BSB663
Image Processing

Pinar Duygulu

Slides are adapted from
Selim Aksoy

Importance of neighborhood

- Both zebras and dalmatians have black and white pixels in similar numbers.
- The difference between the two is the characteristic appearance of small group of pixels rather than individual pixel values.

Adapted from Pinar Duygulu, Bilkent University

Outline

- We will discuss neighborhood operations that work with the values of the image pixels in the neighborhood.

- Spatial domain filtering
- Frequency domain filtering
- Image enhancement
- Finding patterns

Spatial domain filtering

3	3	3
3	3	3
3	3	3

- What is the value of the center pixel?

3	4	3
2	3	3
3	4	2

- What assumptions are you making to infer the center value?

Spatial domain filtering

- Some neighborhood operations work with
  - the values of the image pixels in the neighborhood, and
  - the corresponding values of a subimage that has the same dimensions as the neighborhood.
- The subimage is called a filter (or mask, kernel, template, window).
- The values in a filter subimage are referred to as coefficients, rather than pixels.

Spatial domain filtering

- Operation: modify the pixels in an image based on some function of the pixels in their neighborhood.
- Simplest: linear filtering (replace each pixel by a linear combination of its neighbors).
- Linear spatial filtering is often referred to as “convolving an image with a filter”.

Linear filtering

For a linear spatially invariant system

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

m=0 1 2 …

111	115	113	111	112	111	112	111
135	138	137	139	145	146	149	147
163	168	188	196	206	202	206	207
180	184	206	219	202	200	195	193
189	193	214	216	104	79	83	77
191	201	217	220	103	59	60	68
195	205	216	222	113	68	69	83
199	203	223	228	108	68	71	77

g[m,n]

$\otimes$

-1	2	-1
-1	2	-1
-1	2	-1

h[m,n]

=

?	?	?	?	?	?	?	?
?	-5	9	-9	21	-12	10	?
?	-29	18	24	4	-7	5	?
?	-50	40	142	-88	-34	10	?
?	-41	41	264	-175	-71	0	?
?	-24	37	349	-224	-120	-10	?
?	-23	33	360	-217	-134	-23	?
?	?	?	?	?	?	?	?

f[m,n]

Spatial domain filtering

- Be careful about indices, image borders and padding during implementation.

zero

fixed/clamp

periodic/wrap

reflected/mirror

Border padding examples.

Adapted from CSE 455, U of Washington

Smoothing spatial filters

- Often, an image is composed of
  - some underlying ideal structure, which we want to detect and describe,
  - together with some random noise or artifact, which we would like to remove.
- Smoothing filters are used for blurring and for noise reduction.
- Linear smoothing filters are also called averaging filters.

Smoothing spatial filters

I

10	11	10	0	0	1
9	10	11	1	0	1
10	9	10	0	2	1
11	10	9	10	9	11
9	10	11	9	99	11
10	9	9	11	10	10

F

1/9

1	1	1
1	1	1
1	1	1

O

X	X	X	X	X	X
X	10				X
X					X
X					X
X					X
X	X	X	X	X	X

1/9.(10x1 + 11x1 + 10x1 + 9x1 + 10x1 + 11x1 + 10x1 + 9x1 + 10x1) =
1/9.(90) = 10

Adapted from Octavia Camps, Penn State

Smoothing spatial filters

I

10	11	10	0	0	1
9	10	11	1	0	1
10	9	10	0	2	1
11	10	9	10	9	11
9	10	11	9	99	11
10	9	9	11	10	10

F

1/9

1	1	1
1	1	1
1	1	1

O

X	X	X	X	X	X
X					X
X					X
X					X
X				20	X
X	X	X	X	X	X

1/9.(10x1 + 9x1 + 11x1 + 9x1 + 99x1 + 11x1 + 11x1 + 10x1 + 10x1) =
1/9.(180) = 20

Adapted from Octavia Camps, Penn State

Smoothing spatial filters

- Common types of noise:
  - Salt-and-pepper noise: contains random occurrences of black and white pixels.
  - Impulse noise: contains random occurrences of white pixels.
  - Gaussian noise: variations in intensity drawn from a Gaussian normal distribution.


Original


Salt and pepper noise


Impulse noise


Gaussian noise

Adapted from Linda Shapiro, U of Washington

Adapted from Linda Shapiro, U of Washington

Smoothing spatial filters

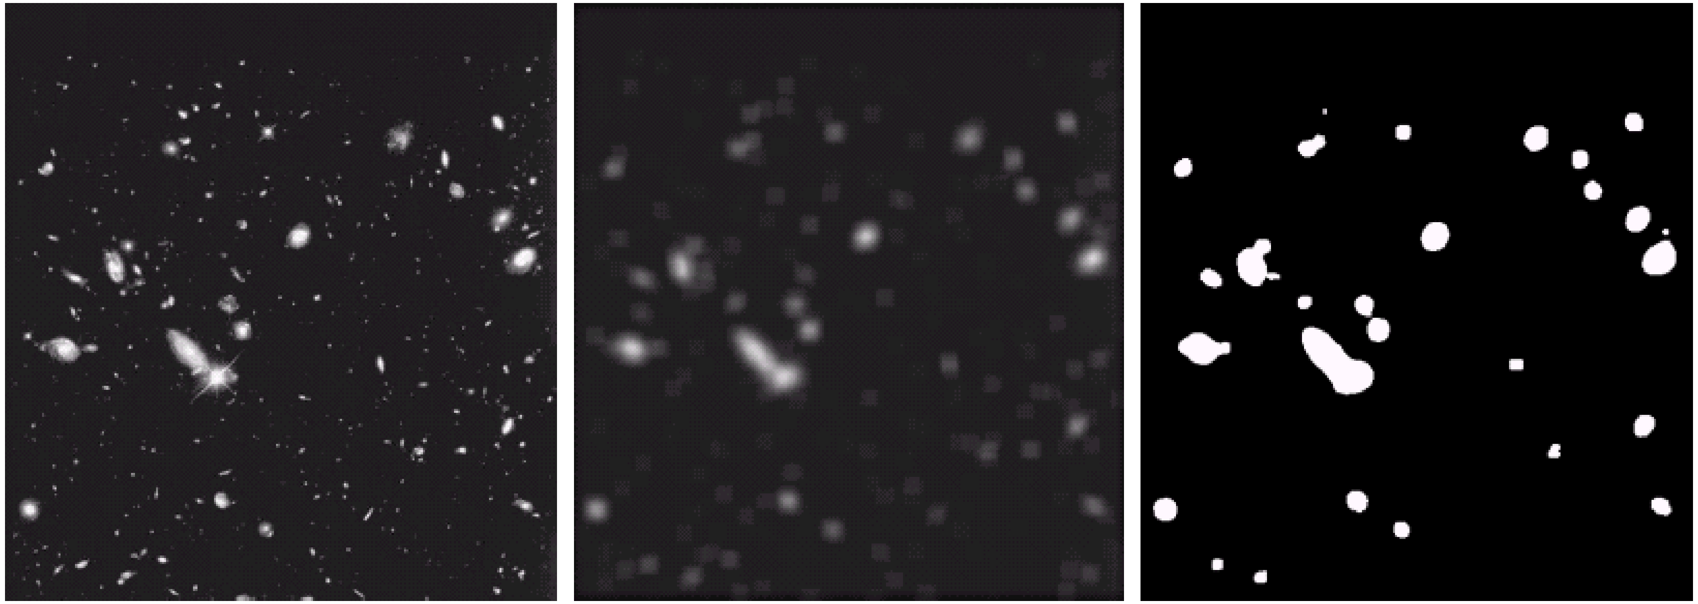

FIGURE 3.36 (a) Image from the Hubble Space Telescope. (b) Image processed by a 15 × 15 averaging mask. (c) Result of thresholding (b). (Original image courtesy of NASA.)

Adapted from Gonzales and Woods

Smoothing spatial filters

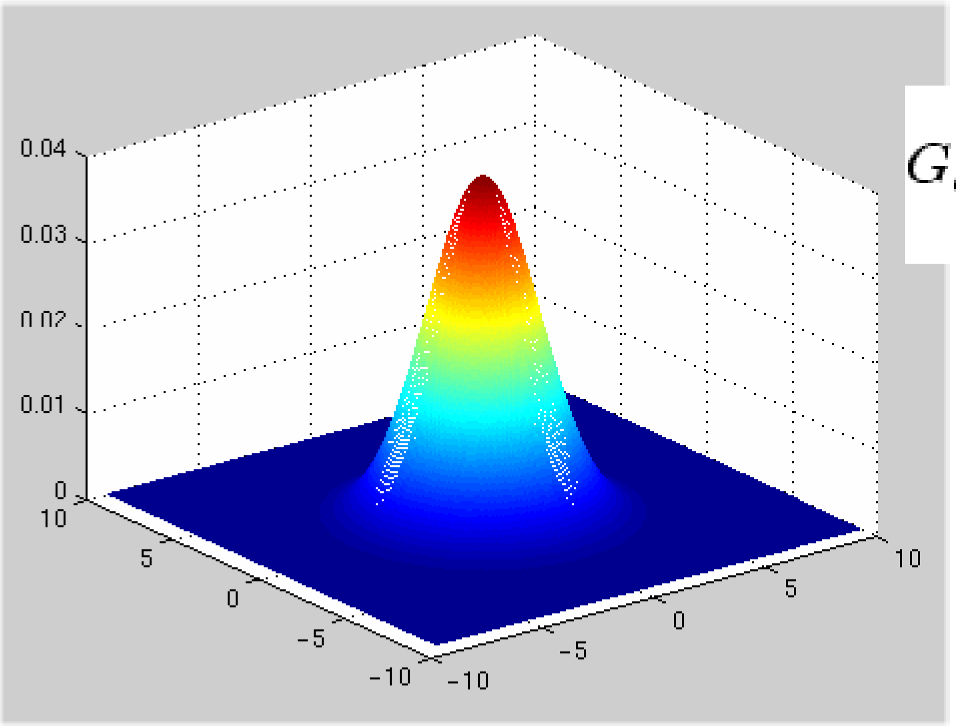

2D Gaussian filter

$$G_\sigma(x,y) = \frac{1}{2\pi\sigma^2} \exp\left(-\frac{(x^2+y^2)}{2\sigma^2}\right)$$

A weighted average that weighs pixels at its center much more strongly than its boundaries.

Adapted from Martial Hebert, CMU

Smoothing spatial filters

- If σ is small: smoothing will have little effect.
- If σ is larger: neighboring pixels will have larger weights resulting in consensus of the neighbors.
- If σ is very large: details will disappear along with the noise.

Effect of σ

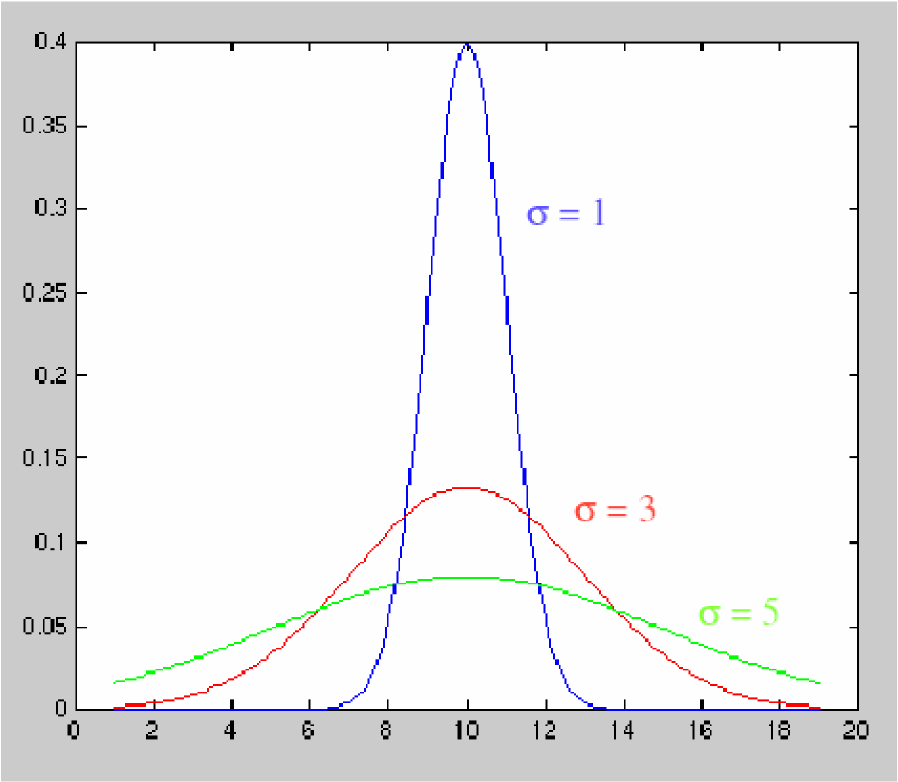

Adapted from Martial Hebert, CMU

Smoothing spatial filters

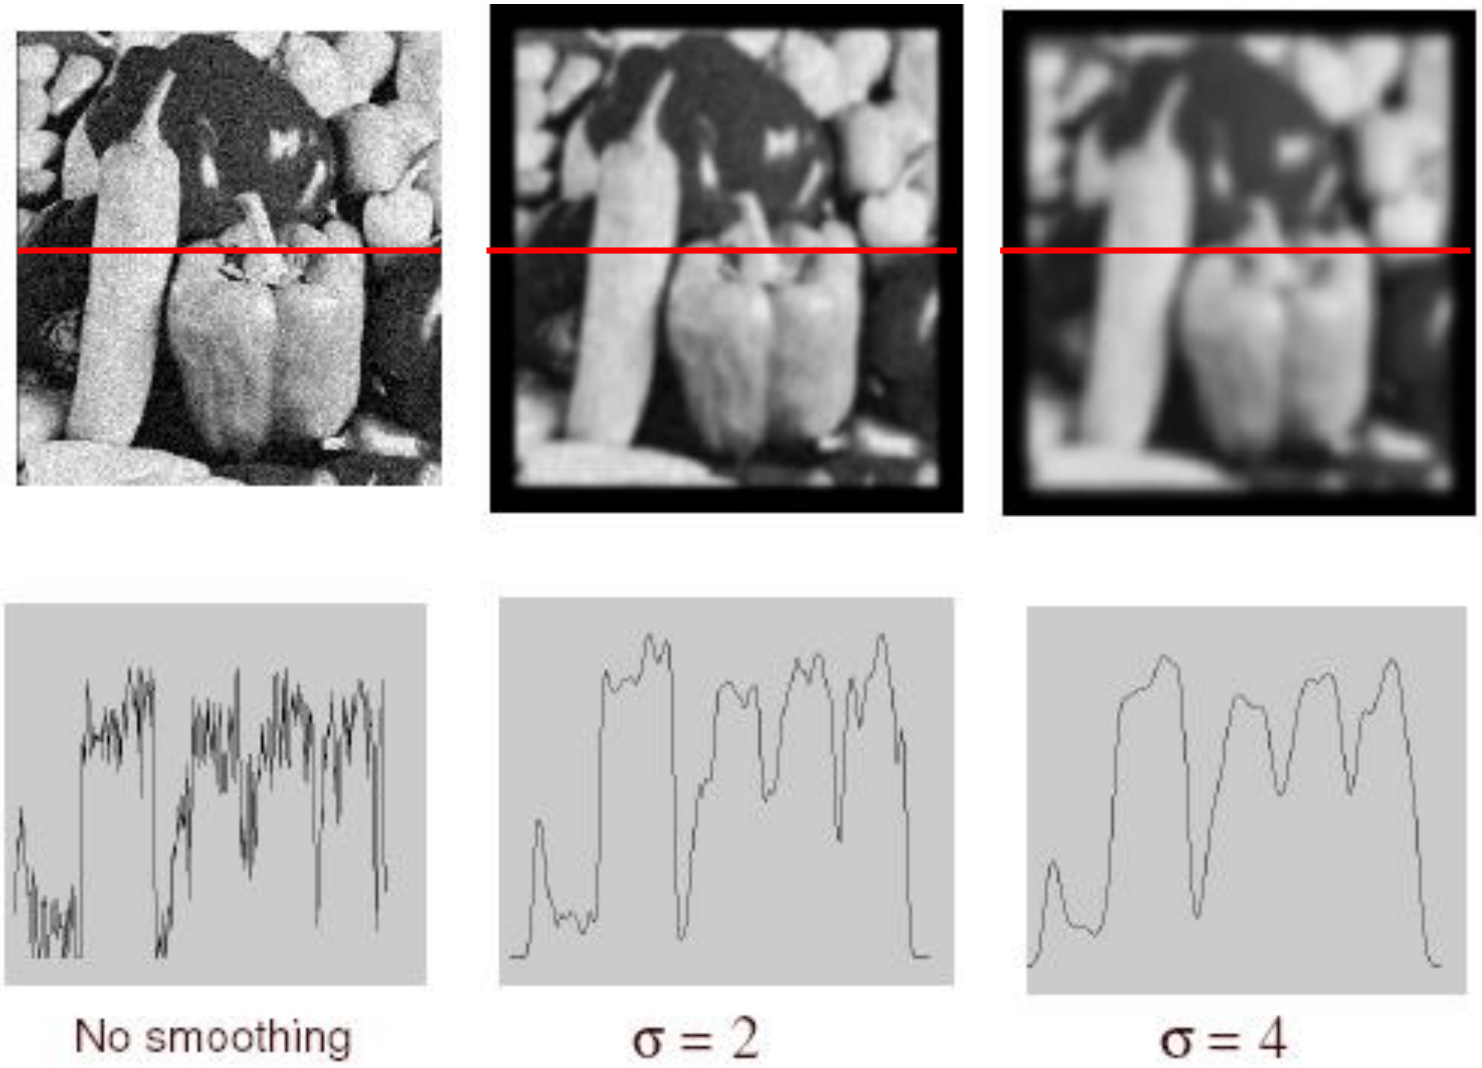

Adapted from Martial Hebert, CMU

Smoothing spatial filters

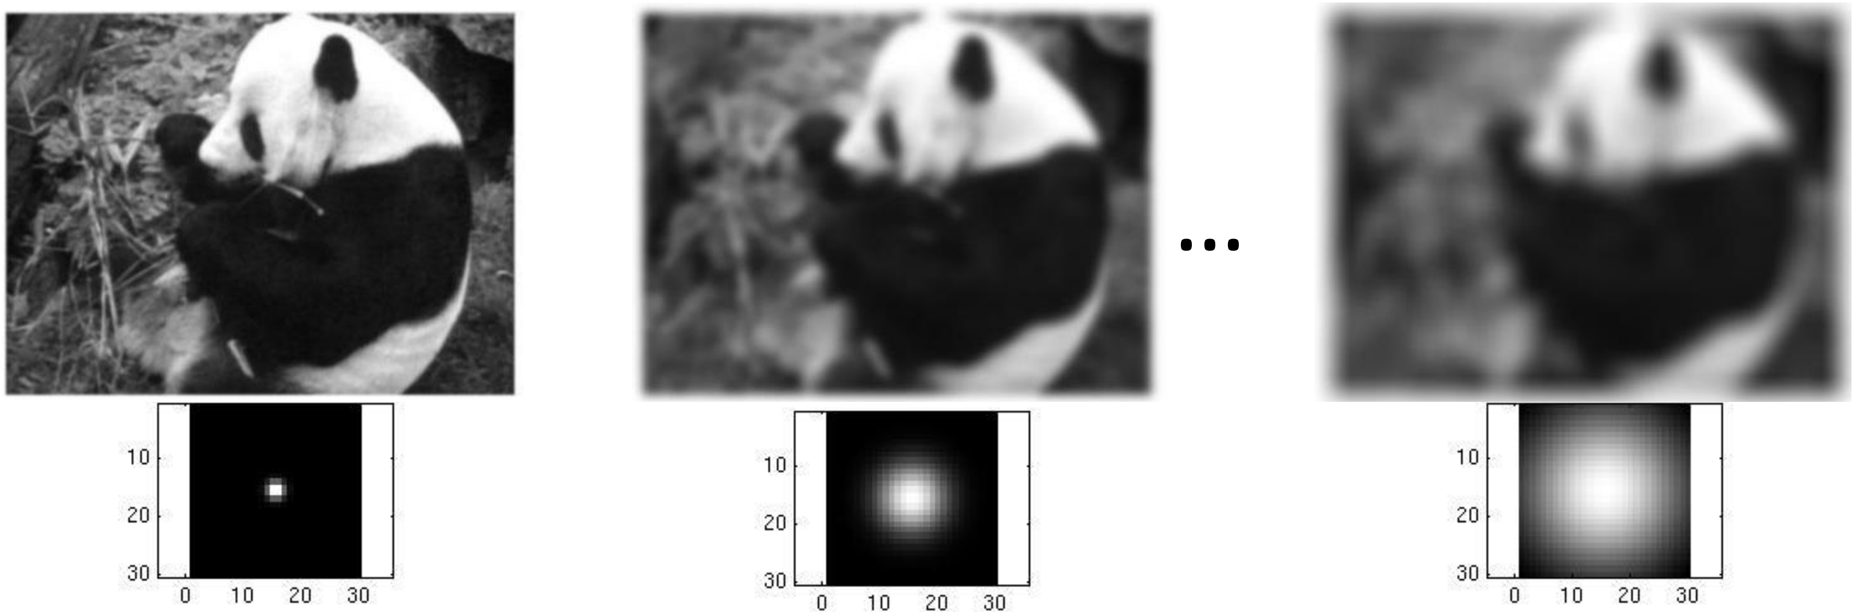

Width of the Gaussian kernel controls the amount of smoothing.

Adapted from K. Grauman

Smoothing spatial filters

Result of blurring using a uniform local model.

Produces a set of narrow horizontal and vertical bars – ringing effect.

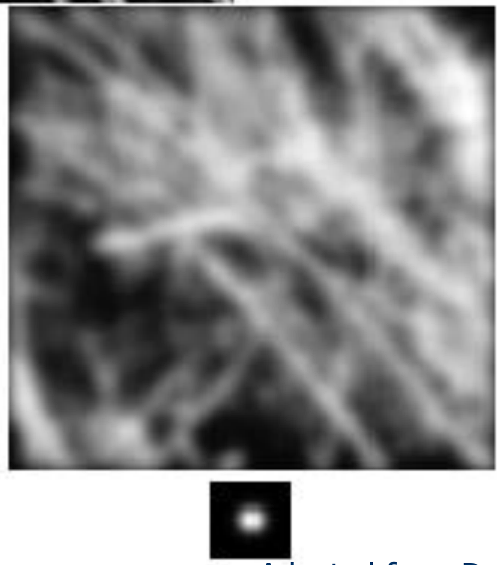

Result of blurring using a Gaussian filter.

Adapted from David Forsyth, UC Berkeley

Smoothing spatial filters

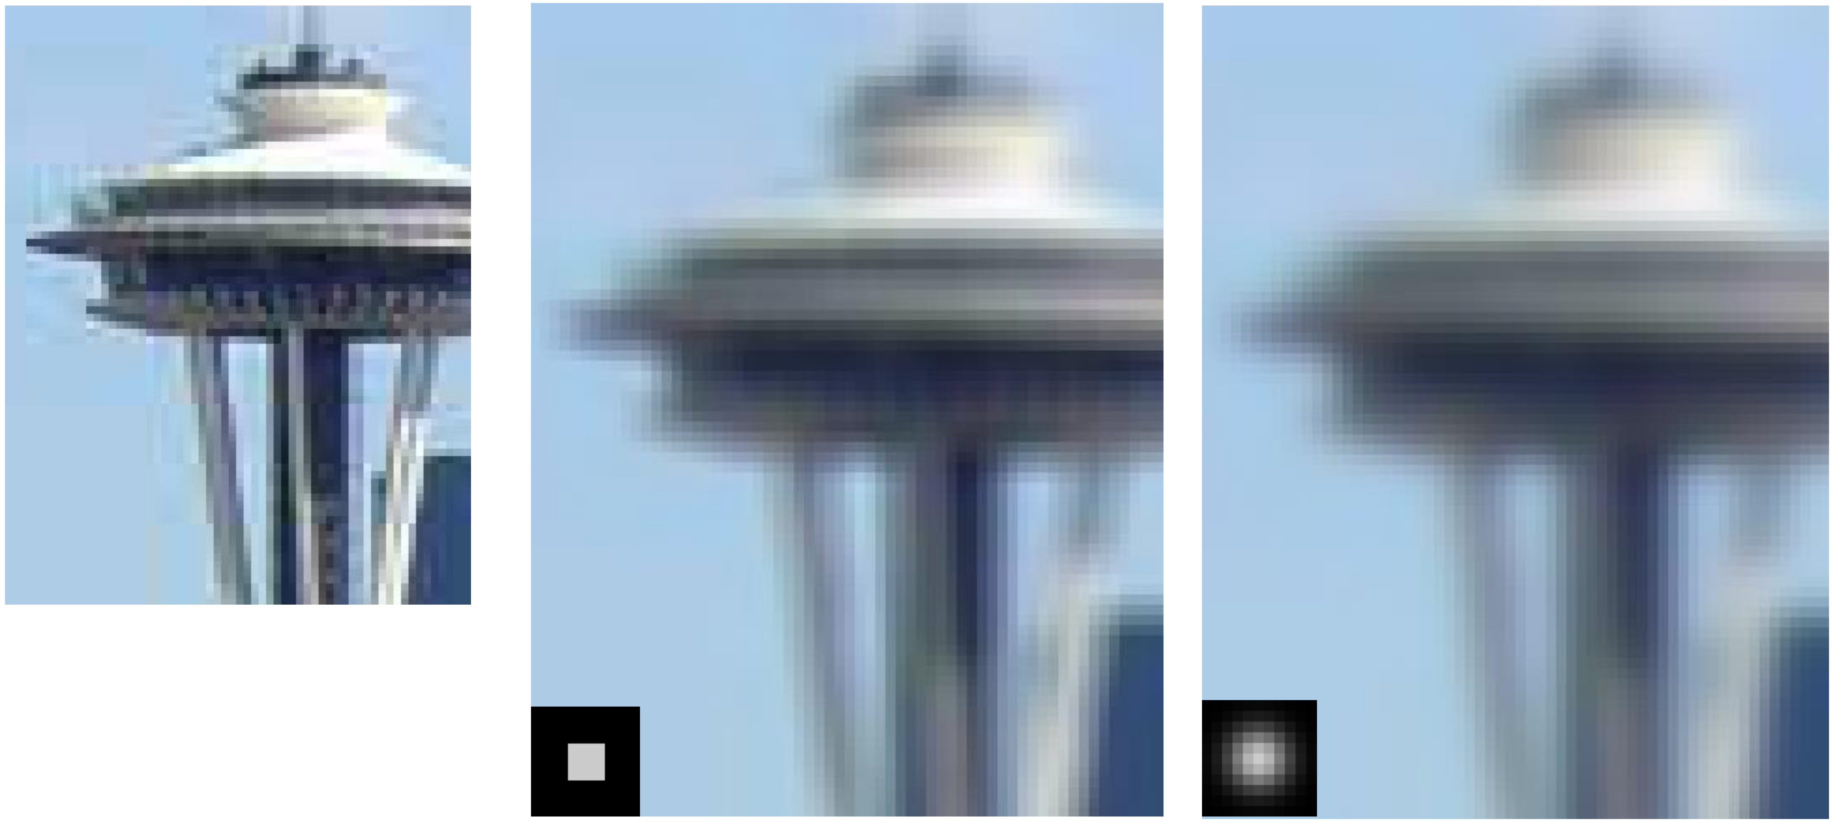

Gaussian versus mean filters

Adapted from CSE 455, U of Washington

Order-statistic filters

- Order-statistic filters are nonlinear spatial filters whose response is based on
  - ordering (ranking) the pixels contained in the image area encompassed by the filter, and then
  - replacing the value of the center pixel with the value determined by the ranking result.
- The best-known example is the median filter.
- It is particularly effective in the presence of impulse or salt-and-pepper noise, with considerably less blurring than linear smoothing filters.

Order-statistic filters

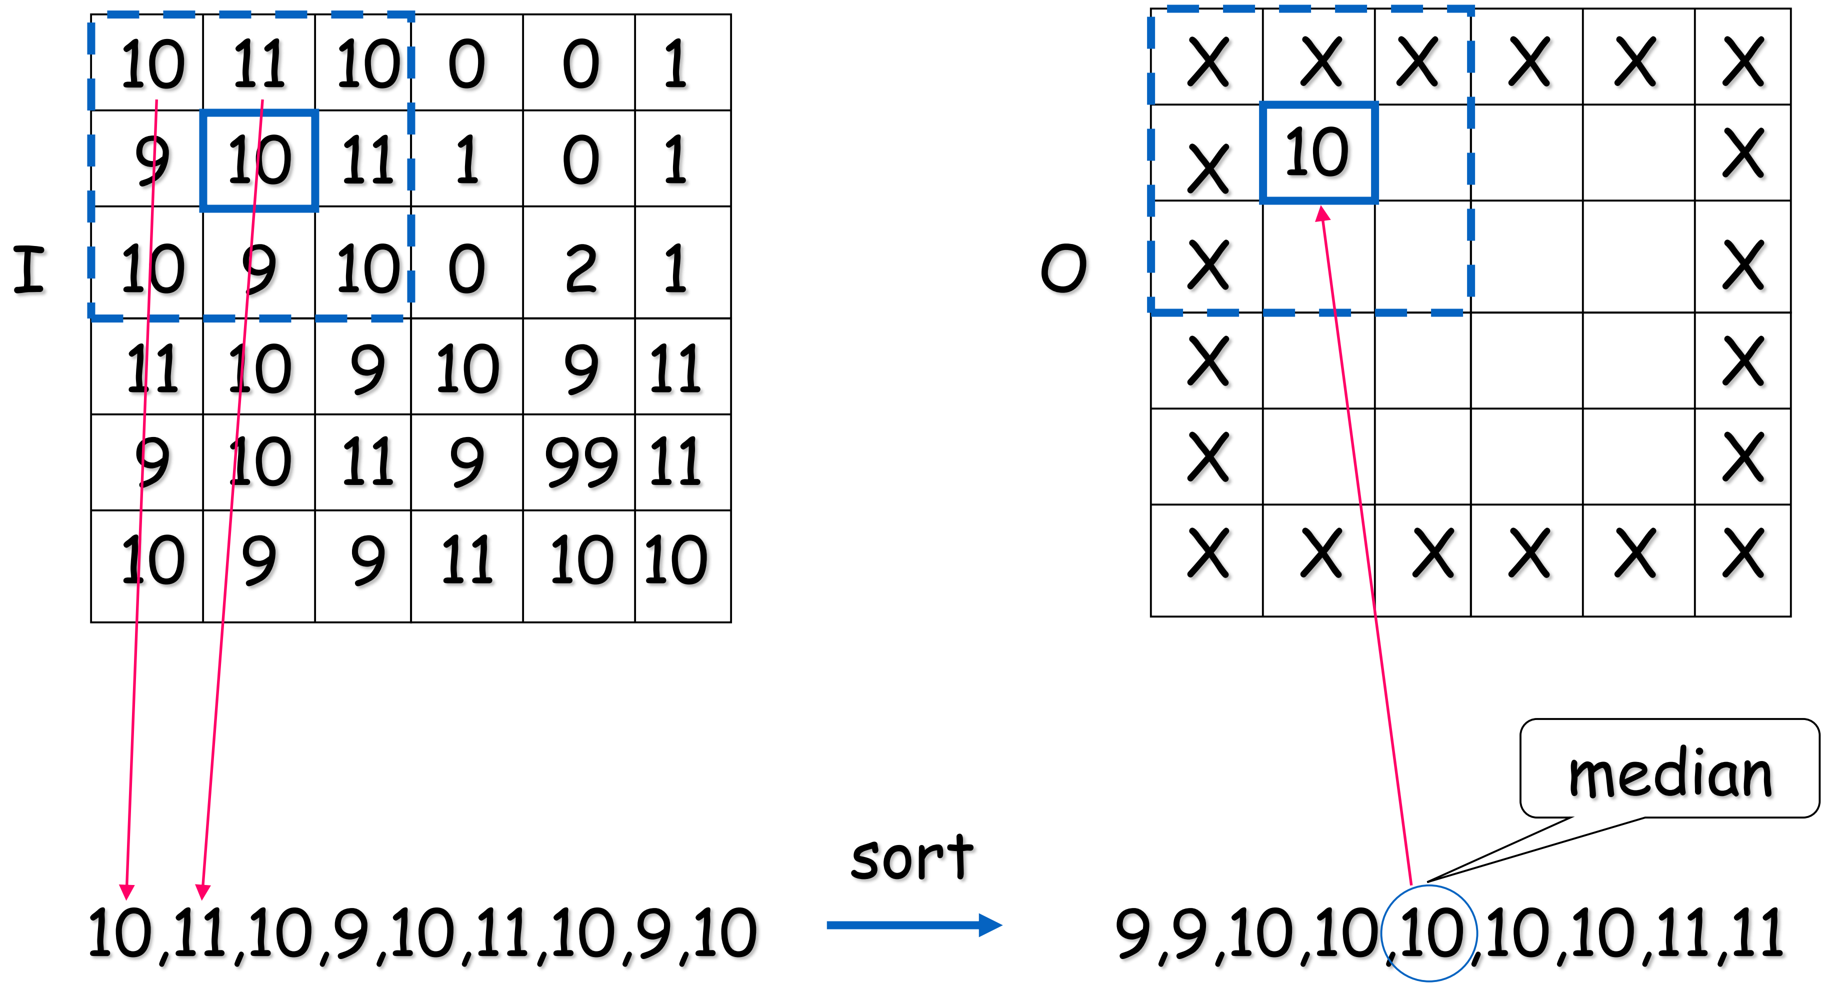

Adapted from Octavia Camps, Penn State

Order-statistic filters

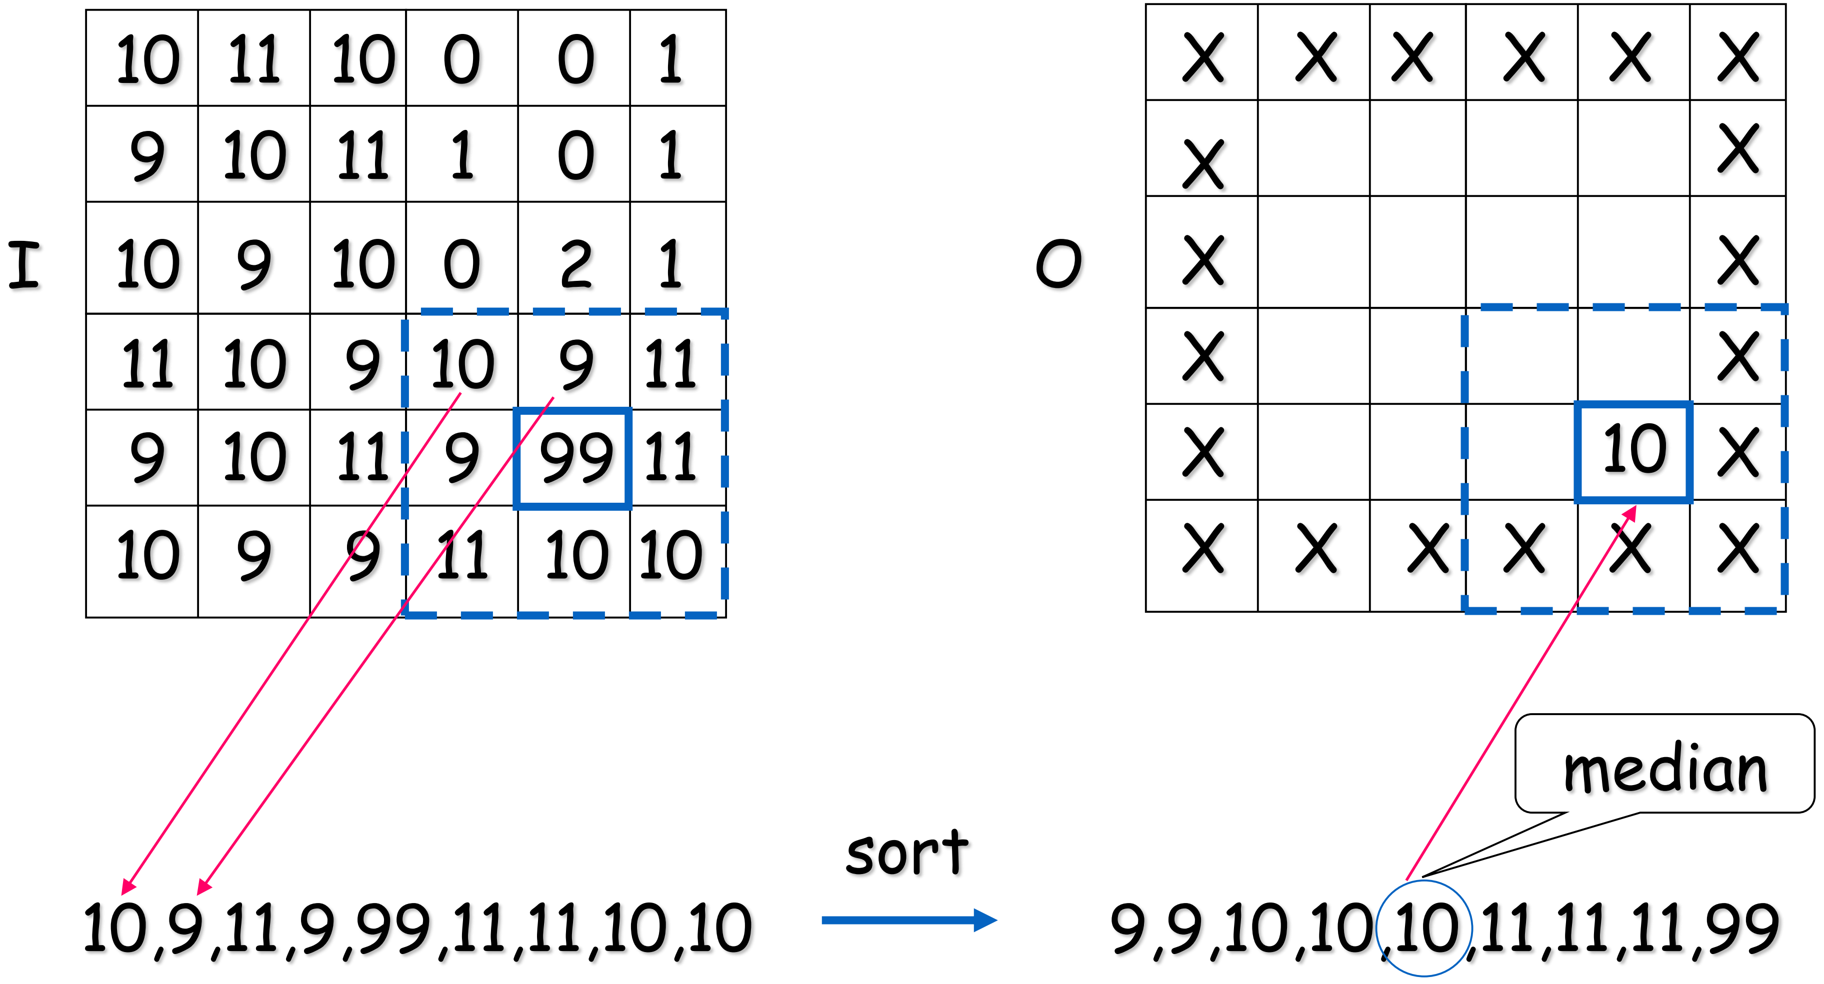

Adapted from Octavia Camps, Penn State

Salt-and-pepper noise

Adapted from Linda Shapiro,
U of Washington

Gaussian noise

Adapted from Linda Shapiro,
U of Washington

Order-statistic filters

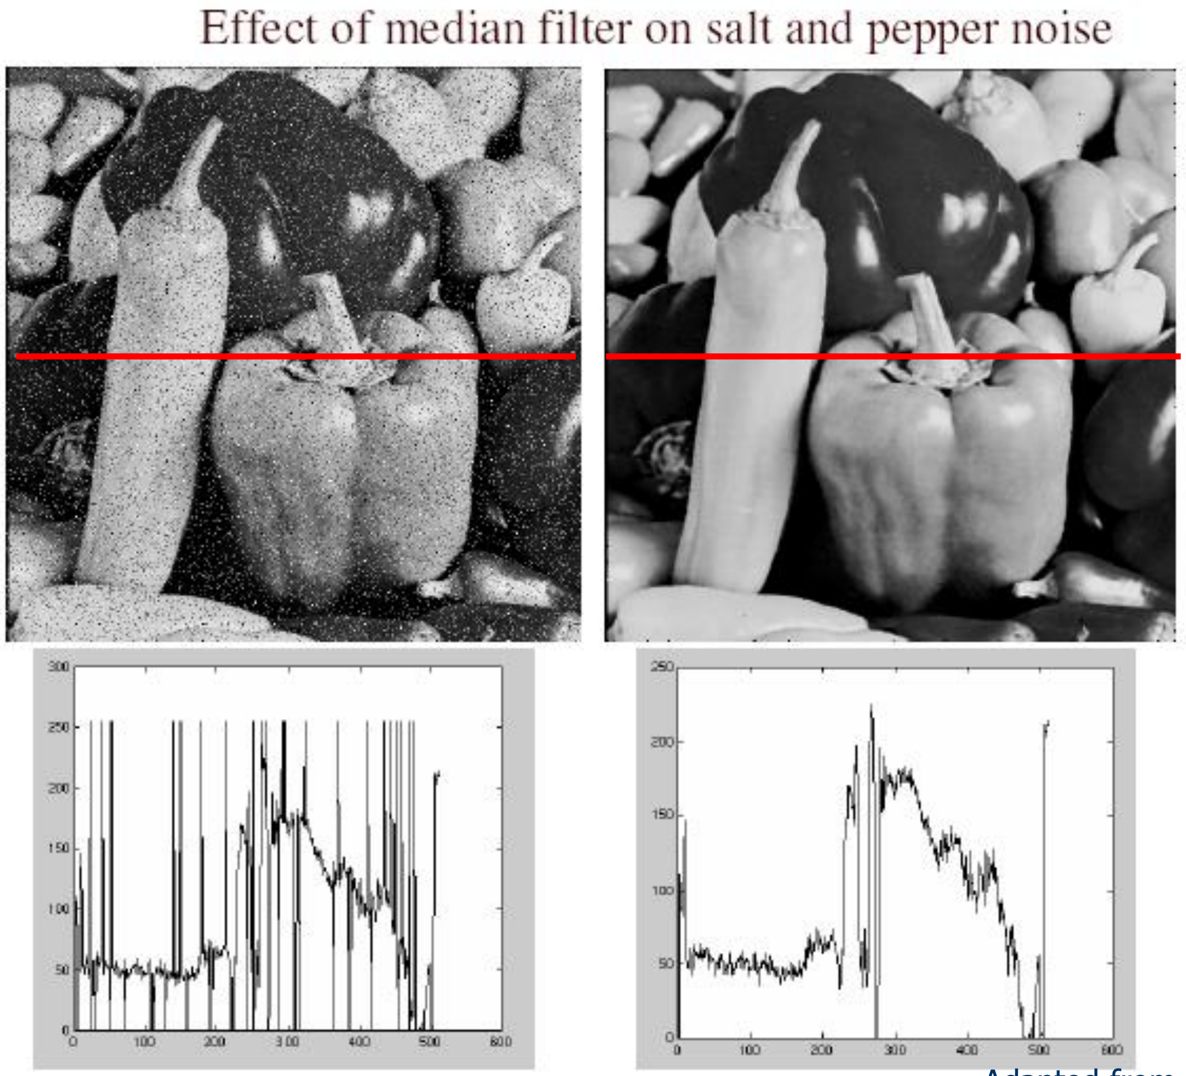

Adapted from Martial Hebert, CMU

Spatially varying filters

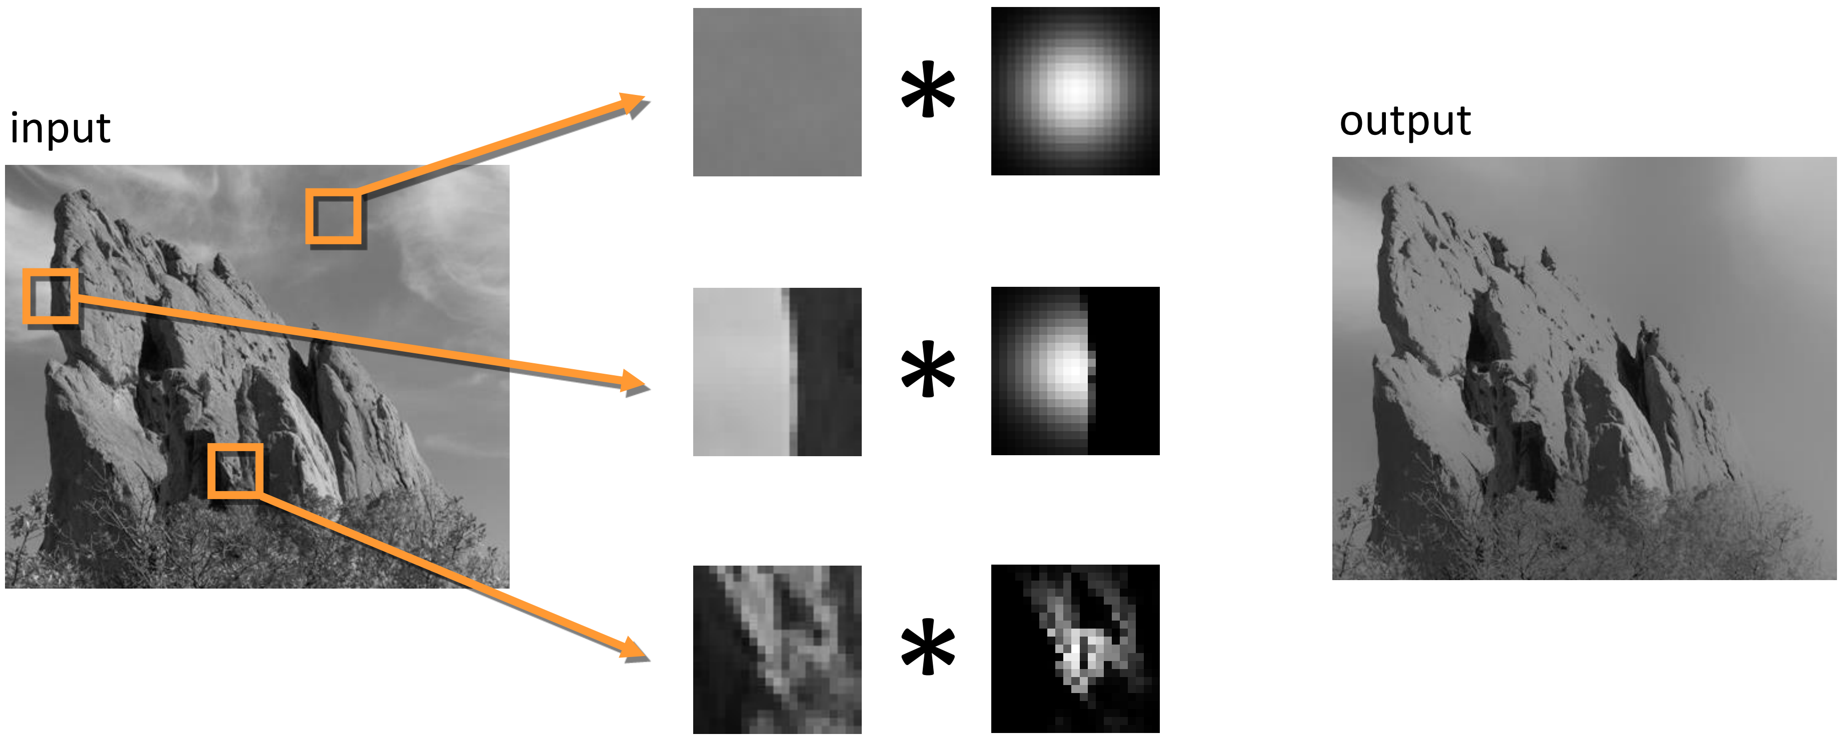

Bilateral filter: kernel depends on the local image content.
See the Szeliski book for the math.

Adapted from Sylvian Paris

Spatially varying filters

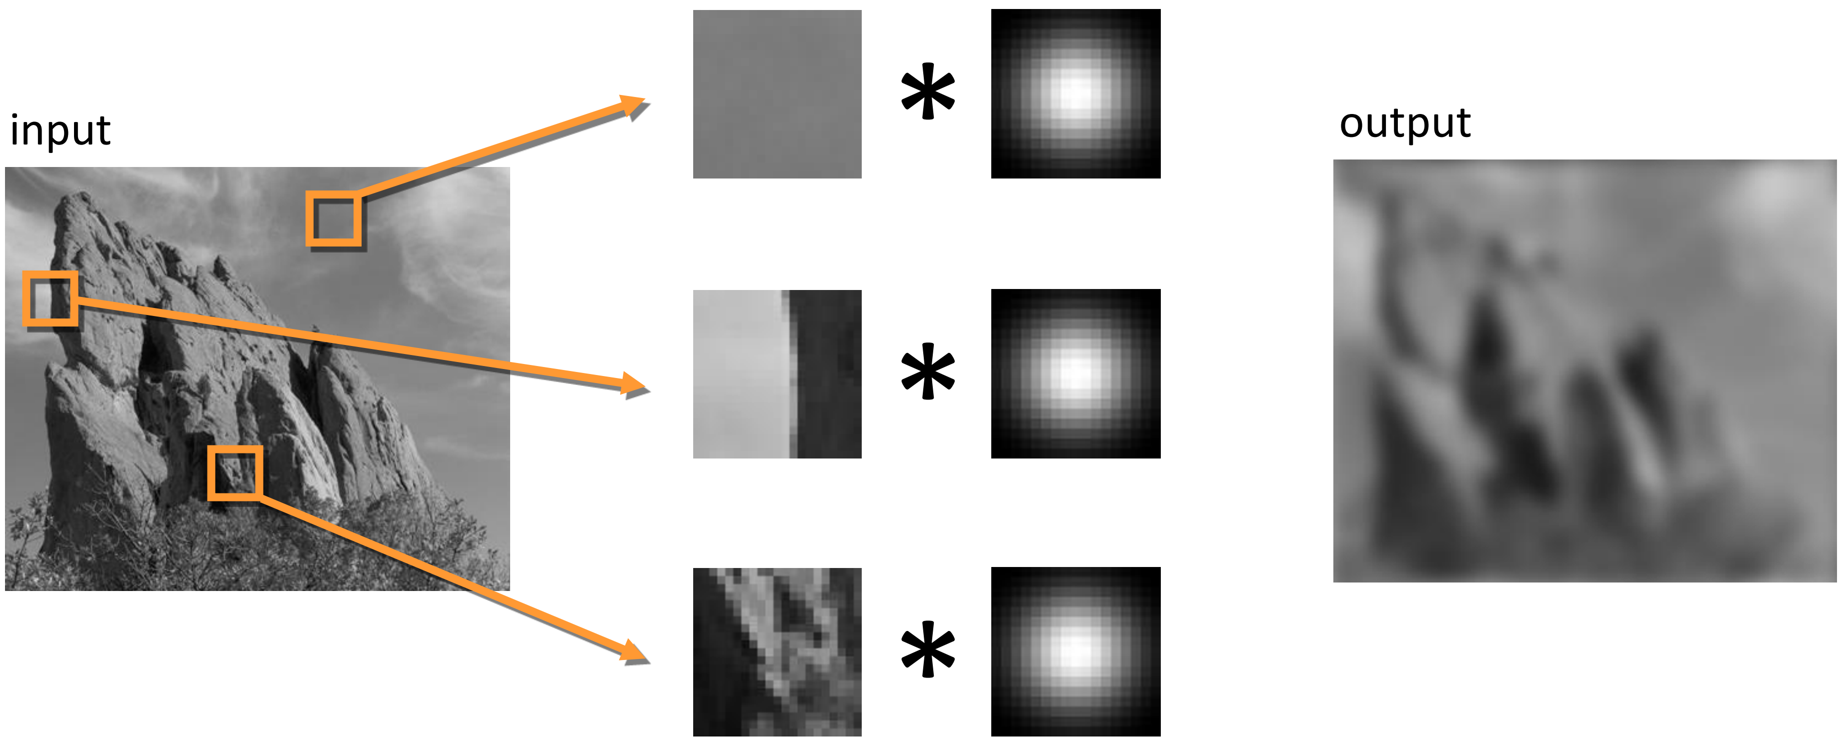

Compare to the result of using the same Gaussian kernel everywhere

Adapted from Sylvian Paris

Sharpening spatial filters

- Objective of sharpening is to highlight or enhance fine detail in an image.
- Since smoothing (averaging) is analogous to integration, sharpening can be accomplished by spatial differentiation.
- First-order derivative of 1D function f(x)

$$f(x+1) - f(x).$$

- Second-order derivative of 1D function f(x)

$$f(x+1) - 2f(x) + f(x-1).$$

Sharpening spatial filters

- For a function $f(x, y)$, the *gradient* at (x, y) is defined as

$$\nabla f = \left[\frac{\partial f}{\partial x} \; \frac{\partial f}{\partial y} \right]^T$$

where its magnitude can be used to implement first-order derivatives.

−1	0	0	−1	
0	1	1	0	

Robert's cross-gradient operators

−1	−2	−1	−1	0	1
0	0	0	−2	0	2
1	2	1	−1	0	1

Sobel gradient operators

Sharpening spatial filters

- *Laplacian* of a function (image) $f(x, y)$ of two variables x and y

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

is a second-order derivative operator.

(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 3.39 (a) Filter mask used to implement the digital Laplacian, as defined in Eq. (3.7-4). (b) Mask used to implement an extension of this equation that includes the diagonal neighbors. (c) and (d) Two other implementations of the Laplacian.

Sharpening spatial filters

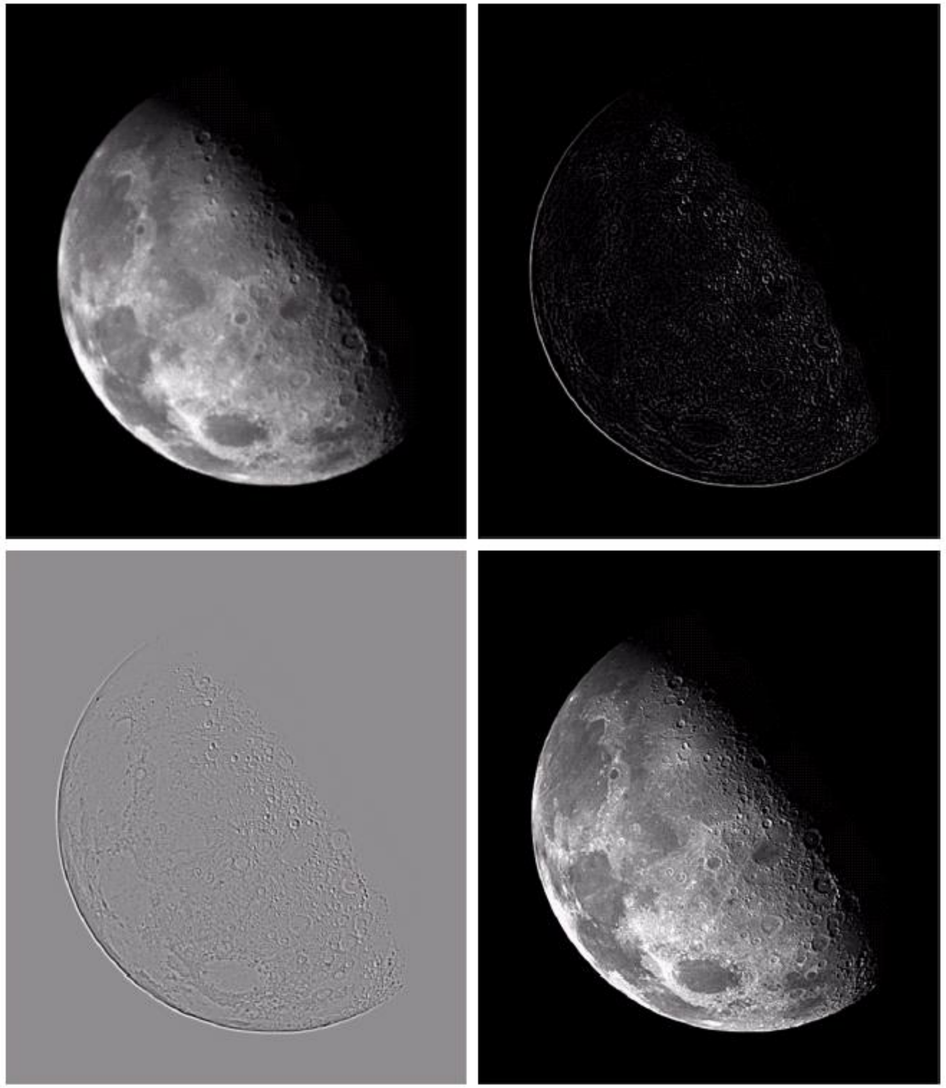

FIGURE 3.40 (a) Image of the North Pole of the moon. (b) Laplacian-filtered image. (c) Laplacian image scaled for display purposes. (d) Image enhanced by using Eq. (3.7-5). (Original image courtesy of NASA.)

Adapted from Gonzales and Woods

Sharpening spatial filters

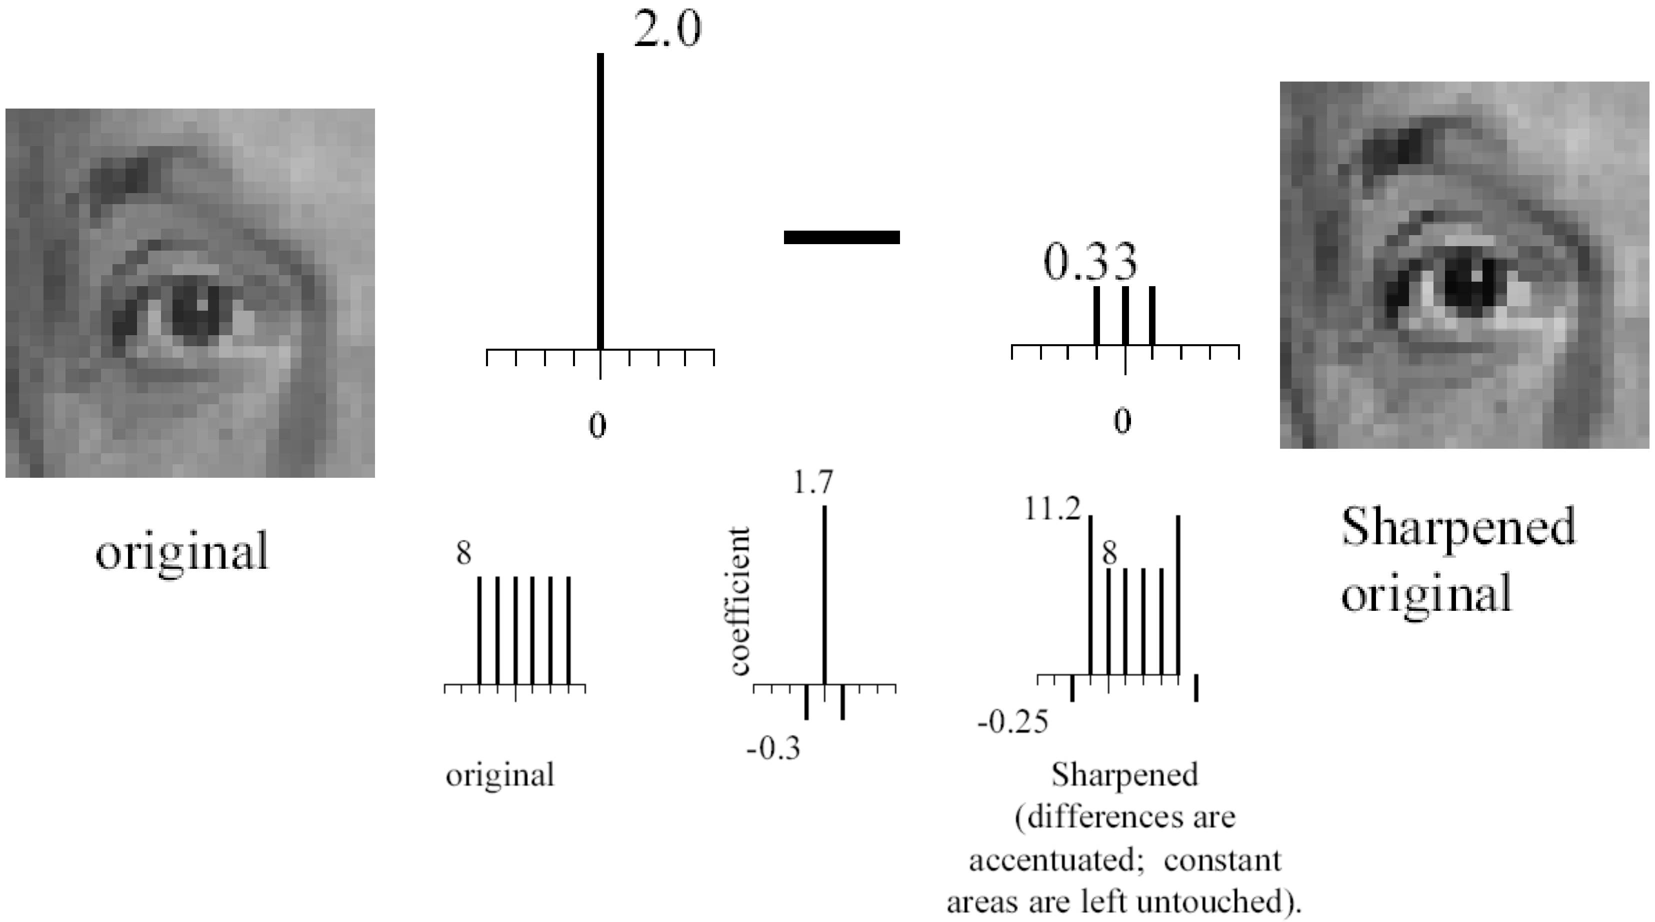

High-boost filtering

Adapted from Darrell and Freeman, MIT

Sharpening spatial filters

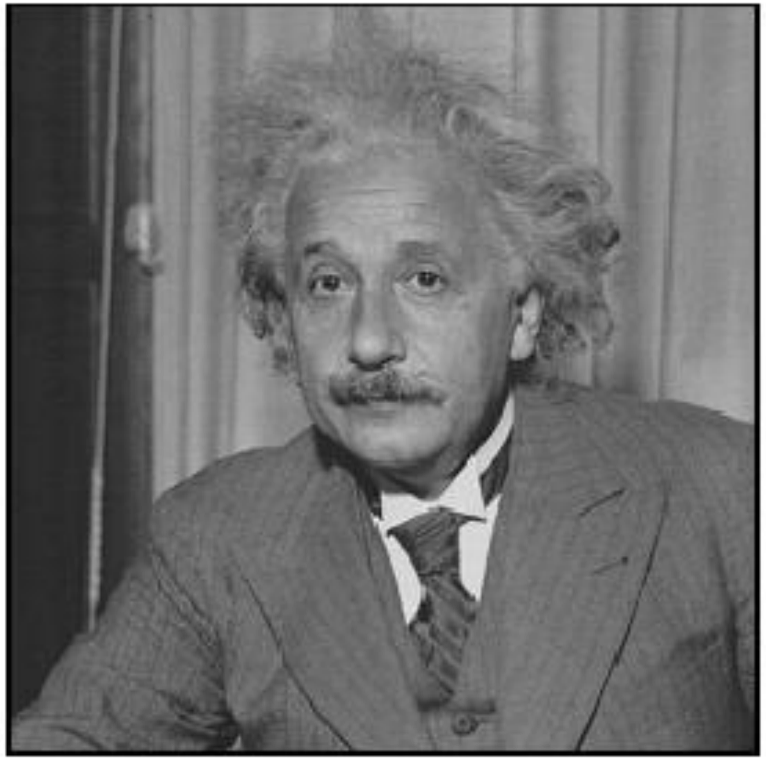

before

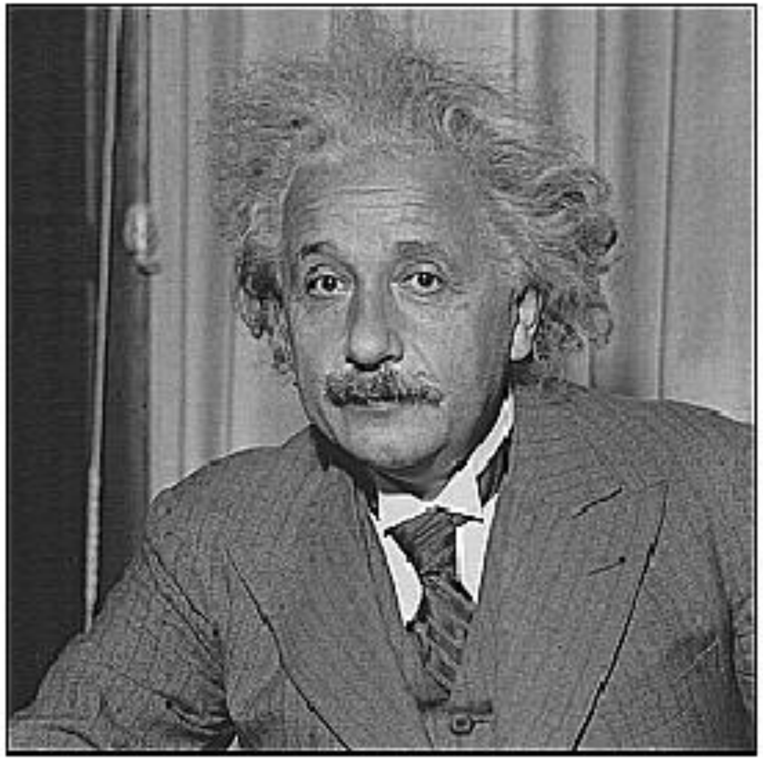

after

Adapted from Darrell and Freeman, MIT

Combining spatial enhancement methods

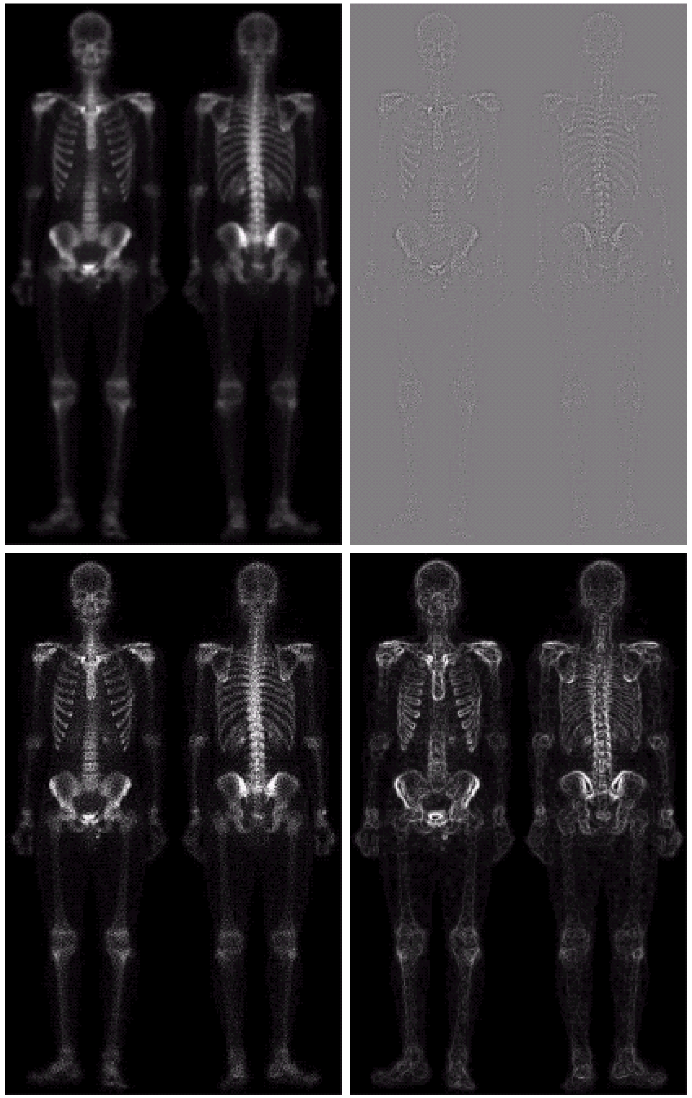

a b
c d

FIGURE 3.46
(a) Image of whole body bone scan. (b) Laplacian of (a). (c) Sharpened image obtained by adding (a) and (b). (d) Sobel of (a).

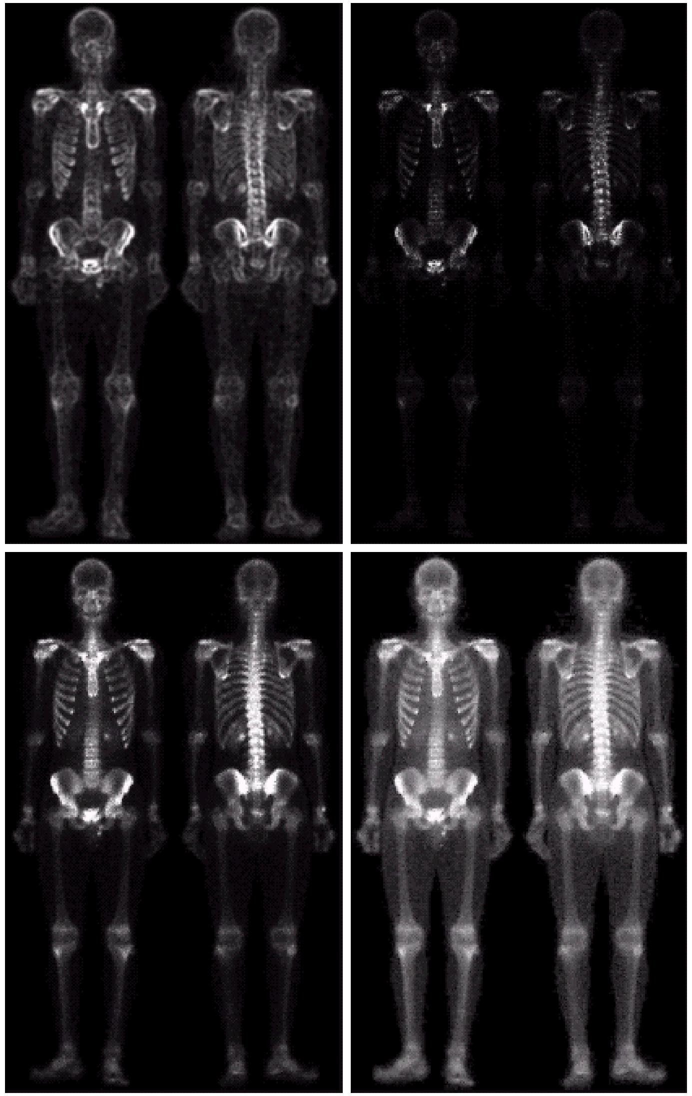

e f
g h

FIGURE 3.46
(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.)

Template matching

- Correlation can also be used for matching.
- If we want to determine whether an image f contains a particular object, we let h be that object (also called a template) and compute the correlation between f and h.
- If there is a match, the correlation will be maximum at the location where h finds a correspondence in f.
- Preprocessing such as scaling and alignment is necessary in most practical applications.

Template matching

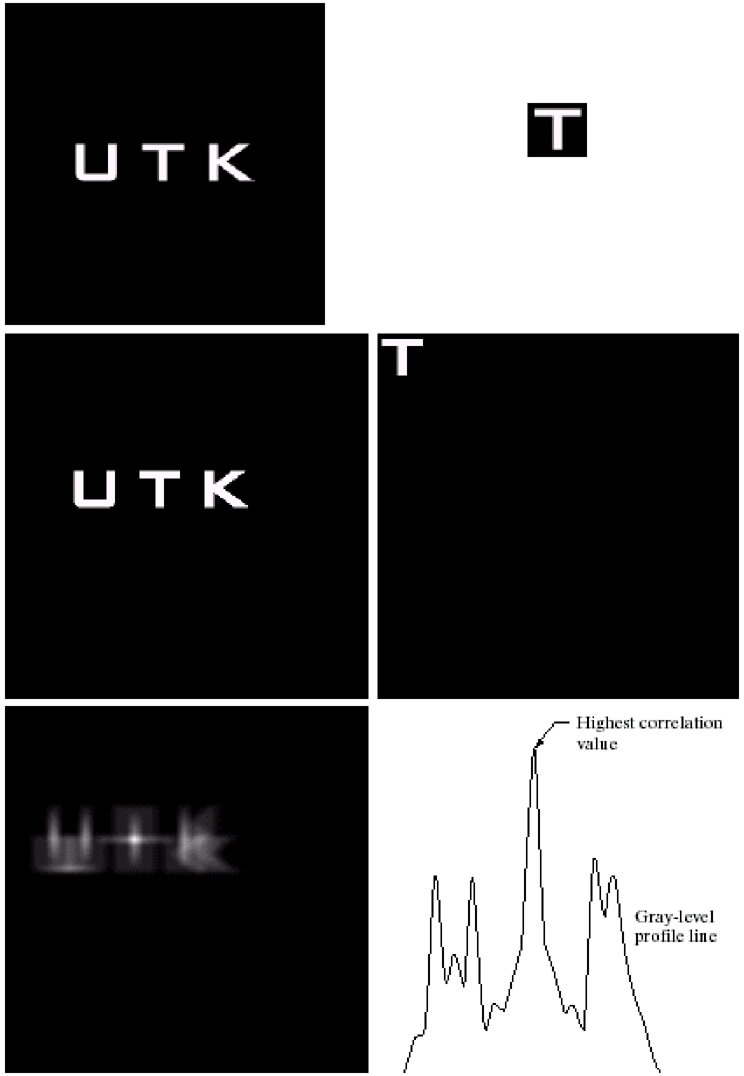

a b
c d
e f

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.

Adapted from Gonzales and Woods

Template matching

Face detection using template matching: face templates.

Template matching

Face detection using template matching: detected faces.

Template matching

Where is Waldo?

http://machinelearningmastery.com/using-opencv-python-and-template-matching-to-play-wheres-waldo/

Resizing images

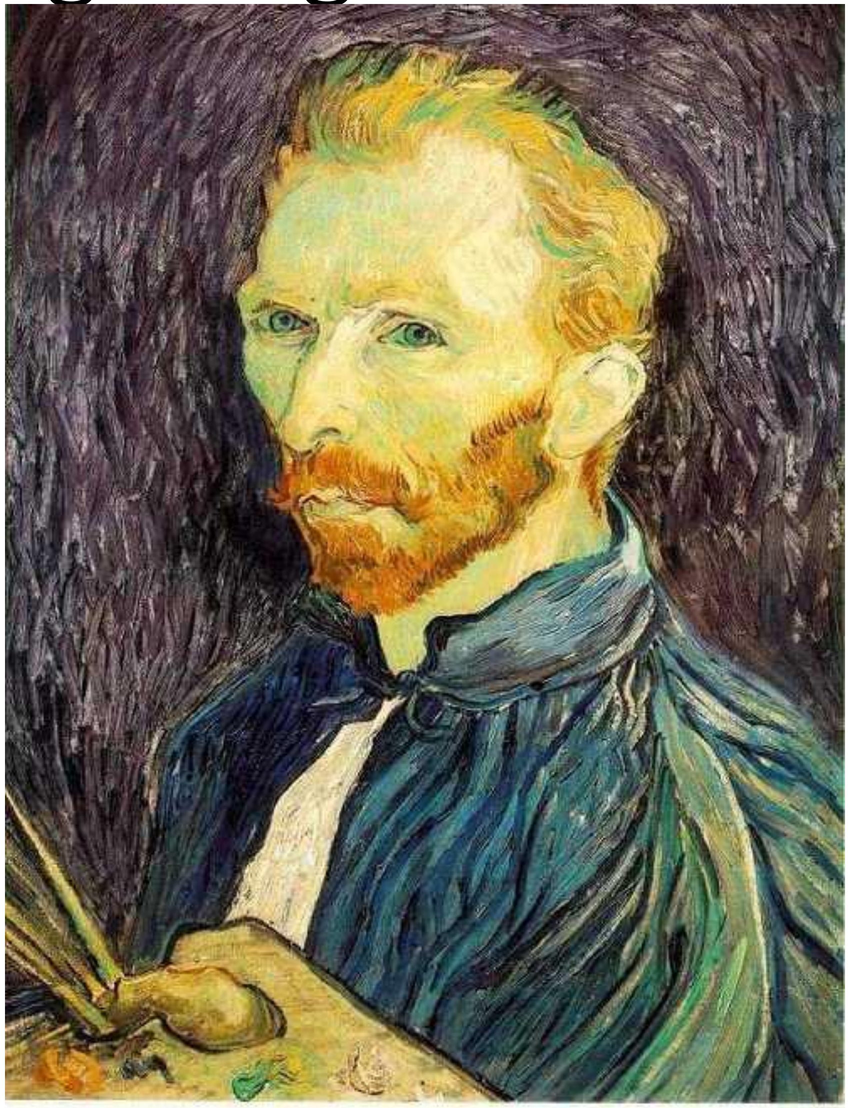

How can we generate a half-sized version of a large image?

Adapted from Steve Seitz, U of Washington

Resizing images

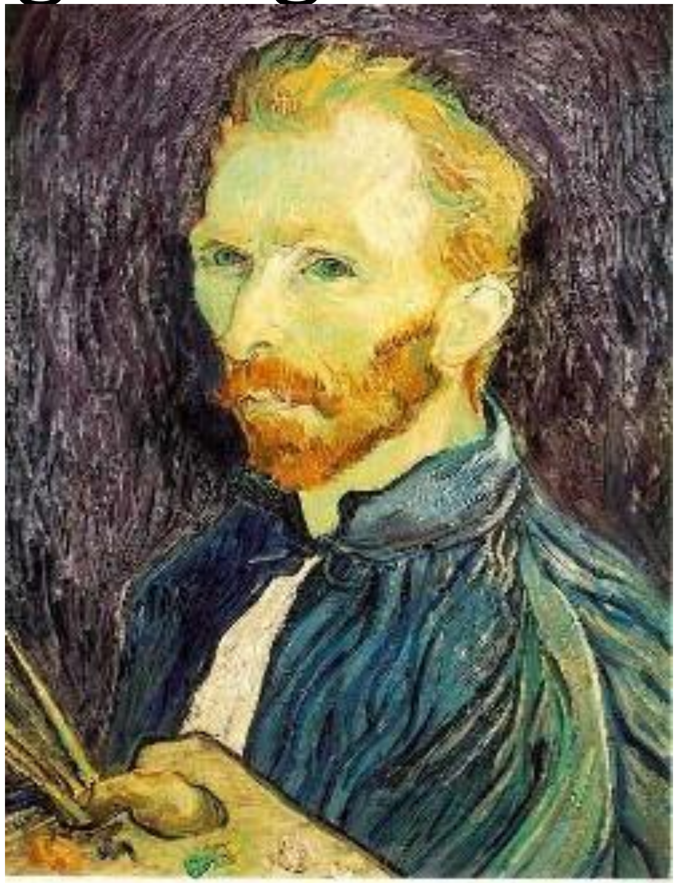

1/4

1/8

Throw away every other row and column to create a 1/2 size image (also called sub-sampling).

Adapted from Steve Seitz, U of Washington

Resizing images

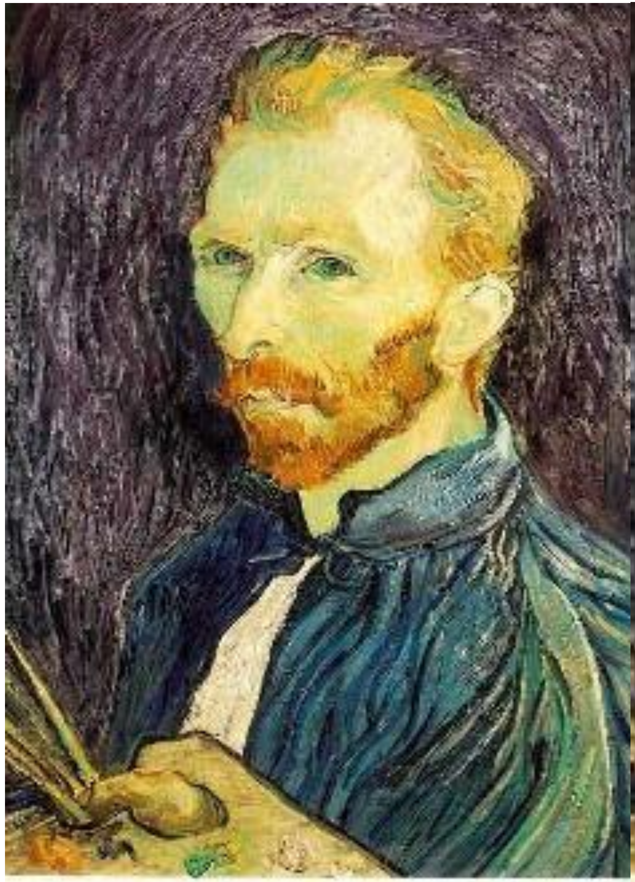

1/2

1/4 (2x zoom)

1/8 (4x zoom)

Does this look nice?

Adapted from Steve Seitz, U of Washington

Resizing images

- We cannot shrink an image by simply taking every k'th pixel.
- Solution: smooth the image, then sub-sample.

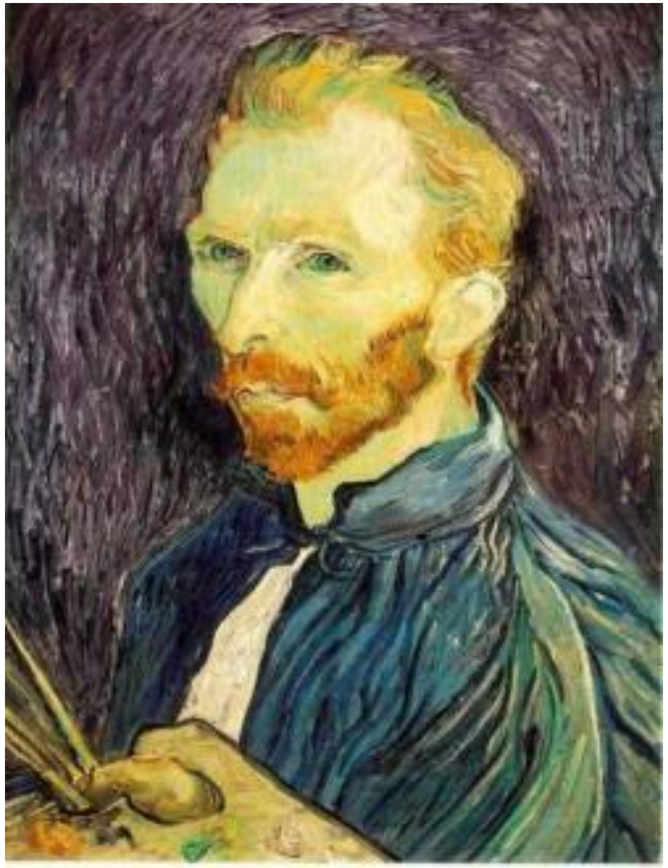

Gaussian 1/2

Gaussian 1/4

Gaussian 1/8

Adapted from Steve Seitz, U of Washington

Resizing images

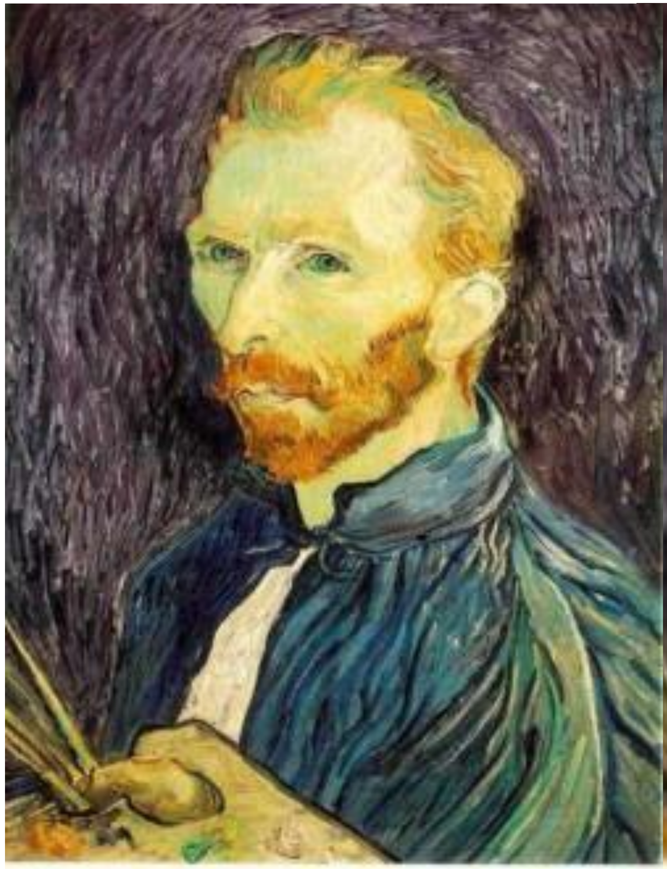

Gaussian 1/2

Gaussian 1/4
(2x zoom)

Gaussian 1/8
(4x zoom)

Adapted from Steve Seitz, U of Washington

Sampling and aliasing

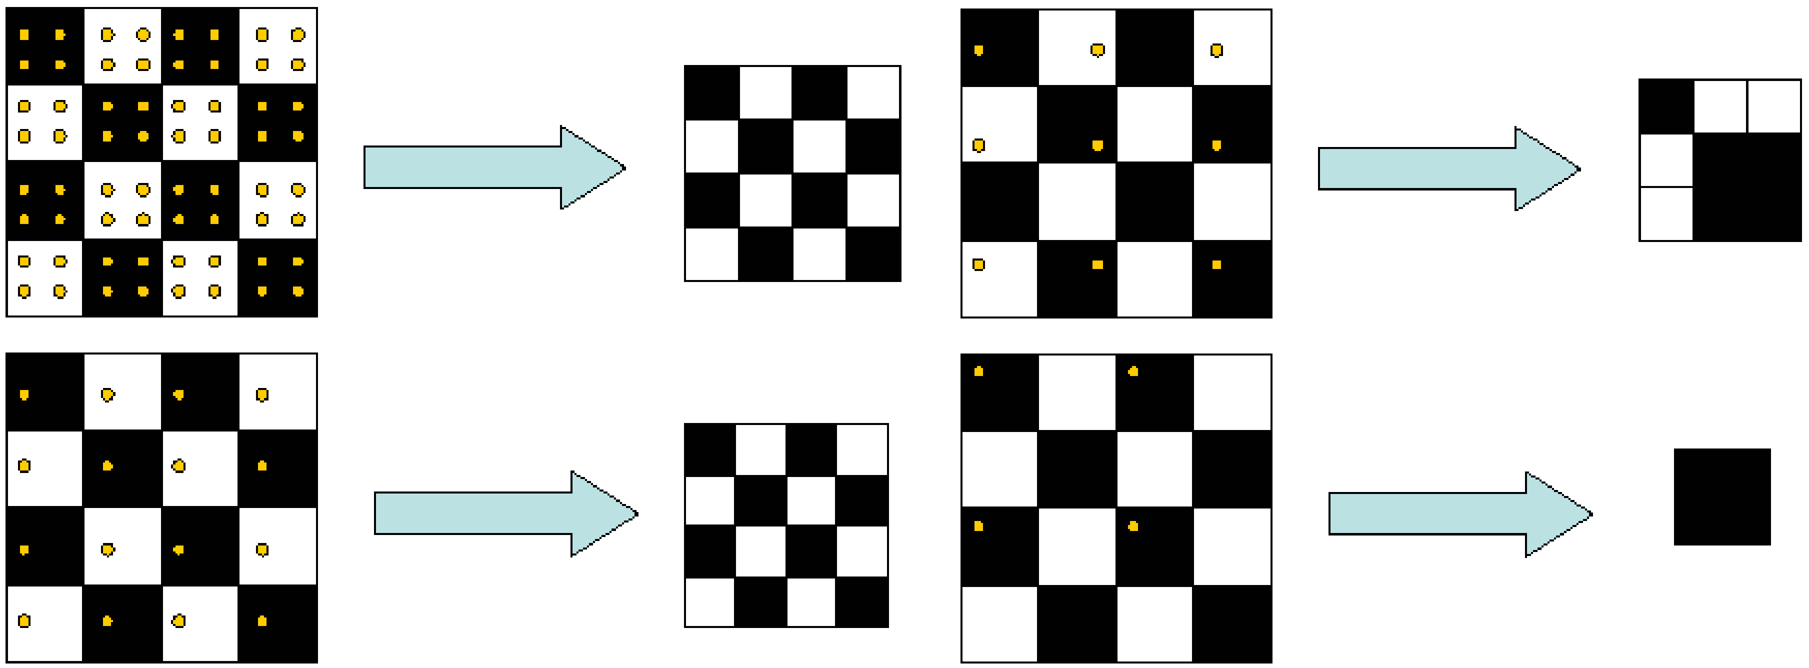

Examples of GOOD sampling Examples of BAD sampling -> Aliasing

Adapted from Steve Seitz, U of Washington

Sampling and aliasing

- Errors appear if we do not sample properly.
- Common phenomenon:
  - High spatial frequency components of the image appear as low spatial frequency components.
- Examples:
  - Wagon wheels rolling the wrong way in movies.
  - Checkerboards misrepresented in ray tracing.
  - Striped shirts look funny on color television.

Sampling and aliasing

Moire patterns in real-world images. Here are comparison images by Dave Etchells of Imaging Resource using the Canon D60 (with an antialias filter) and the Sigma SD-9 (which has no antialias filter). The bands below the fur in the image at right are the kinds of artifacts that appear in images when no antialias filter is used. Sigma chose to eliminate the filter to get more sharpness, but the resulting apparent detail may or may not reflect features in the image.

Adapted from Ali Farhadi

Gaussian pyramids

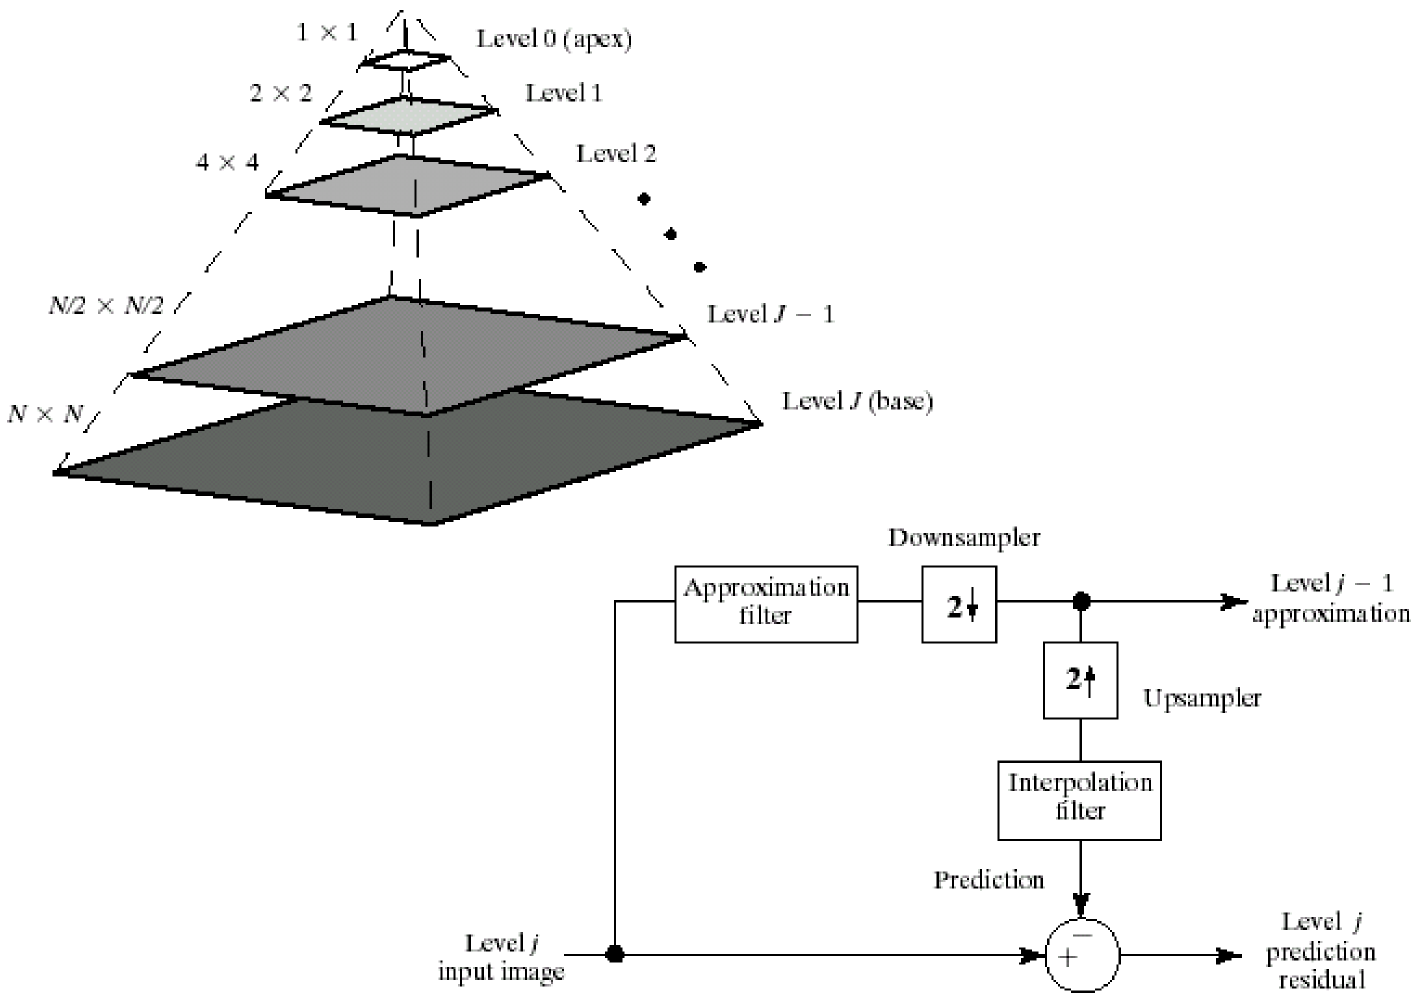

a
b

FIGURE 7.2 (a) A pyramidal image structure and (b) system block diagram for creating it.

Adapted from Gonzales and Woods

Gaussian pyramids

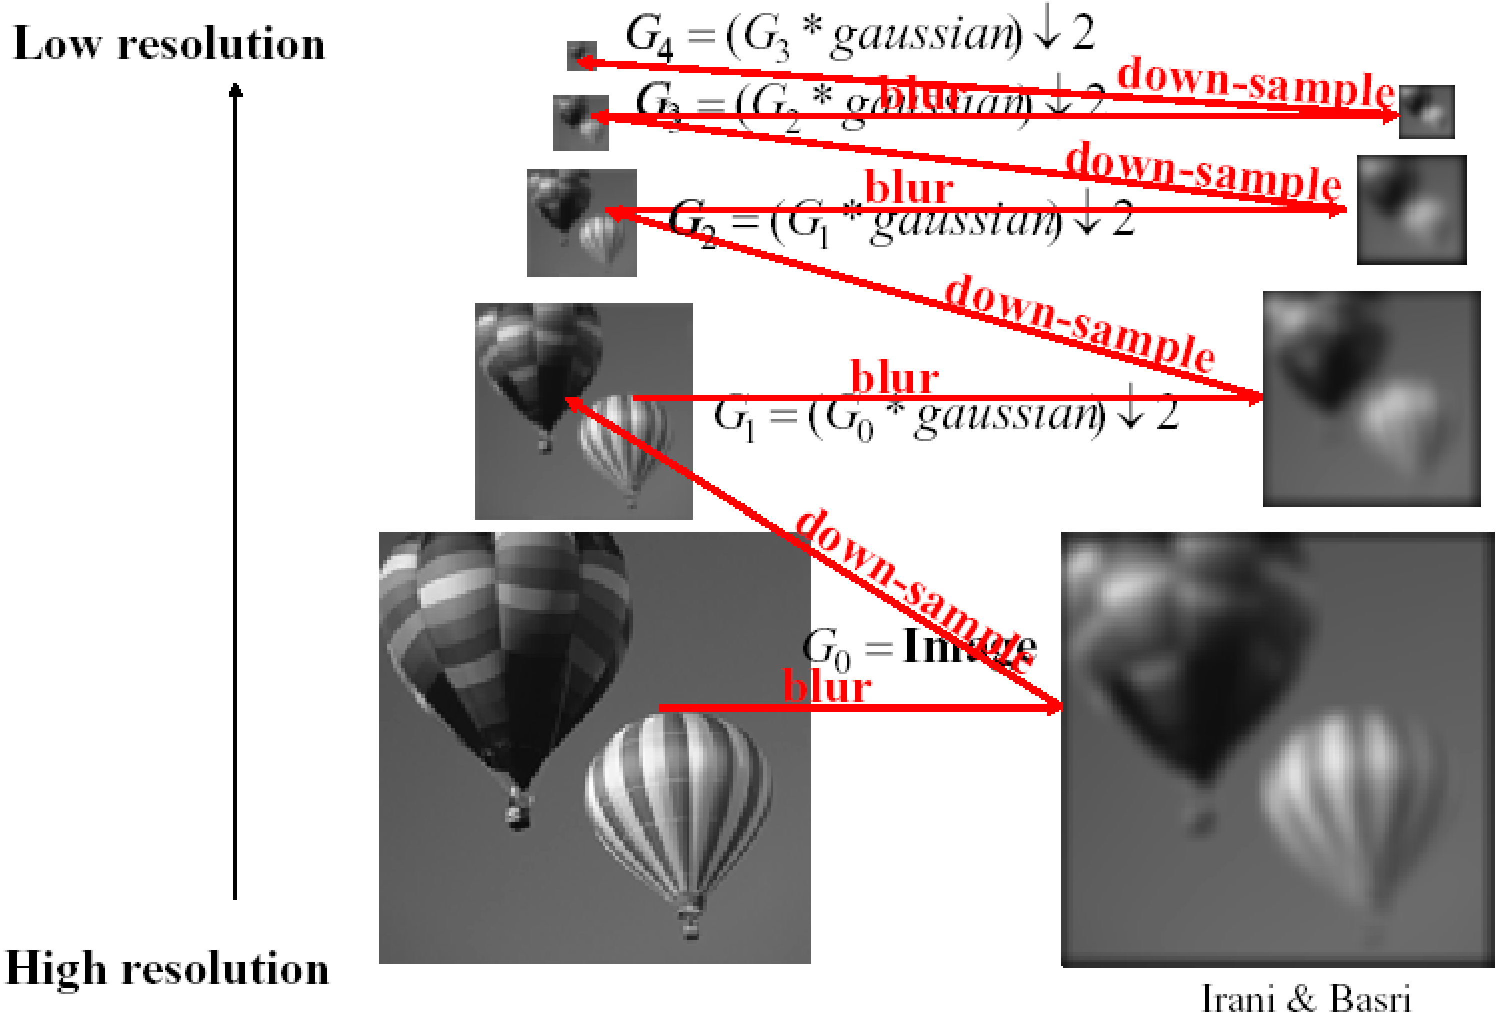

Adapted from Michael Black, Brown University

Gaussian pyramids

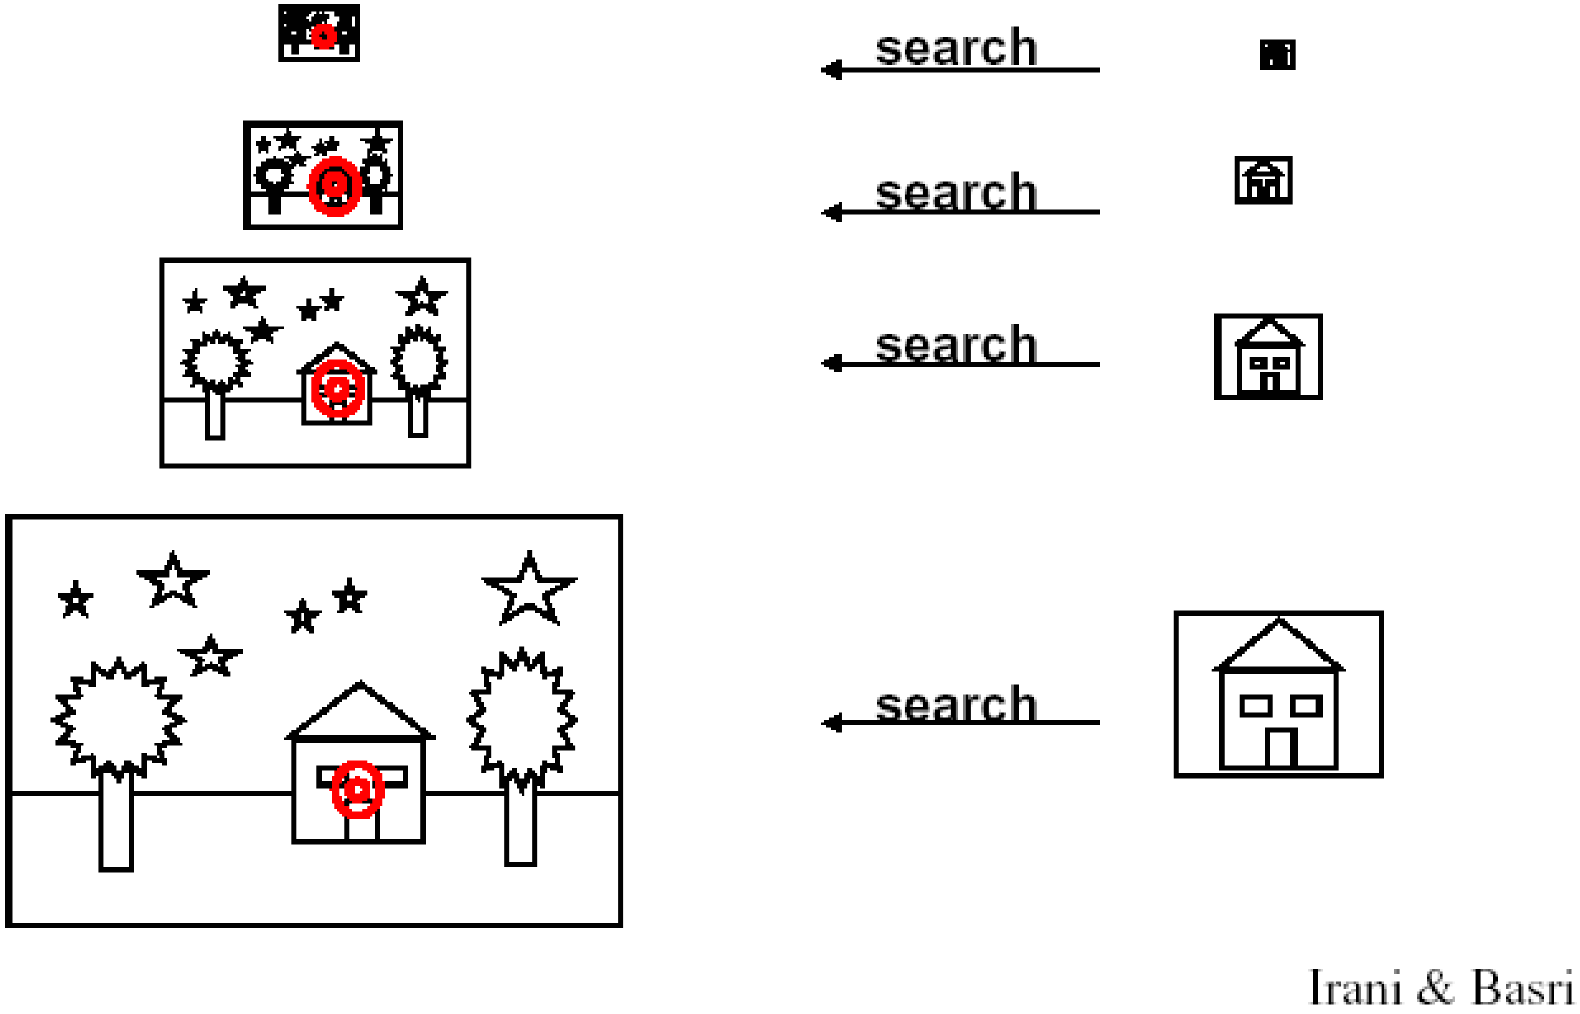

Adapted from Michael Black, Brown University