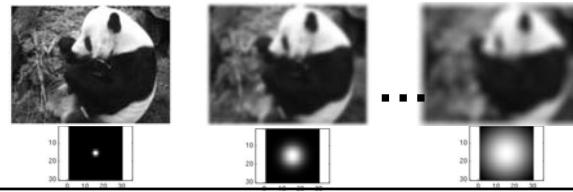


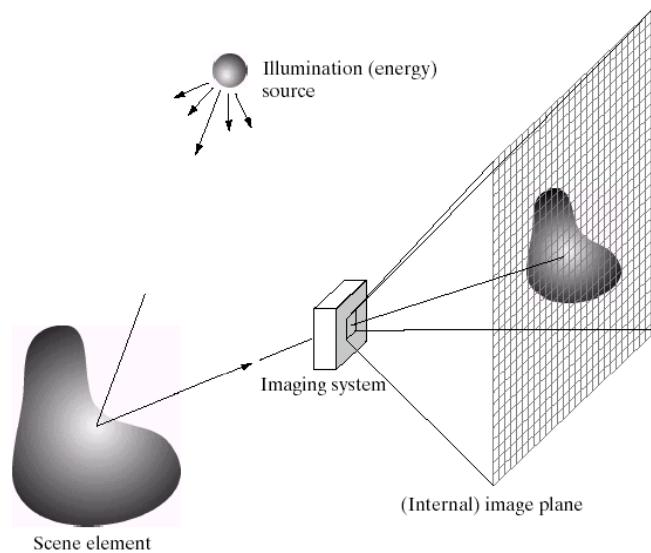
## Plan for today

- 1. Basics in feature extraction: filtering
- 2. Invariant local features
- 3. Specific object recognition methods

## Basics in feature extraction



## Image Formation



Slide credit: Derek Hoiem

## Digital camera

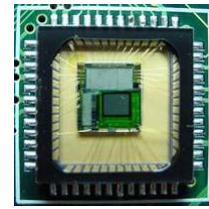
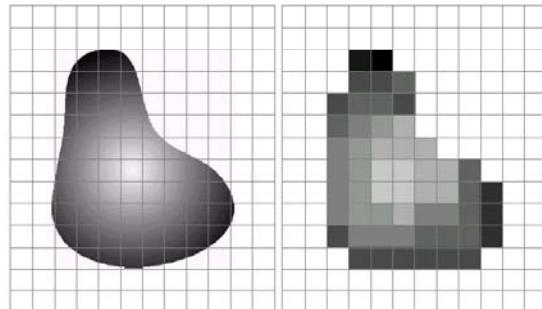


A digital camera replaces film with a sensor array

- Each cell in the array is light-sensitive diode that converts photons to electrons
- <http://electronics.howstuffworks.com/digital-camera.htm>

Slide by Steve Seitz

# Digital images

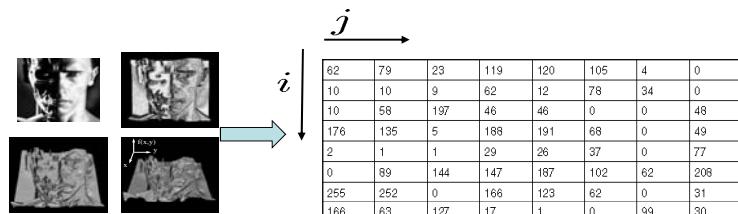
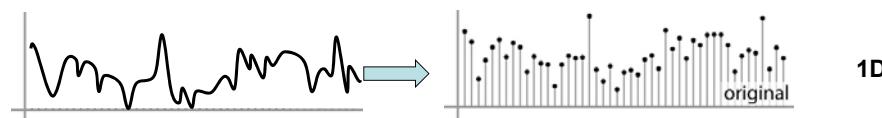
**a**

**FIGURE 2.17** (a) Continuous image projected onto a sensor array. (b) Result of image sampling and quantization.

Slide credit: Derek Hoiem

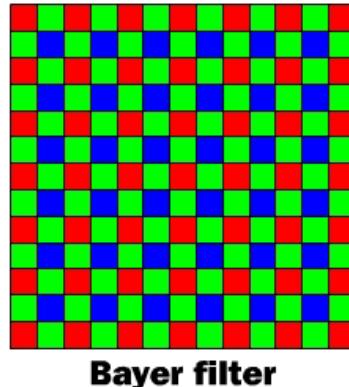
# Digital images

- **Sample** the 2D space on a regular grid
- **Quantize** each sample (round to nearest integer)
- Image thus represented as a matrix of integer values.

**2D****1D**

Adapted from S. Seitz

## Digital color images



**Bayer filter**

© 2000 How Stuff Works

## Digital color images

Color images,  
RGB color  
space



Kristen Grauman

## Images in Matlab

- Images represented as a matrix
- Suppose we have a NxM RGB image called “im”
  - $\text{im}(1, 1, 1)$  = top-left pixel value in R-channel
  - $\text{im}(y, x, b)$  = y pixels down, x pixels to right in the b<sup>th</sup> channel
  - $\text{im}(N, M, 3)$  = bottom-right pixel in B-channel
- `imread(filename)` returns a uint8 image (values 0 to 255)
  - Convert to double format (values 0 to 1) with `im2double`

row      column

		R	G	B						
0.92	0.93	0.94	0.97	0.62	0.37	0.85	0.97	0.93	0.92	0.99
0.95	0.89	0.82	0.89	0.56	0.31	0.75	0.92	0.81	0.95	0.91
0.89	0.72	0.51	0.55	0.51	0.42	0.57	0.41	0.49	0.91	0.92
0.96	0.95	0.88	0.94	0.56	0.46	0.91	0.87	0.90	0.97	0.95
0.71	0.81	0.81	0.87	0.57	0.37	0.80	0.88	0.89	0.79	0.85
0.49	0.62	0.60	0.58	0.50	0.60	0.58	0.50	0.61	0.45	0.33
0.86	0.84	0.74	0.58	0.51	0.39	0.73	0.92	0.91	0.49	0.74
0.96	0.67	0.54	0.85	0.48	0.37	0.88	0.90	0.94	0.82	0.93
0.69	0.49	0.56	0.66	0.43	0.42	0.77	0.73	0.71	0.90	0.99
0.79	0.73	0.90	0.67	0.33	0.61	0.69	0.79	0.73	0.93	0.97
0.91	0.94	0.89	0.49	0.41	0.78	0.78	0.77	0.89	0.99	0.93
0.79	0.73	0.90	0.67	0.33	0.61	0.69	0.79	0.73	0.93	0.97
0.91	0.94	0.89	0.49	0.41	0.78	0.78	0.77	0.89	0.99	0.93
0.79	0.73	0.90	0.67	0.33	0.61	0.69	0.79	0.73	0.93	0.97
0.91	0.94	0.89	0.49	0.41	0.78	0.78	0.77	0.89	0.99	0.93

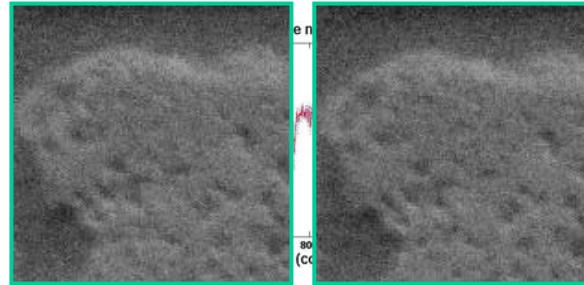
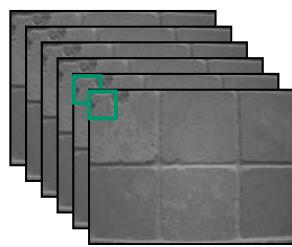
Slide credit: Derek Hoiem

## Main idea: image filtering

- Compute a function of the local neighborhood at each pixel in the image
  - Function specified by a “filter” or mask saying how to combine values from neighbors.
- Uses of filtering:
  - Enhance an image (denoise, resize, etc)
  - Extract information (texture, edges, etc)
  - Detect patterns (template matching)

Adapted from Derek Hoiem

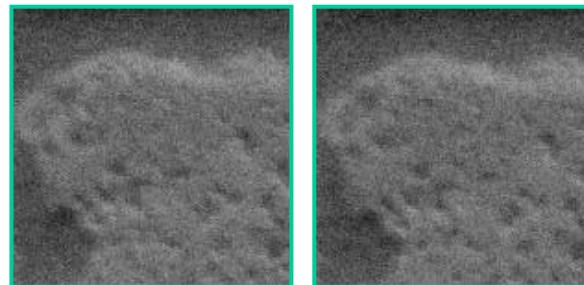
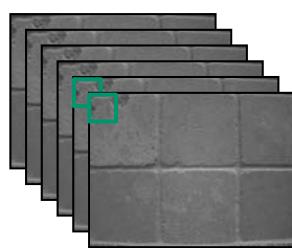
## Motivation: noise reduction



- Even multiple images of the **same static scene** will not be identical.

Kristen Grauman

## Motivation: noise reduction



- Even multiple images of the same static scene will not be identical.
- How could we reduce the noise, i.e., give an estimate of the true intensities?
- **What if there's only one image?**

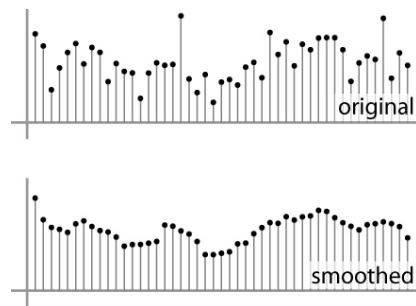
Kristen Grauman

## First attempt at a solution

- Let's replace each pixel with an average of all the values in its neighborhood
- Assumptions:
  - Expect pixels to be like their neighbors
  - Expect noise processes to be independent from pixel to pixel

## First attempt at a solution

- Let's replace each pixel with an average of all the values in its neighborhood
- Moving average in 1D:

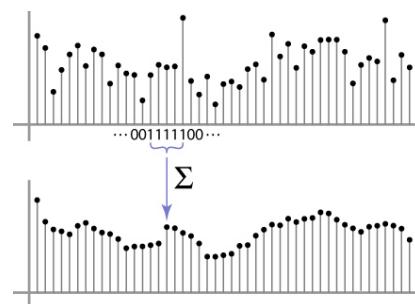


Source: S. Marschner

## Weighted Moving Average

Can add weights to our moving average

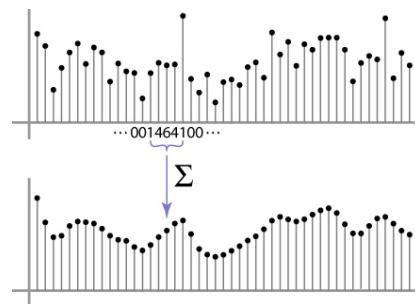
*Weights* [1, 1, 1, 1, 1] / 5



Source: S. Marschner

## Weighted Moving Average

Non-uniform weights [1, 4, 6, 4, 1] / 16



Source: S. Marschner

## Moving Average In 2D

$F[x, y]$

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

$G[x, y]$

0										

Source: S. Seitz

## Moving Average In 2D

$F[x, y]$

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

$G[x, y]$


Source: S. Seitz

## Moving Average In 2D

$F[x, y]$

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	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	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

$G[x, y]$


Source: S. Seitz

## Moving Average In 2D

$F[x, y]$

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	90	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

$G[x, y]$


Source: S. Seitz

## Moving Average In 2D

$F[x, y]$

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

$G[x, y]$

	0	10	20	30	30	30	20	10		
	0	20	40	60	60	60	40	20		
	0	30	60	90	90	90	60	30		
	0	30	50	80	80	90	60	30		
	0	30	50	80	80	90	60	30		
	0	20	30	50	50	60	40	20		
	10	20	30	30	30	30	20	10		
	10	10	10	0	0	0	0	0		

Source: S. Seitz

## Moving Average In 2D

$F[x, y]$

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

$G[x, y]$

	0	10	20	30	30	30	20	10		
	0	20	40	60	60	60	40	20		
	0	30	60	90	90	90	60	30		
	0	30	50	80	80	90	60	30		
	0	30	50	80	80	90	60	30		
	0	20	30	50	50	60	40	20		
	10	20	30	30	30	30	20	10		
	10	10	10	0	0	0	0	0		

Source: S. Seitz

## Correlation filtering

Say the averaging window size is  $2k+1 \times 2k+1$ :

$$G[i, j] = \frac{1}{(2k+1)^2} \underbrace{\sum_{u=-k}^k \sum_{v=-k}^k}_{\substack{\text{Attribute uniform} \\ \text{weight to each pixel}}} F[i + u, j + v] \underbrace{\quad}_{\substack{\text{Loop over all pixels in neighborhood} \\ \text{around image pixel } F[i,j]}}$$

Now generalize to allow **different weights** depending on neighboring pixel's relative position:

$$G[i, j] = \sum_{u=-k}^k \sum_{v=-k}^k H[u, v] \underbrace{F[i + u, j + v]}_{\text{Non-uniform weights}}$$

## Correlation filtering

$$G[i, j] = \sum_{u=-k}^k \sum_{v=-k}^k H[u, v] F[i + u, j + v]$$

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**”  $H[u, v]$  is the prescription for the weights in the linear combination.

## Averaging filter

- What values belong in the kernel  $H$  for the moving average example?

$$\begin{array}{c}
 F[x, y] \quad \otimes \quad H[u, v] \quad G[x, y] \\
 \begin{array}{|c|c|c|c|c|c|c|c|c|c|} \hline
 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ \hline
 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ \hline
 0 & 0 & 0 & 90 & 90 & 90 & 90 & 90 & 0 & 0 \\ \hline
 0 & 0 & 0 & 90 & 90 & 90 & 90 & 90 & 0 & 0 \\ \hline
 0 & 0 & 0 & 90 & 90 & 90 & 90 & 90 & 0 & 0 \\ \hline
 0 & 0 & 0 & 90 & 90 & 90 & 90 & 90 & 0 & 0 \\ \hline
 0 & 0 & 0 & 90 & 90 & 90 & 90 & 90 & 0 & 0 \\ \hline
 0 & 0 & 0 & 90 & 90 & 90 & 90 & 90 & 0 & 0 \\ \hline
 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ \hline
 0 & 0 & 90 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ \hline
 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ \hline
 \end{array} \quad \frac{1}{9} \quad \begin{array}{|c|c|c|} \hline
 1 & 1 & 1 \\ \hline
 1 & ? & 1 \\ \hline
 1 & 1 & 1 \\ \hline
 \end{array} \quad \begin{array}{|c|c|c|c|c|c|c|c|c|c|} \hline
 0 & 10 & 20 & 30 & 30 \\ \hline
 \end{array} \\
 \text{"box filter"} \quad \quad \quad
 \end{array}$$

$$G = H \otimes F$$

## Smoothing by averaging

 depicts box filter:  
white = high value, black = low value



original



filtered

What if the filter size was 5 x 5 instead of 3 x 3?

## Gaussian filter

- What if we want nearest neighboring pixels to have the most influence on the output?

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

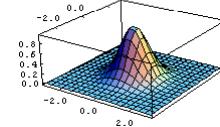
$F[x, y]$

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

$H[u, v]$

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

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



Source: S. Seitz

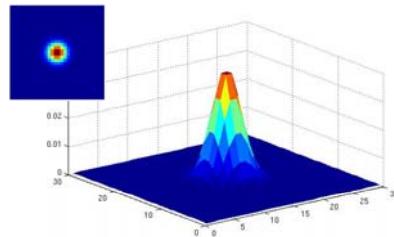
- Removes high-frequency components from the image (“low-pass filter”).

## Smoothing with a Gaussian

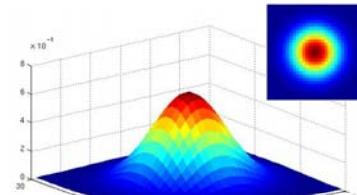


## Gaussian filters

- What parameters matter here?
- **Variance** of Gaussian: determines extent of smoothing



$\sigma = 2$  with  
30 x 30  
kernel

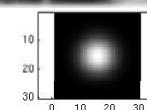
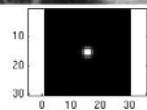


$\sigma = 5$  with  
30 x 30  
kernel

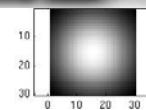
Kristen Grauman

## Smoothing with a Gaussian

Parameter  $\sigma$  is the “scale” / “width” / “spread” of the Gaussian kernel, and controls the amount of smoothing.



...



```
for sigma=1:3:10
    h = fspecial('gaussian', fsize, sigma);
    out = imfilter(im, h);
    imshow(out);
    pause;
end
```

Kristen Grauman

## Properties of smoothing filters

- Smoothing

- Values positive
- Sum to 1 → \_\_\_\_\_
- Amount of smoothing proportional to mask size
- Remove “high-frequency” components; “low-pass” filter

Kristen Grauman

## Predict the outputs using correlation filtering



$$\text{eye patch} * \begin{bmatrix} 0 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 0 \end{bmatrix} = ?$$



$$\text{eye patch} * \begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & 1 \\ 0 & 0 & 0 \end{bmatrix} = ?$$



$$\text{eye patch} * \begin{bmatrix} 0 & 0 & 0 \\ 0 & 2 & 0 \\ 0 & 0 & 0 \end{bmatrix} - \frac{1}{9} \begin{bmatrix} 1 & 1 & 1 \\ 1 & 1 & 1 \\ 1 & 1 & 1 \end{bmatrix} = ?$$

## Practice with linear filters



0	0	0
0	1	0
0	0	0

?

Original

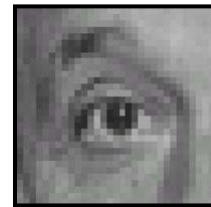
Source: D. Lowe

## Practice with linear filters



0	0	0
0	1	0
0	0	0

Original



Filtered  
(no change)

Source: D. Lowe

## Practice with linear filters



0	0	0
0	0	1
0	0	0

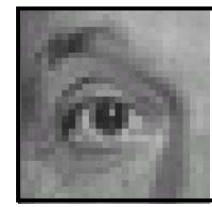
?

Source: D. Lowe

## Practice with linear filters



0	0	0
0	0	1
0	0	0



Original

Shifted left  
by 1 pixel  
with  
correlation

Source: D. Lowe

## Practice with linear filters



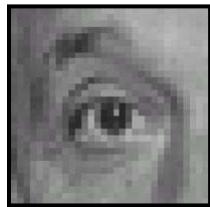
Original

$$\frac{1}{9} \begin{array}{|c|c|c|} \hline 1 & 1 & 1 \\ \hline 1 & 1 & 1 \\ \hline 1 & 1 & 1 \\ \hline \end{array}$$

?

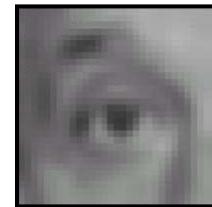
Source: D. Lowe

## Practice with linear filters



Original

$$\frac{1}{9} \begin{array}{|c|c|c|} \hline 1 & 1 & 1 \\ \hline 1 & 1 & 1 \\ \hline 1 & 1 & 1 \\ \hline \end{array}$$



Blur (with a  
box filter)

Source: D. Lowe

## Practice with linear filters



Original

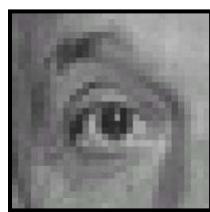
$$\begin{bmatrix} 0 & 0 & 0 \\ 0 & 2 & 0 \\ 0 & 0 & 0 \end{bmatrix}$$

$$- \frac{1}{9} \begin{bmatrix} 1 & 1 & 1 \\ 1 & 1 & 1 \\ 1 & 1 & 1 \end{bmatrix}$$

?

Source: D. Lowe

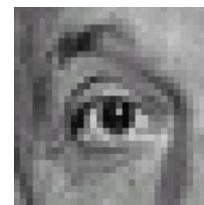
## Practice with linear filters



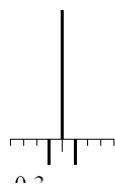
Original

$$\begin{bmatrix} 0 & 0 & 0 \\ 0 & 2 & 0 \\ 0 & 0 & 0 \end{bmatrix}$$

$$- \frac{1}{9} \begin{bmatrix} 1 & 1 & 1 \\ 1 & 1 & 1 \\ 1 & 1 & 1 \end{bmatrix}$$

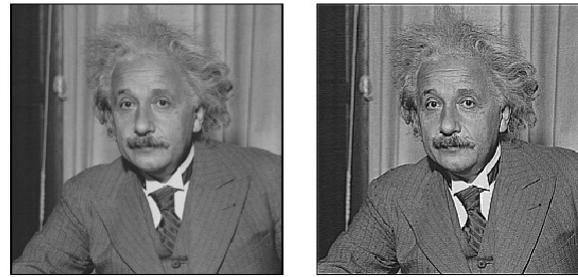


Sharpening filter:  
accentuates differences  
with local average



Source: D. Lowe

## Filtering examples: sharpening



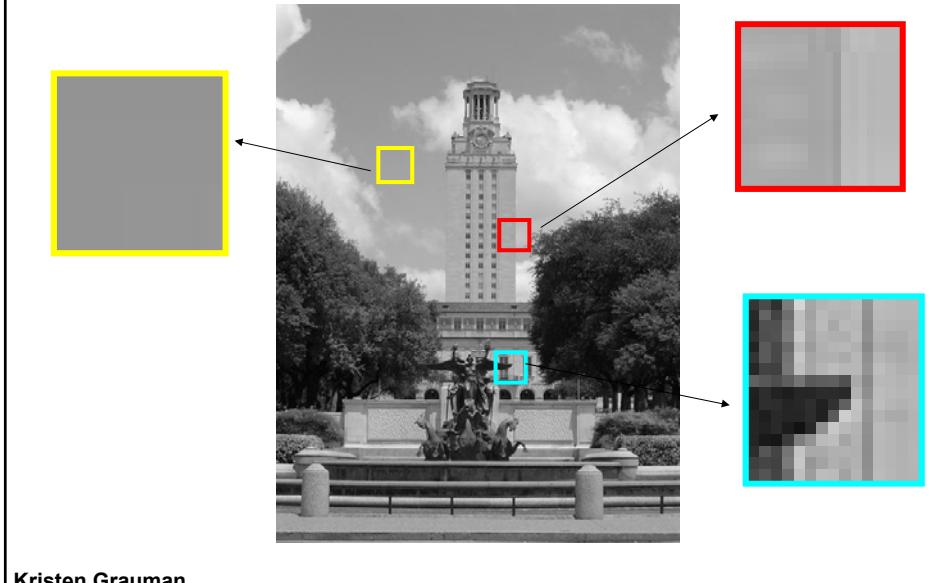
before

after

## Main idea: image filtering

- Compute a function of the local neighborhood at each pixel in the image
  - Function specified by a “filter” or mask saying how to combine values from neighbors.
- Uses of filtering:
  - Enhance an image (denoise, resize, etc)
  - Extract information (texture, edges, etc)
  - Detect patterns (template matching)

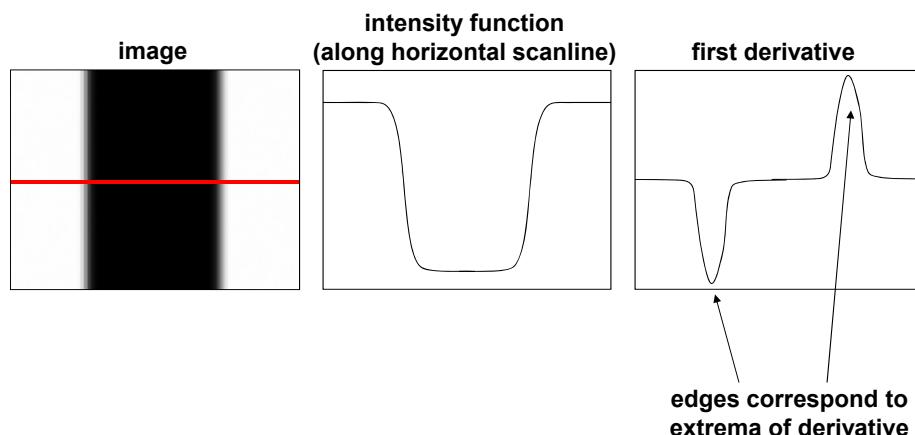
## Why are gradients important?



Kristen Grauman

## Derivatives and edges

An edge is a place of rapid change in the image intensity function.



Source: L. Lazebn

## Derivatives with convolution

For 2D function,  $f(x,y)$ , the partial derivative is:

$$\frac{\partial f(x, y)}{\partial x} = \lim_{\varepsilon \rightarrow 0} \frac{f(x + \varepsilon, y) - f(x, y)}{\varepsilon}$$

For discrete data, we can approximate using finite differences:

$$\frac{\partial f(x, y)}{\partial x} \approx \frac{f(x+1, y) - f(x, y)}{1}$$

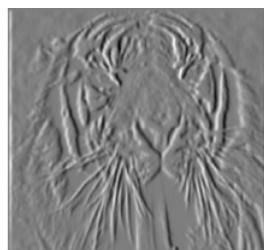
To implement above as convolution, what would be the associated filter?

Kristen Grauman

## Partial derivatives of an image

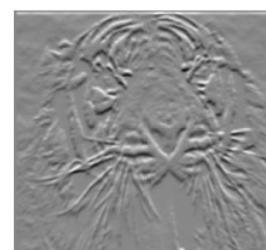
$$\frac{\partial f(x, y)}{\partial x}$$

**-1** **1**



$$\frac{\partial f(x, y)}{\partial y}$$

<b>-1</b>	<b>?</b>	<b>1</b>
<b>1</b>	or	<b>-1</b>



Which shows changes with respect to x?

Kristen Grauman

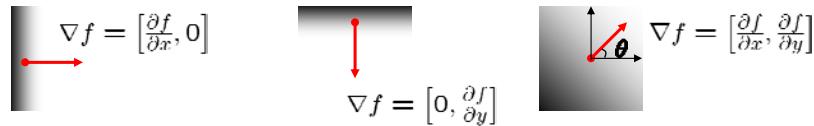
(showing filters for correlation)

## Image gradient

The gradient of an image:

$$\nabla f = \left[ \frac{\partial f}{\partial x}, \frac{\partial f}{\partial y} \right]$$

The gradient points in the direction of most rapid change in intensity



The **gradient direction** (orientation of edge normal) is given by:

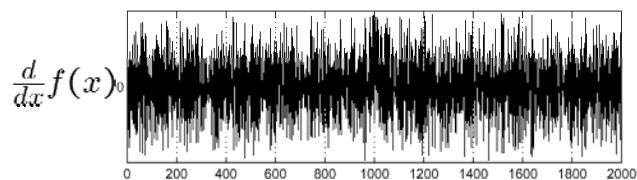
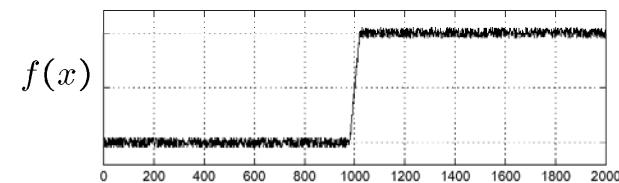
$$\theta = \tan^{-1} \left( \frac{\partial f}{\partial y} / \frac{\partial f}{\partial x} \right)$$

Slide credit Steve Seitz

## Effects of noise

Consider a single row or column of the image

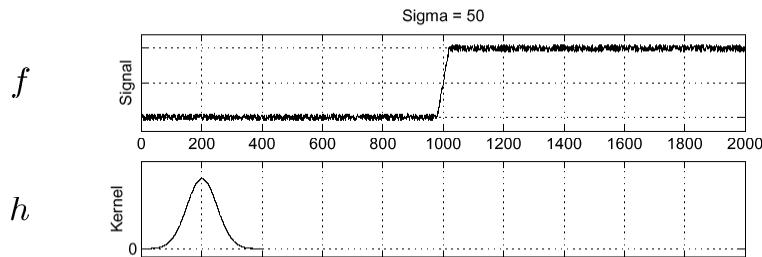
- Plotting intensity as a function of position gives a signal



Where is the edge?

Slide credit Steve Seitz

## Solution: smooth first

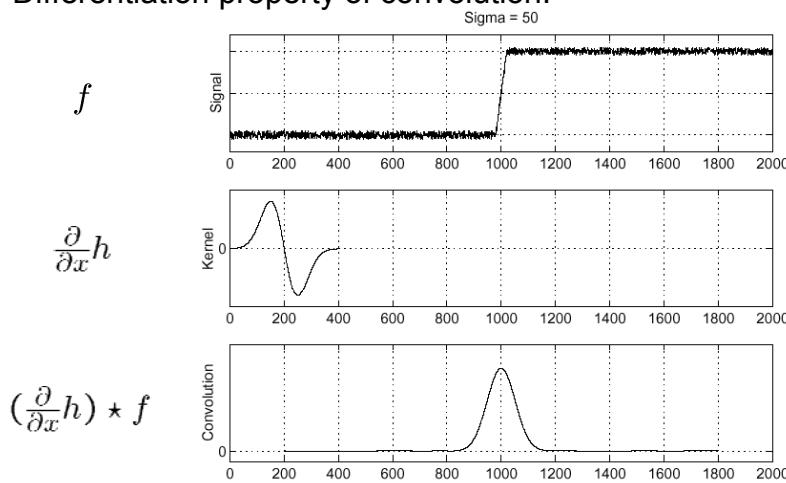


Where is the edge?      Look for peaks in  $\frac{\partial}{\partial x}(h * f)$

## Derivative theorem of convolution

$$\frac{\partial}{\partial x}(h * f) = (\frac{\partial}{\partial x}h) * f$$

Differentiation property of convolution.

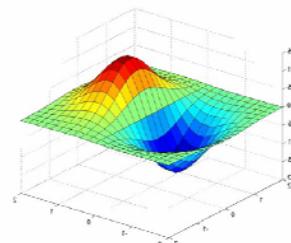


Slide credit Steve Seitz

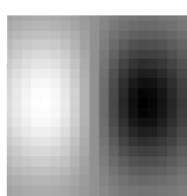
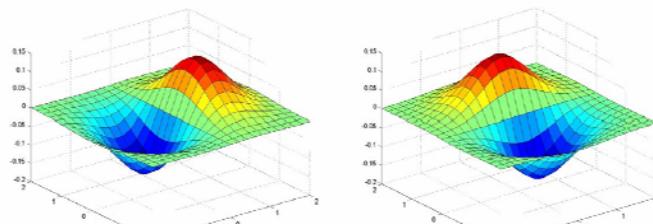
## Derivative of Gaussian filters

$$(I \otimes g) \otimes h = I \otimes (g \otimes h)$$

$$\begin{bmatrix} 0.0030 & 0.0133 & 0.0219 & 0.0133 & 0.0030 \\ 0.0133 & 0.0596 & 0.0983 & 0.0596 & 0.0133 \\ 0.0219 & 0.0983 & 0.1621 & 0.0983 & 0.0219 \\ 0.0133 & 0.0596 & 0.0983 & 0.0596 & 0.0133 \\ 0.0030 & 0.0133 & 0.0219 & 0.0133 & 0.0030 \end{bmatrix} \otimes \begin{bmatrix} 1 & -1 \end{bmatrix}$$



## Derivative of Gaussian filters

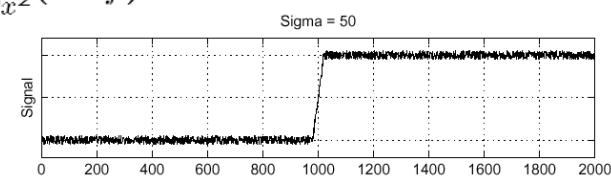


Source: L. Lazebnik

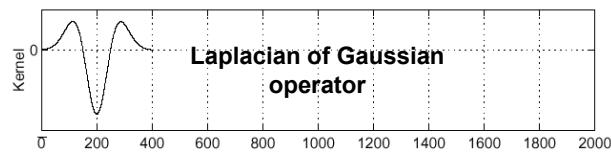
## Laplacian of Gaussian

Consider  $\frac{\partial^2}{\partial x^2}(h \star f)$

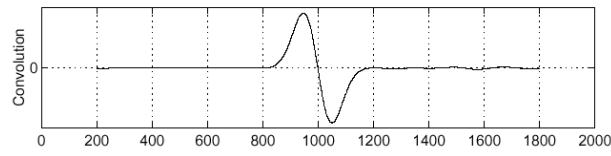
$f$



$\frac{\partial^2}{\partial x^2} h$



$(\frac{\partial^2}{\partial x^2} h) \star f$

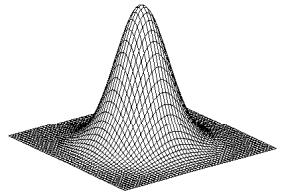


Where is the edge?

Zero-crossings of bottom graph

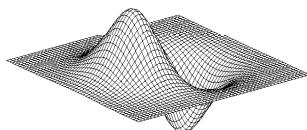
Slide credit: Steve Seitz

## 2D edge detection filters



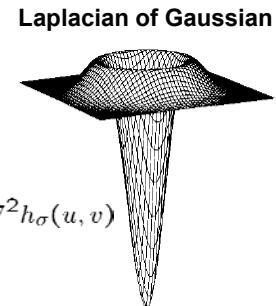
Gaussian

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



derivative of Gaussian

$$\frac{\partial}{\partial x} h_\sigma(u, v) \quad \nabla^2 h_\sigma(u, v)$$



Laplacian of Gaussian

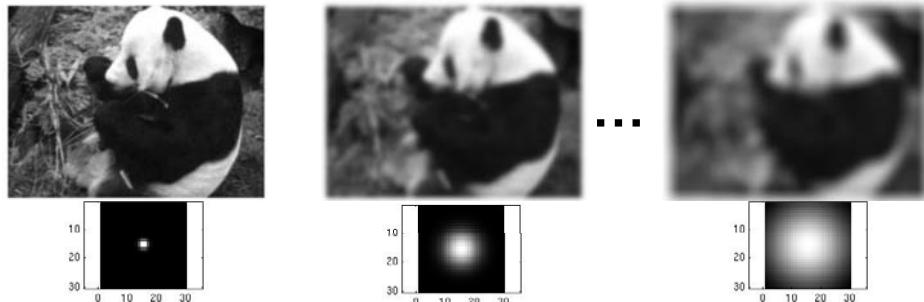
- $\nabla^2$  is the Laplacian operator:

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

Slide credit: Steve Seitz

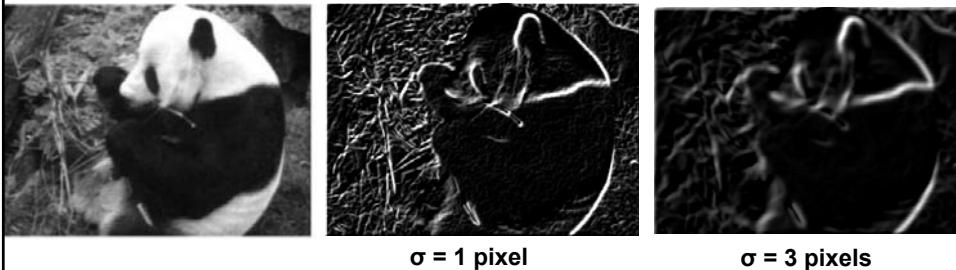
## Smoothing with a Gaussian

Recall: parameter  $\sigma$  is the “scale” / “width” / “spread” of the Gaussian kernel, and controls the amount of smoothing.



Kristen Grauman

## Effect of $\sigma$ on derivatives



The apparent structures differ depending on Gaussian's scale parameter.

Larger values: larger scale edges detected  
 Smaller values: finer features detected

Kristen Grauman

# Mask properties

- Smoothing
  - Values positive
  - Sum to 1 → constant regions same as input
  - Amount of smoothing proportional to mask size
  - Remove “high-frequency” components; “low-pass” filter
  
- Derivatives
  - \_\_\_\_\_ signs used to get high response in regions of high contrast
  - Sum to \_\_\_ → no response in constant regions
  - High absolute value at points of high contrast

Kristen Grauman

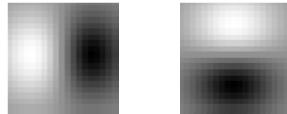
# Main idea: image filtering

- Compute a function of the local neighborhood at each pixel in the image
  - Function specified by a “filter” or mask saying how to combine values from neighbors.
  
- Uses of filtering:
  - Enhance an image (denoise, resize, etc)
  - Extract information (texture, edges, etc)
  - Detect patterns (template matching)

## Template matching

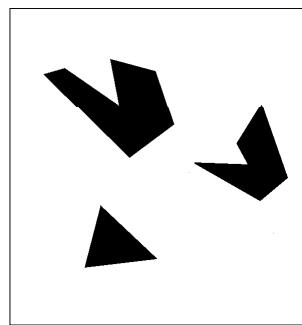
- Filters as **templates**:

Note that filters look like the effects they are intended to find --- “matched filters”



- Use normalized cross-correlation score to find a given pattern (template) in the image.
- Normalization needed to control for relative brightnesses.

## Template matching

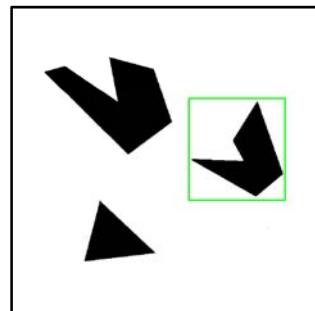


Template (mask)

Scene

### A toy example

## Template matching

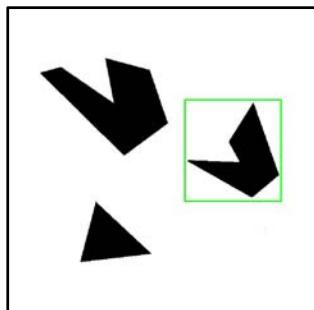


Detected template

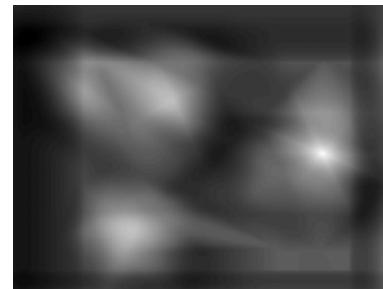


Template

## Template matching



Detected template



Correlation map

## Where's Waldo?

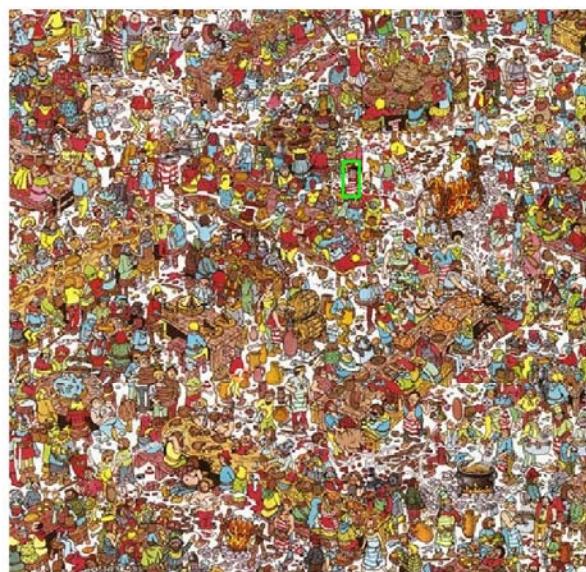


Scene



Template

## Where's Waldo?

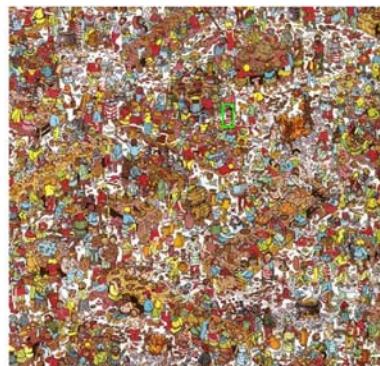


Detected template

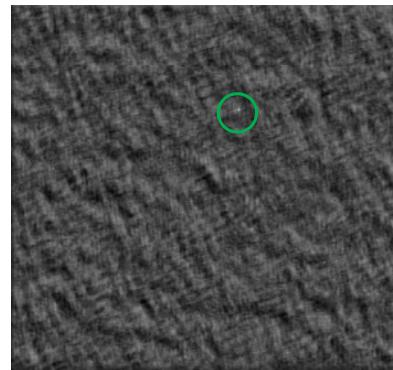


Template

## Where's Waldo?



Detected template



Correlation map

## Template matching



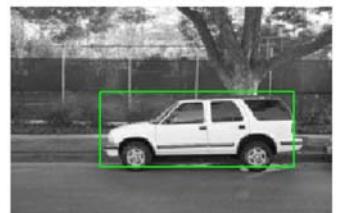
Scene



Template

What if the template is not identical to some subimage in the scene?

## Template matching



Detected template



Template

Match can be meaningful, if scale, orientation, and general appearance is right.

...but we can do better!...

## Summary so far

- Compute a function of the local neighborhood at each pixel in the image
  - Function specified by a “filter” or mask saying how to combine values from neighbors.
- Uses of filtering:
  - Enhance an image (denoise, resize, etc)
  - Extract information (texture, edges, etc)
  - Detect patterns (template matching)