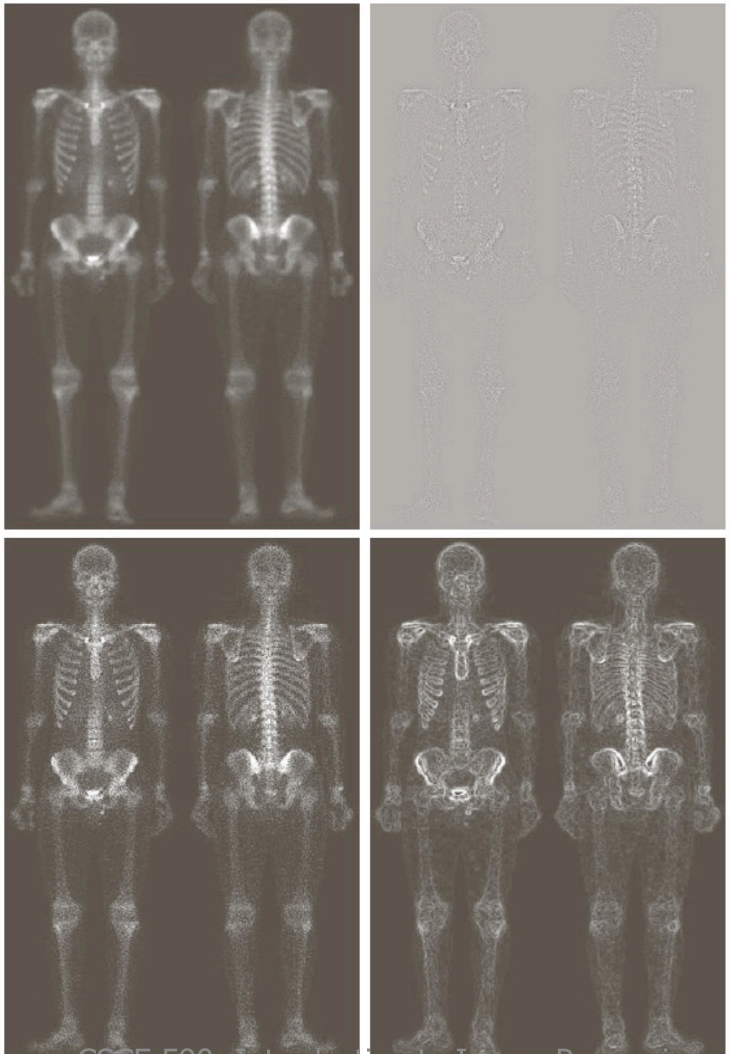

a b
c d

FIGURE 3.43
(a) Image of whole body bone scan.
(b) Laplacian of (a). (c) Sharpened image obtained by adding (a) and (b).
(d) Sobel gradient of (a).

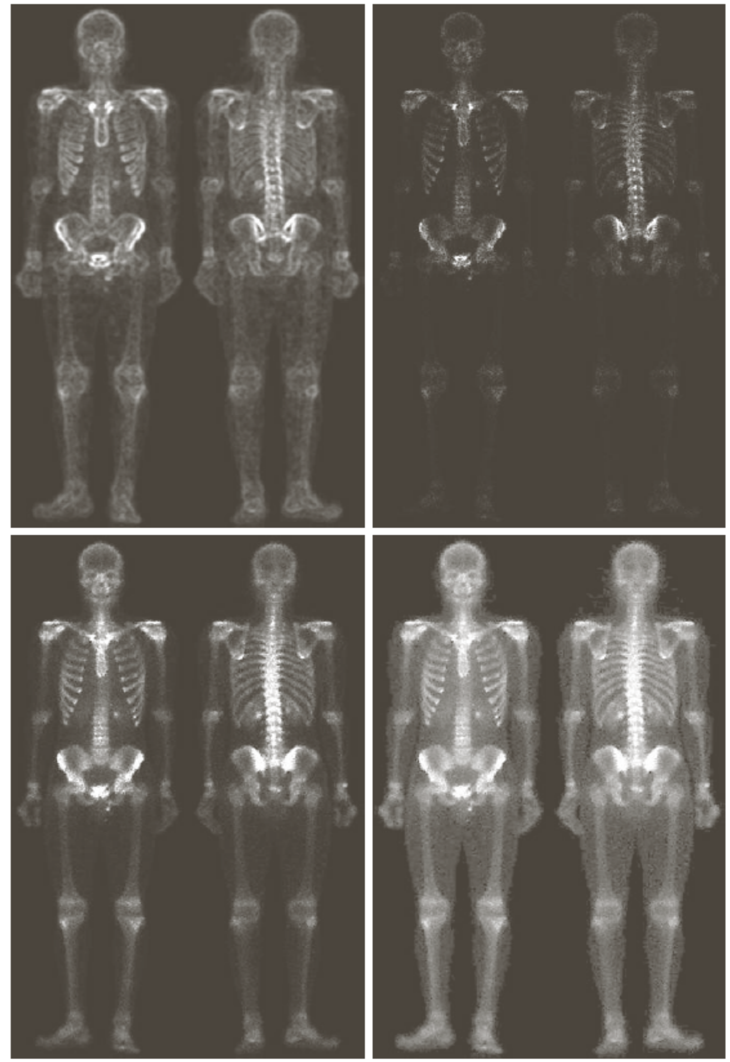

e f
g h

FIGURE 3.43 *(Continued)*
(e) Sobel image smoothed with a 5×5 averaging filter. (f) Mask image formed by the product of (c) and (e).
(g) Sharpened image obtained by the sum of (a) and (f). (h) Final result obtained by applying a power-law transformation to (g). Compare (g) and (h) with (a). (Original image courtesy of G.E. Medical Systems.)

Slides courtesy of Prof. Yan Tong

Order-Statistic (Nonlinear) Filtering

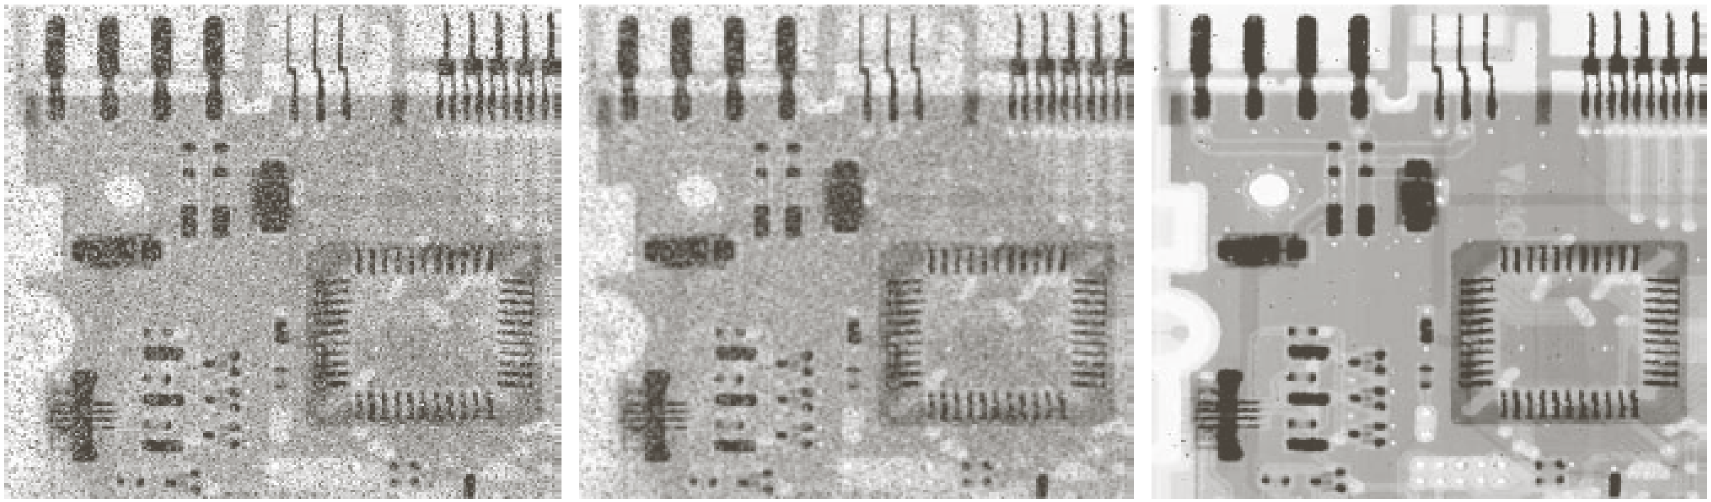

a b c

FIGURE 3.35 (a) X-ray image of circuit board corrupted by salt-and-pepper noise. (b) Noise reduction with a 3 × 3 averaging mask. (c) Noise reduction with a 3 × 3 median filter. (Original image courtesy of Mr. Joseph E. Pascente, Lixi, Inc.)

Order-statistic filtering – rank the pixel values in the filter window and assign the center pixel according to the property of the filter

- Median
- Min/max

Slides courtesy of Prof. Yan Tong

Questions?