

CS4442/9542b  
Artificial Intelligence II  
prof. Olga Veksler

*Lecture 13*

*Computer Vision*

**Introduction, Filtering**

Some slides from: D. Jacobs, D. Lowe, S. Seitz, A. Efros, X. Li, R. Fergus, J. Hayes, S. Lazebnik, D. Hoiem, S. Marschner

# Outline

- Very Brief Intro to Computer Vision
- Digital Images
- Image Filtering
  - noise reduction

# Every Picture Tells a Story

- Goal of computer vision is to write computer programs that can interpret images
  - bridge the gap between the pixels and the story



**what we see**

1	2	0	2	2	1
9	2	2	7	1	2
2	8	2	3	2	2
4	2	2	7	2	8
2	2	2	6	0	2
8	3	2	5	2	2
7	2	4	2	1	9

**what computers see**

# Origin of Computer Vision: MIT Summer Project

MASSACHUSETTS INSTITUTE OF TECHNOLOGY  
PROJECT MAC

Artificial Intelligence Group  
Vision Memo. No. 100.

July 7, 1966

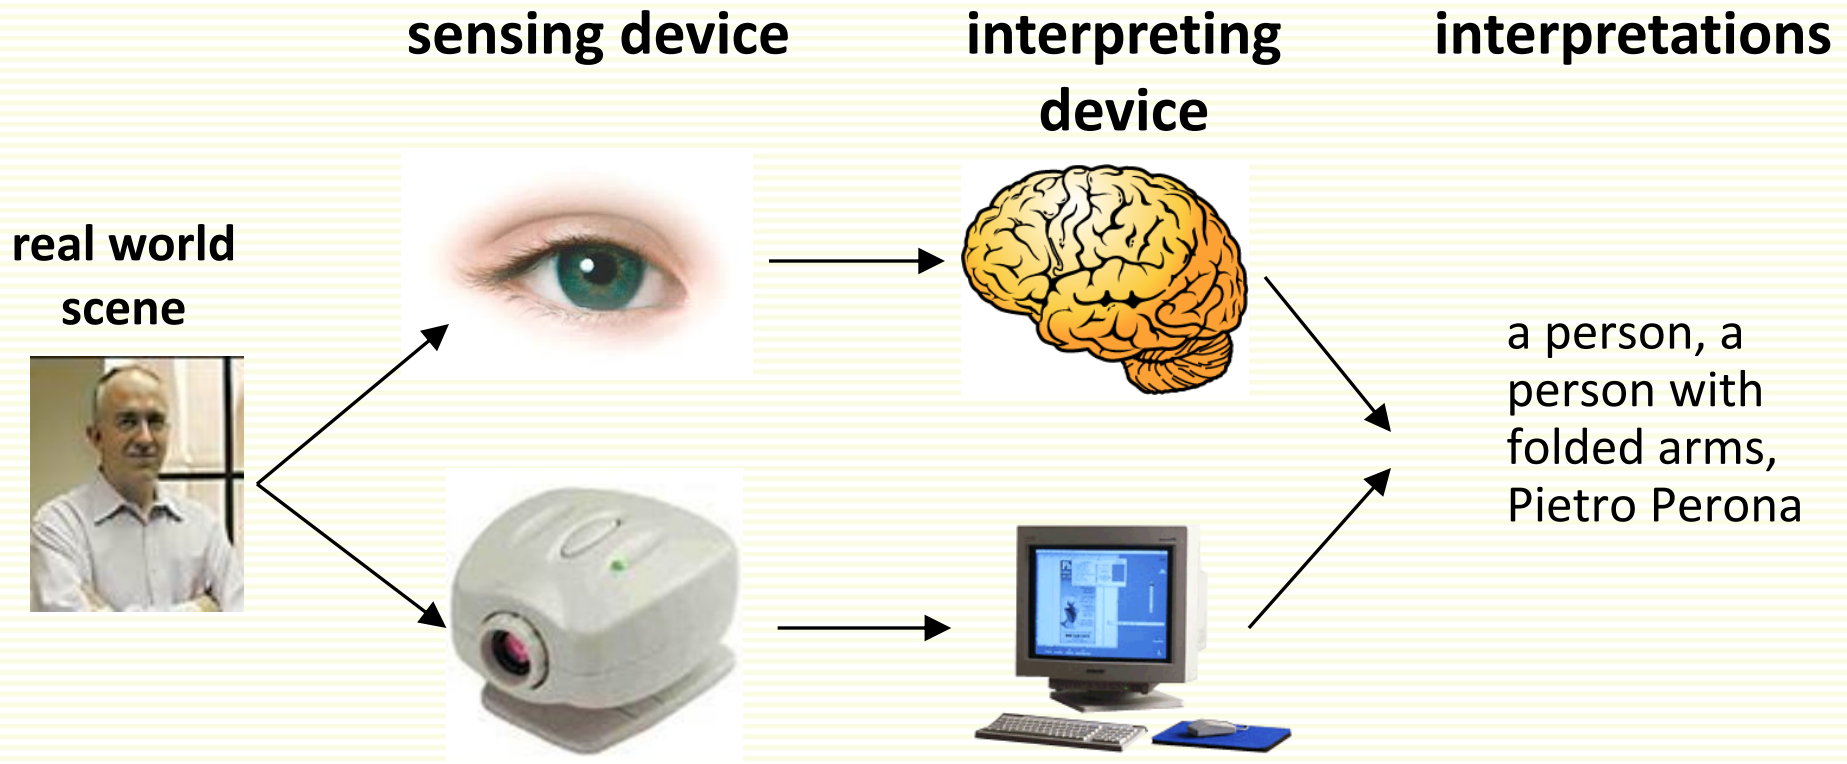
THE SUMMER VISION PROJECT

Seymour Papert

The summer vision project is an attempt to use our summer workers effectively in the construction of a significant part of a visual system. The particular task was chosen partly because it can be segmented into sub-problems which will allow individuals to work independently and yet participate in the construction of a system complex enough to be a real landmark in the development of "pattern recognition".

# The problem

- Want to make a computer understand images
- We know it is possible, we do it effortlessly!



# Just Copy Human Visual System?

- People try to but we don't yet have a sufficient understanding of how our visual system works
- $O(10^{11})$  neurons used in vision
  - about 1/3 of human brain
- Latest CPUs have only  $O(10^8)$  transistors
  - most are cache memory
- Very different architectures:
  - Brain is slow but parallel
  - Computer is fast but mainly serial
- Bird vs Airplane
  - Same underlying principles
  - Very different hardware






# Why Computer Vision Matters




Safety



Health



Security



Comfort



Fun



Personal Photos

# “Early Vision” Problems



- Edge extraction



- Corner extraction



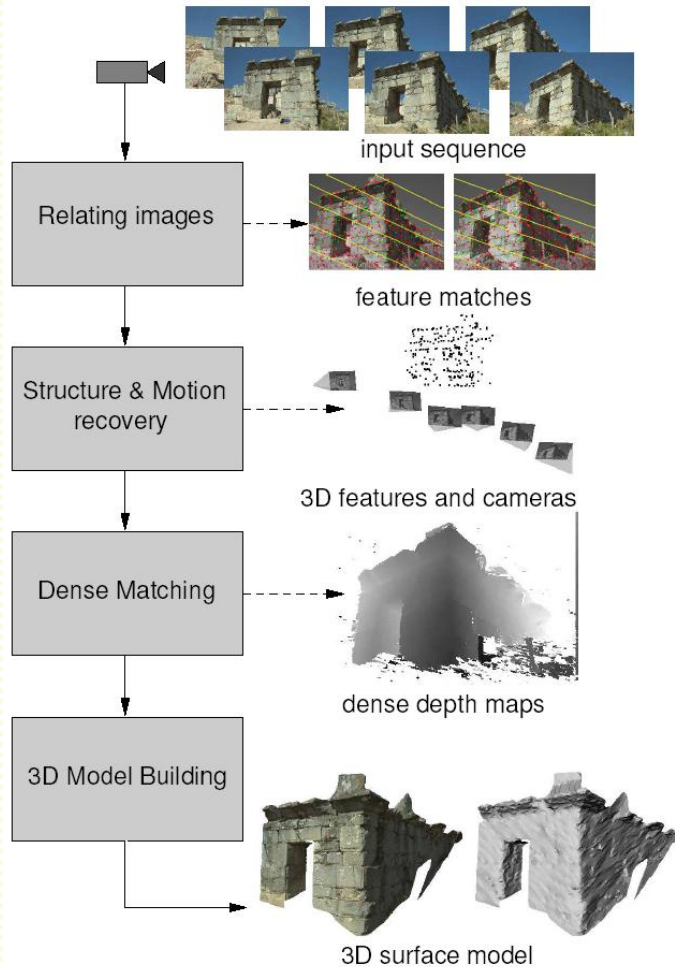
- Blob extraction





# “Mid-level Vision” Problems

- 3D Structure extraction



- Motion and tracking



- Segmentation




# “High-level Vision” Problems

- Face Detection



- Action Recognition



walk

skate

- Object Recognition




- Scene Recognition




# Vision is inferential: Illumination


- Vision is hard: even the simple problem of color perception is inferential



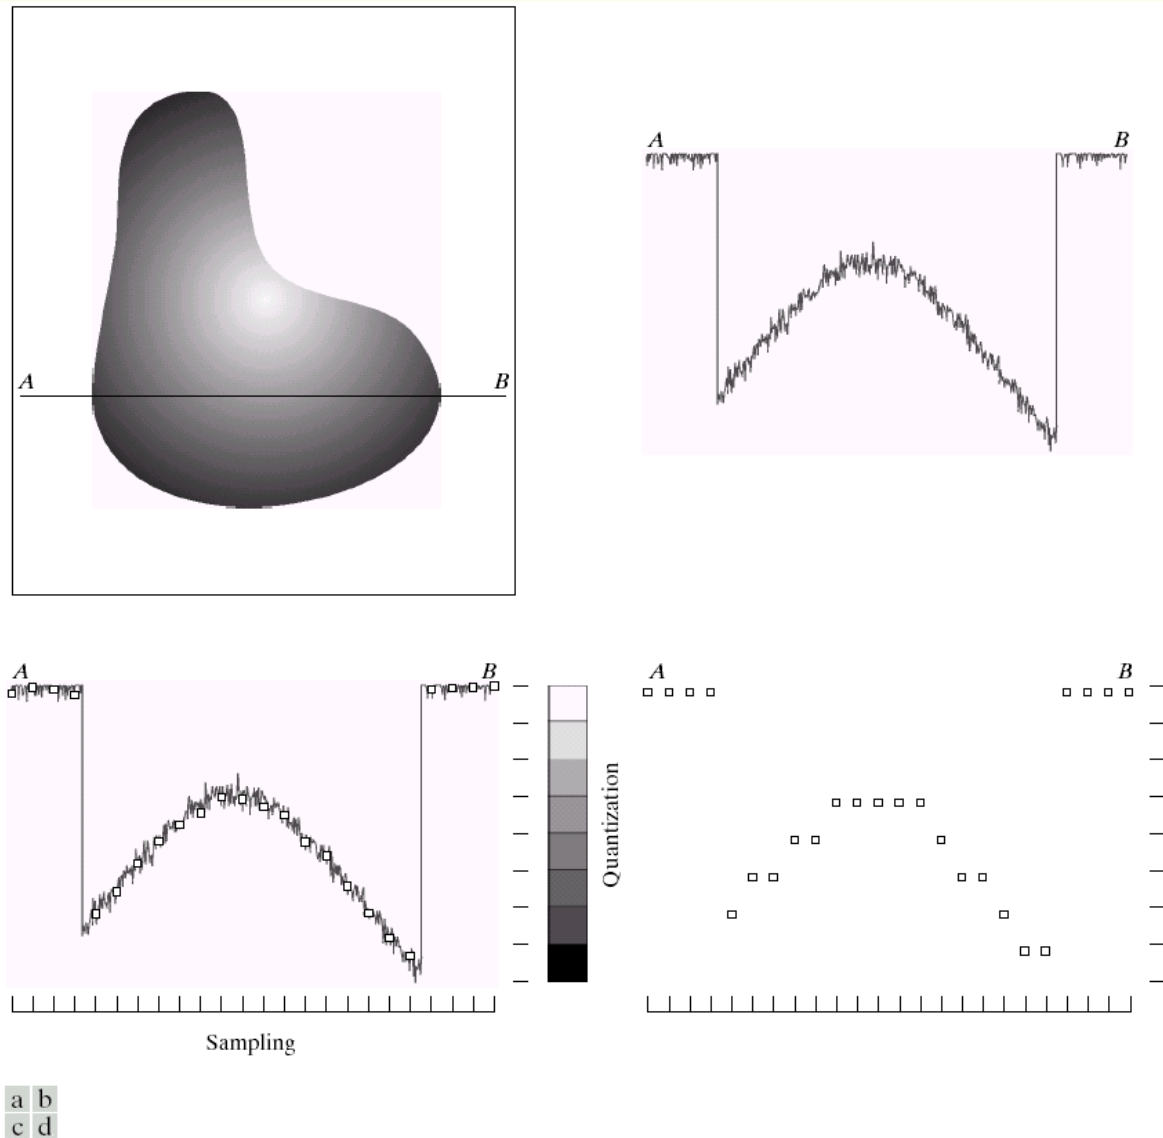
Edward H. Adelson



# Image Formation




# Sampling and Quantization




**FIGURE 2.16** Generating a digital image. (a) Continuous image. (b) A scan line from *A* to *B* in the continuous image, used to illustrate the concepts of sampling and quantization. (c) Sampling and quantization. (d) Digital scan line.



# Sensor Array



**real world object**



**after quantization and sampling**



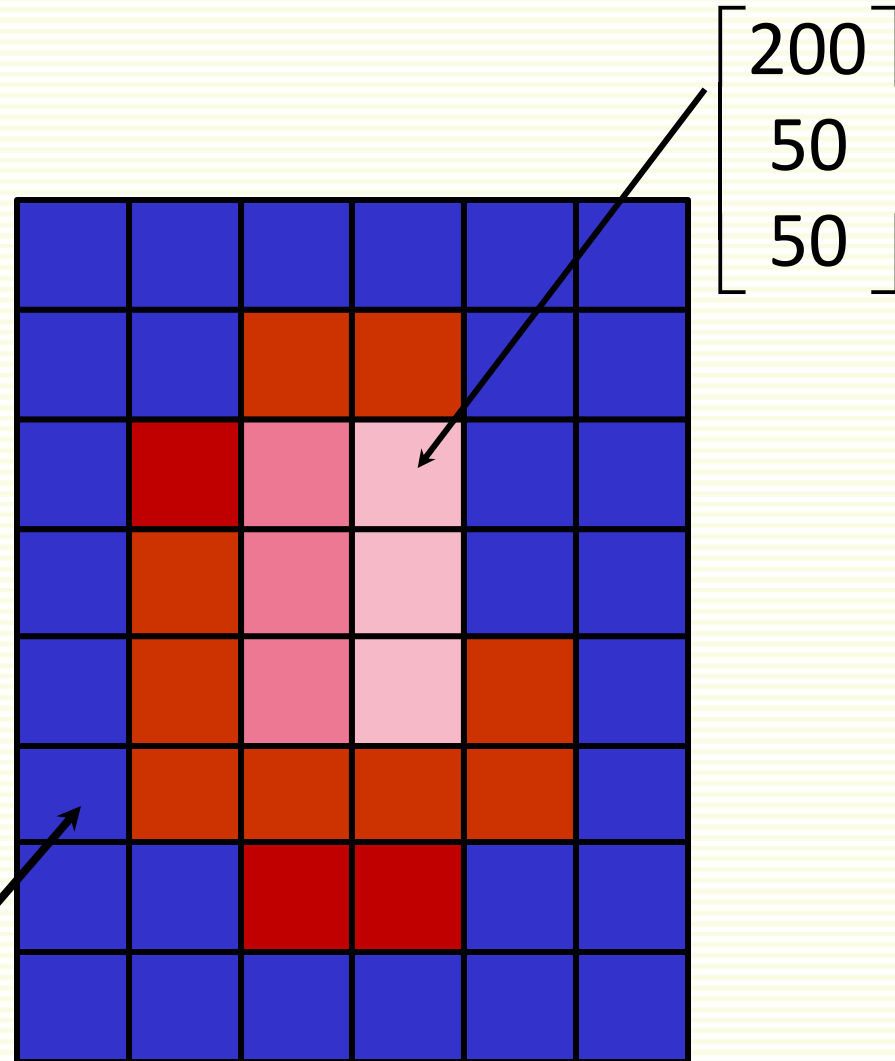


# Digital Color Image

- Color image is three functions pasted together
- Write this as a vector-valued function:



$$f(x,y) = \begin{bmatrix} r(x,y) \\ g(x,y) \\ b(x,y) \end{bmatrix}$$

$$\begin{bmatrix} 0 \\ 10 \\ 120 \end{bmatrix}$$



# Digital Color Image


- Can consider color image as 3 separate images: R, G, B



R



G



B

# Image filtering

- Given  $f(x,y)$  filtering computes a new image  $h(x,y)$
- As a function of local neighborhood at each position  $(x,y)$

$$h(x,y) = f(x,y) + f(x-1,y) \times f(x,y-1)$$

- Linear filtering: function is a weighted sum (or difference) of pixel values

$$h(x,y) = f(x,y) + 2 \times f(x-1,y-1) - 3 \times f(x+1,y+1)$$

- Many applications:
  - Enhance images
    - denoise, resize, increase contrast, ...
  - Extract information from images
    - Texture, edges, distinctive points ...
  - Detect patterns
    - Template matching

1	2	4	2	8
9	2	2	7	5
2	8	1	3	9
4	3	2	7	2
2	2	2	6	1
8	3	2	5	4




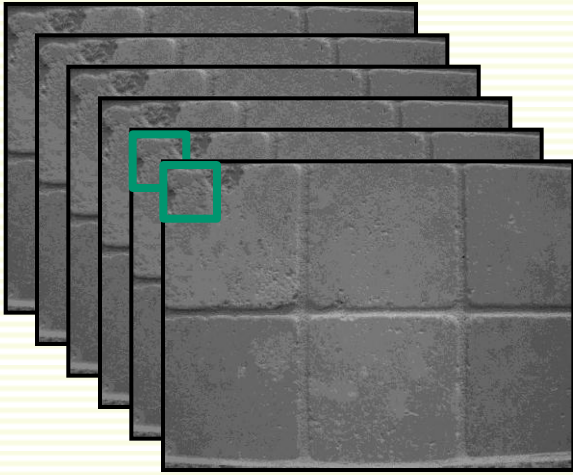
$$h(1,3) = 3 + 4 \times 8 = 35$$

$$h(4,5) = 4 + 5 \times 1 = 9$$

$$h(3,1) = 7 + 2 \times 4 - 3 \times 9 = -12$$

# Filtering for Noise Reduction: Motivation

- Multiple images of even the **same static scene** are not identical






# Common Types of Noise




**original image**



**Impulse noise:** random occurrences of white pixels





**Salt and pepper noise:** random occurrences of black and white pixels




**Gaussian noise:** variations in intensity drawn from a Gaussian distribution





# Gaussian Noise Most Commonly Assumed




**original image**



**$G(0,25)$  noise**



# Noise Reduction




- Noise can be reduced by averaging
- If we had multiple images, simply average them:

$$f_{\text{final}}(x,y) = (f_1(x,y) + f_2(x,y) + \dots + f_n(x,y))/n$$

- **But usually there is only one image!**


# First Attempt at a Solution

- Replace each pixel with an average of all the values in its neighborhood
- Assumptions:
  - expect a pixel to have intensities similar to its neighbors
  - Noise is independent at each pixel




# First attempt at a solution

- Replace each pixel with an average of all the values in its neighborhood (= 5 pixels, say)
- Moving average in 1D:




# First attempt at a solution

- Replace each pixel with an average of all the values in its neighborhood (= 5 pixels, say)
- Moving average in 1D:



# First attempt at a solution


- Replace each pixel with an average of all the values in its neighborhood (= 5 pixels, say)
- Moving average in 1D:






# First attempt at a solution

- Replace each pixel with an average of all the values in its neighborhood (= 5 pixels, say)
- Moving average in 1D:



# First attempt at a solution

- Replace each pixel with an average of all the values in its neighborhood (= 5 pixels, say)
- Moving average in 1D:













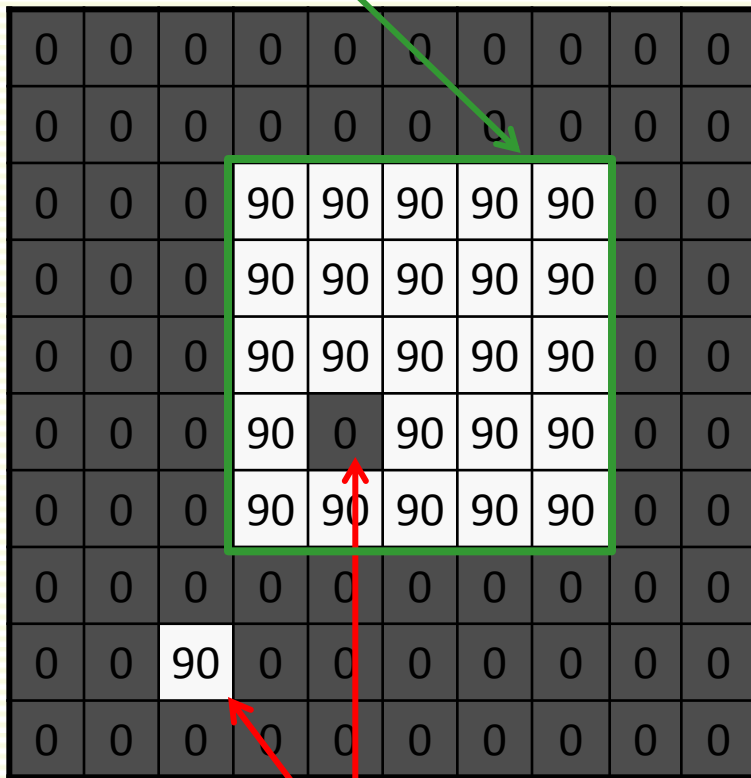




# Moving Average In 2D

$f(x,y)$


sharp border



sticking out

$g(x,y)$


border washed out



not sticking out

# Correlation Filtering

- Write as equation, averaging window  $(2k+1) \times (2k+1)$



$$g(x,y) = \frac{1}{(2k+1)^2} \sum_{u=-k}^k \sum_{v=-k}^k f(x+u, y+v)$$

normalizing factor

loop over all pixels in neighborhood around pixel  $f(i,j)$

$$g(x,y) = \sum_{u=-k}^k \sum_{v=-k}^k \frac{1}{(2k+1)^2} f(x+u, y+v)$$

uniform weight for each pixel

# Correlation Filtering

$$g(x, y) = \sum_{u=-k}^k \sum_{v=-k}^k \frac{1}{(2k+1)^2} f(x+u, y+v)$$

uniform weight for  
each pixel

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

$H[u, v]$

- Generalize by allowing different weights for different pixels in the neighborhood

$$g(x, y) = \sum_{u=-k}^k \sum_{v=-k}^k H[u, v] f(x+u, y+v)$$

non-uniform weight  
for each pixel

1/3	1/9	1/3
1/9	1/4	1/9
1/3	1/9	1/3

$H[u, v]$

# Correlation filtering

$$g(x, y) = \sum_{u=-k}^k \sum_{v=-k}^k H[u, v] f(x + u, y + v)$$

- This is called **cross-correlation**, denoted  $g = H \otimes f$
- Filtering an image: replace each pixel with a linear combination of its neighbors
- The filter **kernel** or **mask**  $H$  gives the weights in linear combination

# Averaging Filter

- What is kernel  $H$  for the moving average example?

$f(x,y)$

0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0
0	0	0	90	90	90	90	90	0	0
0	0	0	90	90	90	90	90	0	0
0	0	0	90	90	90	90	90	0	0
0	0	0	90	0	90	90	90	0	0
0	0	0	90	90	90	90	90	0	0
0	0	0	0	0	0	0	0	0	0
0	0	90	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0

$H[u,v]$

1	1	1
1	1	1
1	1	1

$\frac{1}{9}$

box filter


$g(x,y)$

	0	10	20	30	30				

$$g = H \otimes f$$

# Smoothing by Averaging

- Pictorial representation of box filter: 
  - white means large value, black means low value



original



filtered

- What if the mask is larger than 3x3 ?




# Effect of Average Filter


**Gaussian noise**

**Salt and Pepper noise**


**7 × 7**



**9 × 9**



**11 × 11**



# Gaussian Filter

- Nearest neighboring pixels to have the most influence
  - helps to lessen the effect of boundary smoothing

$f(x,y)$


0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0
0	0	0	90	90	90	90	90	0	0
0	0	0	90	90	90	90	90	0	0
0	0	0	90	90	90	90	90	0	0
0	0	0	90	0	90	90	90	0	0
0	0	0	90	90	90	90	90	0	0
0	0	0	0	0	0	0	0	0	0
0	0	90	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0

$H[u,v]$

$\frac{1}{16}$	1	2	1
	2	4	2
	1	2	1

This kernel  $H$  is an approximation of a 2d Gaussian function:



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



# Gaussian Filters: Mask Size



- Gaussian has infinite domain, discrete filters use finite mask
  - set mask size to exclude non-useful (effectively zero) weights

$\sigma = 5$  with 30 x 30 mask



blue weights  
are so small  
they are  
effectively 0



$\sigma = 5$  with 10 x 10 mask





# Gaussian filters: Variance

- Variance ( $\sigma$ ) contributes to the extent of smoothing
  - larger  $\sigma$  gives less rapidly decreasing weights
  - can construct a larger mask with non-negligible weights



$\sigma = 2$  with 30 x 30 kernel




$\sigma = 5$  with 30 x 30 kernel



$\sigma = 8$  with 30 x 30 kernel



# Matlab




**im**

**outim**

```
>> hsize = 10;  
>> sigma = 5;  
>> h = fspecial('gaussian', hsize, sigma);
```

```
>> mesh(h),
```




```
>> imagesc(h);
```




```
>> outim = imfilter(im, h); % correlation  
>> imshow(outim);
```

# Average vs. Gaussian Filter



mean filter



Gaussian filter



# More Average vs. Gaussian Filter

mean filter



Gaussian filter



$5 \times 5$



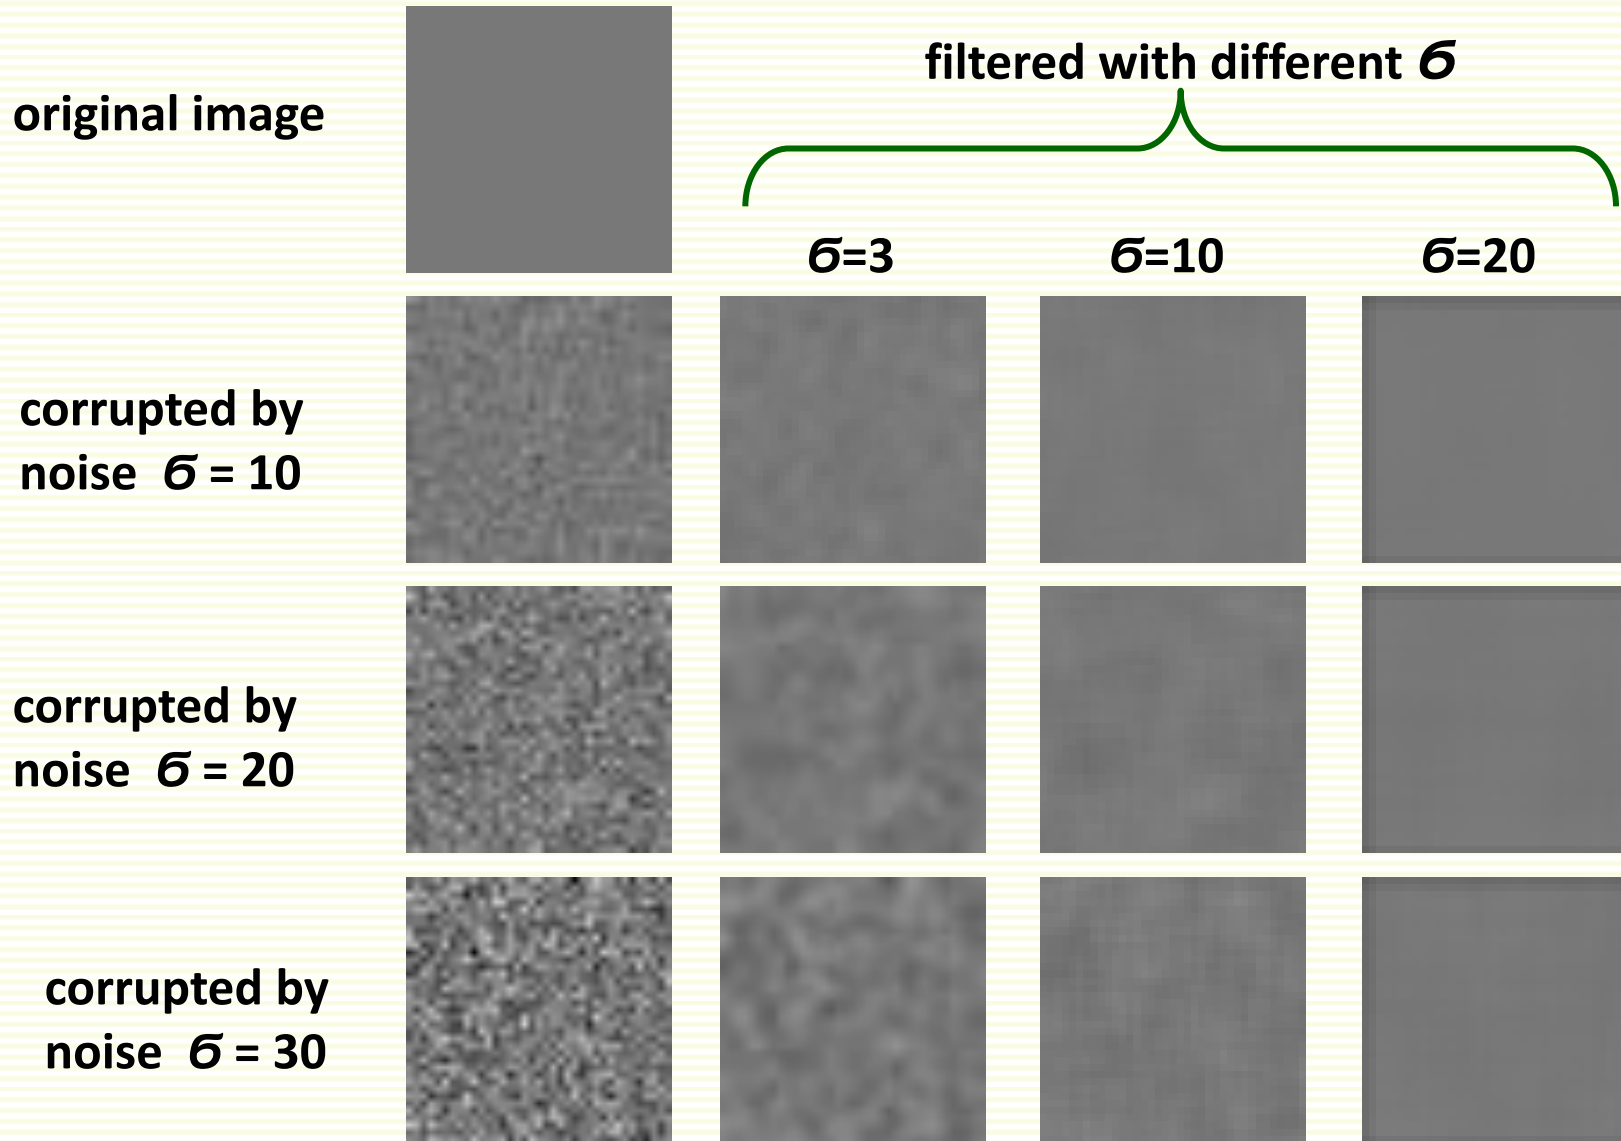
$15 \times 15$



$31 \times 31$




# Gaussian Filter with different $\sigma$






# Boundary Issues


- What is the size of the output?
- MATLAB: output size / “shape” options
  - *shape* = ‘full’: output size is sum of sizes of  $f$  and  $g$
  - *shape* = ‘same’: output size is same as  $f$
  - *shape* = ‘valid’: output size is difference of sizes of  $f$  and  $g$



full




same



valid

# Boundary issues



- What about near the edge?
  - the filter window falls off the edge of the image
  - need to extrapolate image




clip filter (black)



copy edge



wrap around



reflect across edge

# Properties of Smoothing Filters

- Values positive
- Sum to 1
  - constant regions same as input
  - overall image brightness stays unchanged
- Amount of smoothing proportional to mask size
  - larger mask means more extensive smoothing

# Filtering an Impulse Signal

- What is the result of filtering the impulse signal (image) with arbitrary kernel  $H$ ?

0	0	0	0	0	0	0
0	0	0	0	0	0	0
0	0	0	0	0	0	0
0	0	0	1	0	0	0
0	0	0	0	0	0	0
0	0	0	0	0	0	0
0	0	0	0	0	0	0

$f(x,y)$

$\otimes$

a	b	c
d	e	f
g	h	i

$H[u,v]$

=


$g(x,y)=?$

# Filtering an Impulse Signal

- What is the result of filtering the impulse signal (image) with arbitrary kernel  $H$ ?

0	0	0	0	0	0	0
0	0	0	0	0	0	0
0	0	0	0	0	0	0
0	0	0	1	0	0	0
0	0	0	0	0	0	0
0	0	0	0	0	0	0
0	0	0	0	0	0	0

$f(x,y)$

$\otimes$

a	b	c
d	e	f
g	h	i

$H[u,v]$

=


		i	h	g		
		f	e	d		
		c	b	a		

$g(x,y)=?$

# Convolution

- **Convolution:**

- Flip the mask in both dimensions
  - bottom to top, right to left
- Then apply cross-correlation



$$g(x,y) = \sum_{u=-k}^k \sum_{v=-k}^k H[u,v] f(x-u, y-v)$$




***H***



***H***

flipped



***f***

- Notation for convolution:  $g = H * f$

# Convolution vs. Correlation

- Convolution:  $g = H * f$

$$g(x, y) = \sum_{u=-k}^k \sum_{v=-k}^k H[u, v] f(x - u, y - v)$$

- Correlation:  $g = H \otimes f$

$$g(x, y) = \sum_{u=-k}^k \sum_{v=-k}^k H[u, v] f(x + u, y + v)$$

- For Gaussian or box filter, how the outputs differ?
- If the input is an impulse signal, how the outputs differ?



# Practice with Correlation Filtering



original



0	0	0
0	1	0
0	0	0

= ?

# Practice with Correlation Filtering



original



0	0	0
0	1	0
0	0	0



filtered (no change)

# Practice with Correlation Filtering



original



0	0	0
0	0	1
0	0	0

= ?

# Practice with Correlation Filtering



original



0	0	0
0	0	1
0	0	0



shifted left  
by 1 pixel with  
correlation

# Practice with Correlation Filtering



Original



$\frac{1}{9}$

1	1	1
1	1	1
1	1	1

= ?

# Practice with Correlation Filtering




original



$\frac{1}{9}$

1	1	1
1	1	1
1	1	1



blur (with a box filter)

# Practice with Correlation Filtering

apply one mask  
after the other,  
or subtract masks  
and apply one  
resulting mask

-1/9	-1/9	-1/9
-1/9	17/9	-1/9
-1/9	-1/9	-1/9



original



0	0	0
0	2	0
0	0	0

-

$\frac{1}{9}$

1	1	1
1	1	1
1	1	1

=

?



# Practice with Correlation Filtering



original



0	0	0
0	2	0
0	0	0



$\frac{1}{9}$


1	1	1
1	1	1
1	1	1



sharpened

# Practice with Correlation Filtering

- Why sharpens?

 $\otimes$ 

0	0	0
0	2	0
0	0	0

 -  $\frac{1}{9}$ 

1	1	1
1	1	1
1	1	1




original  $f$

+



original  $f$

-




smoothed

=



original  $f$

+




detail

=




sharpened

# Sharpening Example



before



after

# Separability

- Sometimes filter is separable, can split into two steps:
  - Convolve all rows with 1D filter
  - Convolve all columns with 1D filter
- Both box and Gaussian filters are separable
- Great for efficiency!

# Box Filter

$$\begin{array}{|c|c|c|} \hline 1/9 & 1/9 & 1/9 \\ \hline 1/9 & 1/9 & 1/9 \\ \hline 1/9 & 1/9 & 1/9 \\ \hline \end{array} = \begin{array}{|c|} \hline 1/3 \\ \hline 1/3 \\ \hline 1/3 \\ \hline \end{array} * \begin{array}{|c|c|c|} \hline 1/3 & 1/3 & 1/3 \\ \hline \end{array}$$

$H$                        $H_c$                        $H_r$

0	0	0	0	0	0
0	90	90	90	90	0
0	90	90	90	90	0
0	90	90	90	90	0
0	90	90	90	90	0
0	90	90	90	90	0
0	0	0	0	0	0

$*H =$

0	0	0	0	0	0
0	40	60	60	40	0
0	60	90	90	60	0
0	60	90	90	60	0
0	40	60	60	40	0
0	0	0	0	0	0

0	0	0	0	0	0
0	90	90	90	90	0
0	90	90	90	90	0
0	90	90	90	90	0
0	90	90	90	90	0
0	90	90	90	90	0
0	0	0	0	0	0

$*H_c * H_r =$

0	0	0	0	0	0
0	60	60	60	60	0
0	90	90	90	90	0
0	90	90	90	90	0
0	60	60	60	60	0
0	0	0	0	0	0

$*H_r =$

0	0	0	0	0	0
0	40	60	60	40	0
0	60	90	90	60	0
0	60	90	90	60	0
0	40	60	60	40	0
0	0	0	0	0	0

# Gaussian Filter: Example

- To convolve image with this:

$$\frac{1}{115}$$

2	4	5	4	2
4	9	12	9	4
5	12	15	12	5
4	9	12	9	4
2	4	5	4	2

$H$

- First convolve each row with:

$$\frac{1}{10.7}$$

1.3	3.2	3.8	3.2	1.3
-----	-----	-----	-----	-----

$H_r$

- Then each column with:

$$\frac{1}{10.7}$$


1.3	3.2	3.8	3.2	1.3
-----	-----	-----	-----	-----

$H_c$

# Gaussian Filter: Example

- Straightforward convolution with  $5 \times 5$  kernel
  - 25 multiplications, 24 additions per pixel
- Smart convolution
  - 10 multiplications, 9 additions per pixel
- Savings are even larger for larger kernels
  - for  $n \times n$  kernel, straightforward convolution is  $O(n^2)$
  - Smart convolution is  $O(n)$  per pixel

# Median Filters



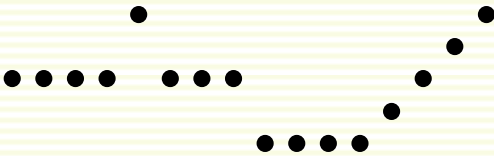
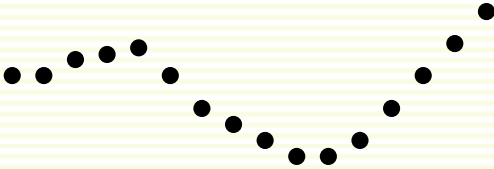
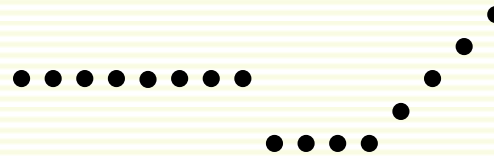
Median of  $\{1,2,25,3,24,22,20,21,23\} = \{1,2,3,20,21,22,23,24,25\}$  is 21

- A **Median Filter** selects median intensity in the window
- No new intensities are introduced
- Median filter preserves sharp details better than mean filter, it is not so prone to oversmoothing
- Better for salt and pepper, impulse (spiky) noise
- Is a median filter a kind of convolution?




# Median Filter

- Median filter is edge preserving

input:	 A plot of an input signal. It consists of a horizontal line of 10 dots. The first 5 dots are at a low level, the 6th dot is at a high level, and the next 4 dots are at a medium level. This is followed by a gap, then a horizontal line of 4 dots at a low level, and finally a diagonal line of 4 dots increasing from low to high.
average:	 A plot of the average of the input signal. The dots form a smooth, continuous curve that follows the general shape of the input signal but without the sharp spike or the distinct steps.
median:	 A plot of the median of the input signal. The dots form a step function that closely matches the original input signal, with the spike and the distinct levels preserved.


# Median filter

**Salt and pepper noise**



**row of noisy image**

**median filtered**



**row of filtered image**

# Comparison: Salt and Pepper Noise Image


Gaussian filter

median filter


$3 \times 3$



$5 \times 5$



$7 \times 7$




# Comparison: Gaussian Noise Image


Gaussian filter

median filter


$3 \times 3$



$5 \times 5$



$7 \times 7$





# Filtering Fun: Face of Faces





Salvador Dalí, *Gala Contemplating the Mediterranean Sea, which at 30 meters becomes the portrait of Abraham Lincoln*, 1976

# Summary

- Image “noise”
- Linear filters and convolution useful for
  - Enhancing images (smoothing, removing noise)
    - Box filter
    - Gaussian filter
    - Impact of scale / width of smoothing filter
  - Detecting features (next time)
- Separable filters more efficient
- Median filter: a non-linear filter, edge-preserving