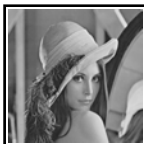


Image gradients and edges

Wed, Jan 26

Prof. Kristen Grauman

UT-Austin


Last time

- Various models for image “noise”
- Linear filters and convolution useful for
  - Image smoothing, removing noise
    - Box filter
    - Gaussian filter
    - Impact of scale / width of smoothing filter
- Separable filters more efficient
- Median filter: a non-linear filter, edge-preserving

Review

Filter f = 1/9 x [1 1 1 1 1 1 1 1 1]


original image h


filtered

Review

Filter f = 1/9 x $[1\ 1\ 1\ 1\ 1\ 1\ 1\ 1\ 1]^T$


original image h


filtered

Review

What happens if we have a smoothing filter that is *unnormalized* (does not sum to one)?

Recall: 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

Edge detection

- **Goal**: map image from 2d array of pixels to a set of curves or line segments or contours.
- **Why?**

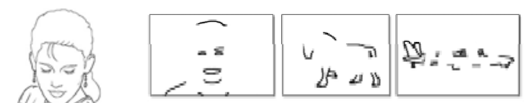

Figure from J. Shotton et al., PAMI 2007

- **Main idea**: look for strong gradients, post-process

What causes an edge?

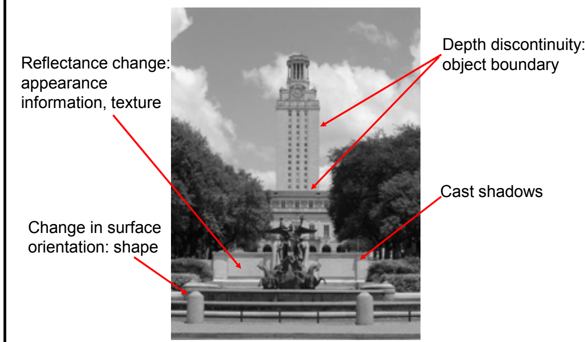

Kristen Grauman, UT-Austin

Edges/gradients and invariance

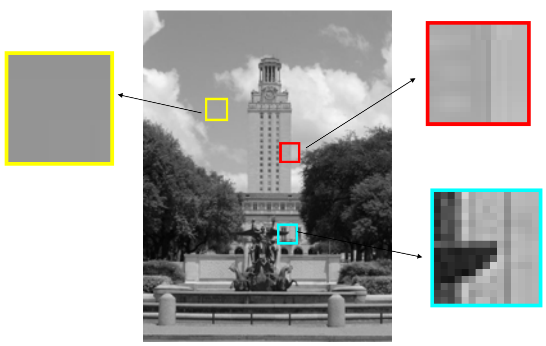

Kristen Grauman, UT-Austin

Derivatives and edges

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

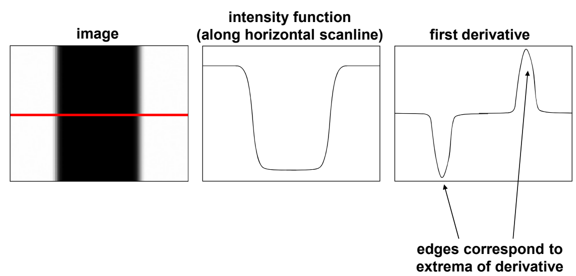

Source: L. Lazebnik

Derivatives with convolution

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

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

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?

Partial derivatives of an image

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

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

-1	1

-1
1

? or

1
-1

Which shows changes with respect to x?

(showing filters for correlation)

Assorted finite difference filters

Prewitt: $M_x = \begin{bmatrix} -1 & 0 & 1 \\ -1 & 0 & 1 \\ -1 & 0 & 1 \end{bmatrix}$; $M_y = \begin{bmatrix} 1 & 1 & 1 \\ 0 & 0 & 0 \\ -1 & -1 & -1 \end{bmatrix}$

Sobel: $M_x = \begin{bmatrix} -1 & 0 & 1 \\ -2 & 0 & 2 \\ -1 & 0 & 1 \end{bmatrix}$; $\boxed{M_y = \begin{bmatrix} 1 & 2 & 1 \\ 0 & 0 & 0 \\ -1 & -2 & -1 \end{bmatrix}}$

Roberts: $M_x = \begin{bmatrix} 0 & 1 \\ -1 & 0 \end{bmatrix}$; $M_y = \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix}$

```
>> My = fspecial('sobel');
>> outim = imfilter(double(im), My);
>> imagesc(outim);
>> colormap gray;
```

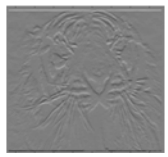

Image gradient

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

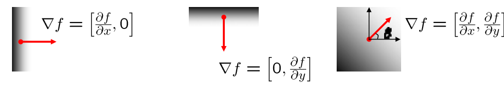

The **gradient direction** (orientation of edge normal) is given by:

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

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

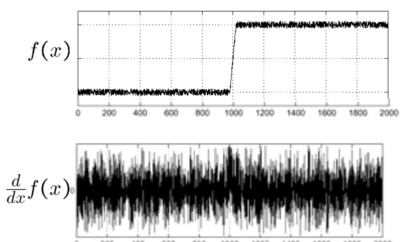

Where is the edge?

Slide credit Steve Seitz

Solution: smooth first

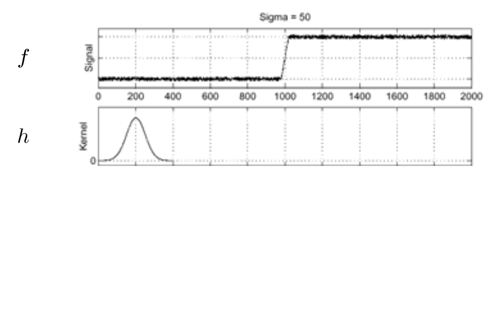

Where is the edge? Look for peaks in $\frac{\partial}{\partial x}(h \star f)$

Derivative theorem of convolution

$$\frac{\partial}{\partial x}(h \star f) = \left(\frac{\partial}{\partial x}h\right) \star f$$

Differentiation property of convolution.

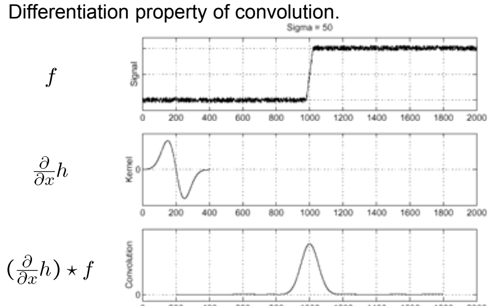

Slide credit Steve Seitz

Derivative of Gaussian filters

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

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

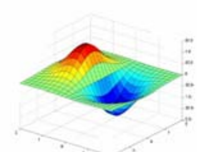

Derivative of Gaussian filters

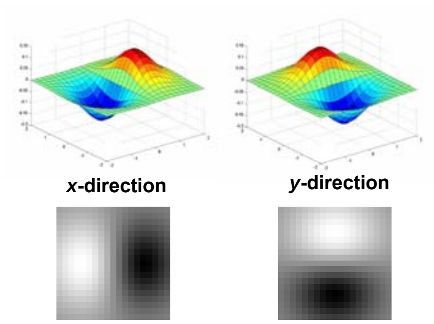

Source: L. Lazebnik

Laplacian of Gaussian

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

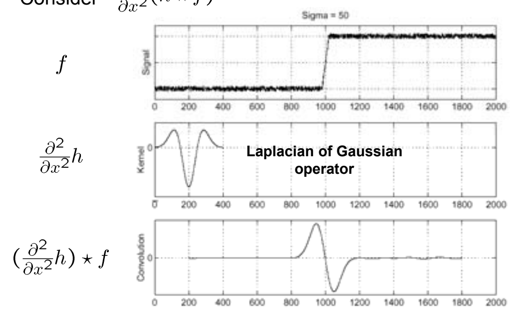

Where is the edge? Zero-crossings of bottom graph

Slide credit: Steve Seitz

2D edge detection filters

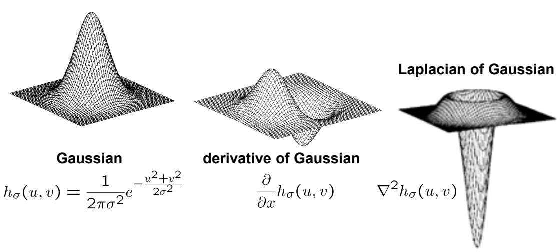

Gaussian $h_\sigma(u,v) = \frac{1}{2\pi\sigma^2} e^{-\frac{u^2+v^2}{2\sigma^2}}$

derivative of Gaussian $\frac{\partial}{\partial x} h_\sigma(u,v)$

$\nabla^2 h_\sigma(u,v)$

- ∇^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 σ is the "scale" / "width" / "spread" of the Gaussian kernel, and controls the amount of smoothing.

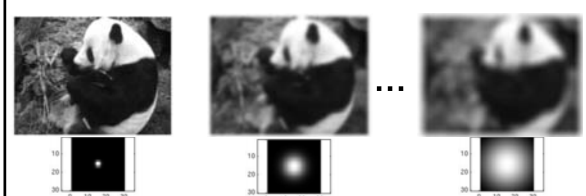

Kristen Grauman, UT-Austin

Effect of σ on derivatives

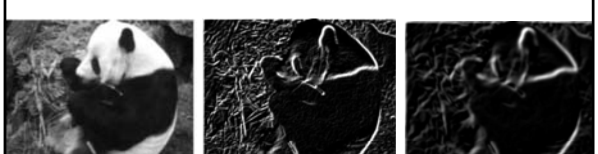

σ = 1 pixel **σ = 3 pixels**

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

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

Kristen Grauman, UT-Austin

So, what scale to choose?

It depends what we're looking for.

Kristen Grauman, UT-Austin

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

Seam carving: main idea

[Shai & Avidan, SIGGRAPH 2007]

Seam carving: main idea

Content-aware resizing

Traditional resizing

[Shai & Avidan, SIGGRAPH 2007]

Seam carving: main idea

Seam carving: main idea

Content-aware resizing

Intuition:

- Preserve the most "interesting" content
  → Prefer to remove pixels with low gradient energy
- To reduce or increase size in one dimension, remove irregularly shaped "seams"
  → Optimal solution via dynamic programming.

Seam carving: main idea

$$Energy(f) = \sqrt{\left(\frac{\partial f}{\partial x}\right)^2 + \left(\frac{\partial f}{\partial y}\right)^2}$$

- Want to remove seams where they won't be very noticeable:
  - Measure "energy" as gradient magnitude
- Choose seam based on **minimum total energy path** across image, subject to 8-connectedness.

Seam carving: algorithm

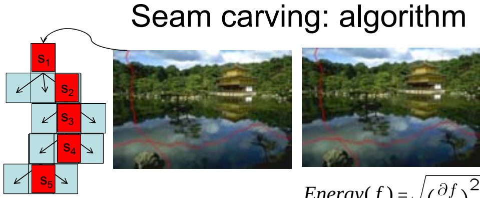

$$Energy(f) = \sqrt{\left(\frac{\partial f}{\partial x}\right)^2 + \left(\frac{\partial f}{\partial y}\right)^2}$$

Let a vertical seam **s** consist of h positions that form an 8-connected path.

Let the cost of a seam be: $Cost(\mathbf{s}) = \sum_{i=1}^{h} Energy(f(s_i))$

Optimal seam minimizes this cost: $\mathbf{s}^* = \min_{\mathbf{s}} Cost(\mathbf{s})$

Compute it efficiently with dynamic programming.

How to identify the minimum cost seam?

- First, consider a **greedy** approach:

$$\begin{matrix} 1 & 3 & \boxed{0} \\ 2 & \boxed{8} & 9 \\ 5 & \boxed{2} & 6 \end{matrix}$$

Energy matrix (gradient magnitude)

Seam carving: algorithm

- Compute the cumulative minimum energy for all possible connected seams at each entry *(i,j)*:

$$\mathbf{M}(i,j) = Energy(i,j) + \min\big(\mathbf{M}(i-1,j-1), \mathbf{M}(i-1,j), \mathbf{M}(i-1,j+1)\big)$$

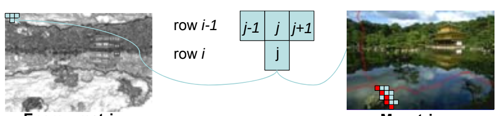

Energy matrix (gradient magnitude)

M matrix: cumulative min energy (for vertical seams)

- Then, min value in last row of **M** indicates end of the minimal connected vertical seam.
- Backtrack up from there, selecting min of 3 above in **M**.

Example

$$\mathbf{M}(i,j) = Energy(i,j) + \min\big(\mathbf{M}(i-1,j-1), \mathbf{M}(i-1,j), \mathbf{M}(i-1,j+1)\big)$$

$$\begin{matrix} 1 & 3 & 0 \\ 2 & 8 & 9 \\ 5 & 2 & 6 \end{matrix}$$

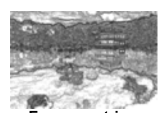

Energy matrix (gradient magnitude)

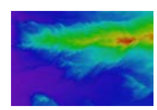

M matrix (for vertical seams)

Example

$$\mathbf{M}(i,j) = Energy(i,j) + \min\big(\mathbf{M}(i-1,j-1), \mathbf{M}(i-1,j), \mathbf{M}(i-1,j+1)\big)$$

$$\begin{matrix} \boxed{1} & 3 & 0 \\ \boxed{2} & 8 & 9 \\ 5 & \boxed{2} & 6 \end{matrix} \qquad \begin{matrix} \boxed{1} & 3 & 0 \\ \boxed{3} & 8 & 9 \\ 8 & \boxed{5} & 14 \end{matrix}$$

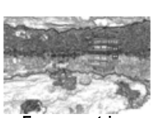

Energy matrix (gradient magnitude)

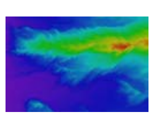

M matrix (for vertical seams)

Real image example

Original Image

Energy Map

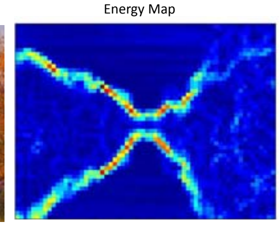

Blue = low energy
Red = high energy

Real image example

Kristen Grauman, UT-Austin

Other notes on seam carving

- Analogous procedure for horizontal seams
- Can also insert seams to *increase* size of image in either dimension
  - Duplicate optimal seam, averaged with neighbors
- Other energy functions may be plugged in
  - E.g., color-based, interactive,…
- Can use combination of vertical and horizontal seams

Kristen Grauman, UT-Austin

Example results from prior classes

(a) Original input

(b) Content-aware resizing

(c) Image from 'imresize'

Results from Eunho Yang

Results from Suyog Jain

Original image

Conventional resize

Seam carving result

Results from Martin Becker

Original image

Conventional resize

Seam carving result

Results from Martin Becker

Original image (599 by 799)

Conventional resize (399 by 599)

Seam carving (399 by 599)

Results from Jay Hennig

Removal of a marked object

(a) Selected an area. (b) Object is removed.

(c) Selected an area. (d) Object is removed.

Results from Donghyuk Shin

Removal of a marked object

Results from Eunho Yang

“Failure cases” with seam carving

By Donghyuk Shin

“Failure cases” with seam carving

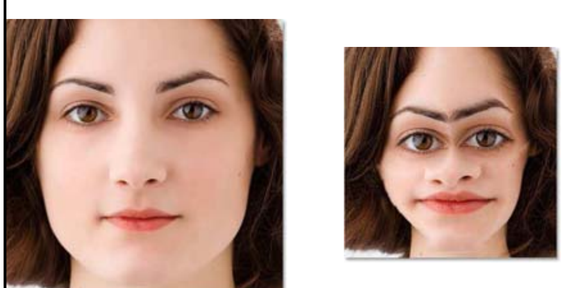

By Suyog Jain

Gradients -> edges

Primary edge detection steps:

1. Smoothing: suppress noise
2. Edge enhancement: filter for contrast
3. Edge localization

   Determine which local maxima from filter output are actually edges vs. noise

   - Threshold, Thin

Thresholding

- Choose a threshold value t
- Set any pixels less than t to zero (off)
- Set any pixels greater than or equal to t to one (on)

Original image

Gradient magnitude image

Thresholding gradient with a lower threshold

Thresholding gradient with a higher threshold

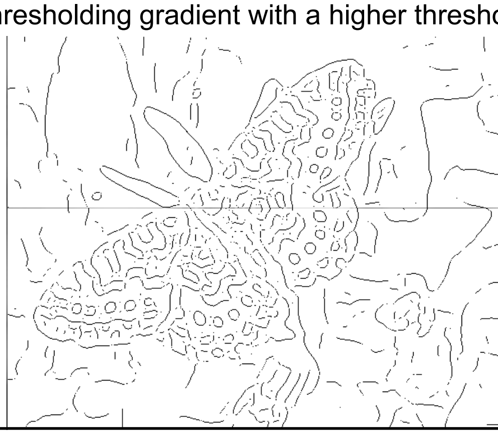

Canny edge detector

- Filter image with derivative of Gaussian
- Find magnitude and orientation of gradient
- **Non-maximum suppression**:
  - Thin wide "ridges" down to single pixel width
- **Linking and thresholding** (**hysteresis**):
  - Define two thresholds: low and high
  - Use the high threshold to start edge curves and the low threshold to continue them

- MATLAB: **`edge(image, 'canny');`**
- **`>>help edge`**

Source: D. Lowe, L. Fei-Fei

The Canny edge detector

original image (Lena)

Slide credit: Steve Seitz

The Canny edge detector

norm of the gradient

The Canny edge detector

thresholding

The Canny edge detector

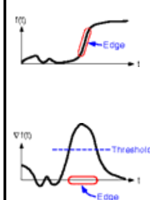

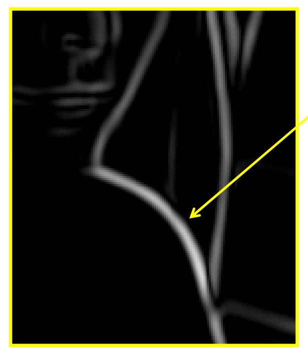

How to turn these thick regions of the gradient into curves?

Non-maximum suppression

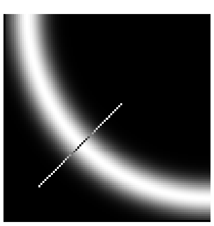

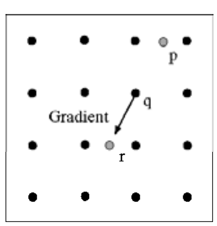

Check if pixel is local maximum along gradient direction, select single max across width of the edge

- requires checking interpolated pixels p and r

The Canny edge detector

Problem: pixels along this edge didn't survive the thresholding

thinning
(non-maximum suppression)

Hysteresis thresholding

- Use a high threshold to start edge curves, and a low threshold to continue them.

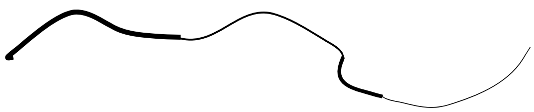

Source: Steve Seitz

Hysteresis thresholding

original image

high threshold (strong edges) | low threshold (weak edges) | hysteresis threshold

Source: L. Fei-Fei

Hysteresis thresholding

high threshold (strong edges) | low threshold (weak edges) | hysteresis threshold

Source: L. Fei-Fei

Recap: Canny edge detector

- Filter image with derivative of Gaussian
- Find magnitude and orientation of gradient
- **Non-maximum suppression**:
  - Thin wide "ridges" down to single pixel width
- **Linking and thresholding** (**hysteresis**):
  - Define two thresholds: low and high
  - Use the high threshold to start edge curves and the low threshold to continue them
- MATLAB: `edge(image, 'canny');`
- `>>help edge`

Source: D. Lowe, L. Fei-Fei

Low-level edges vs. perceived contours

Background | Texture | Shadows

Kristen Grauman, UT-Austin

Low-level edges vs. perceived contours

image | human segmentation | gradient magnitude

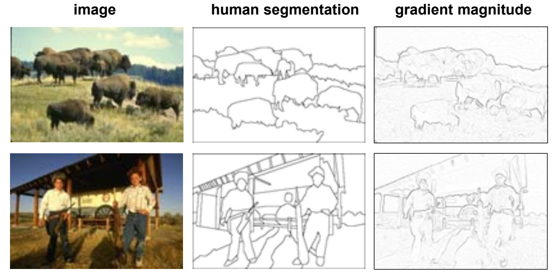

Berkeley segmentation database:
http://www.eecs.berkeley.edu/Research/Projects/CS/vision/grouping/segbench/

Source: L. Lazebnik

Learn from humans which combination of features is most indicative of a “good” contour?

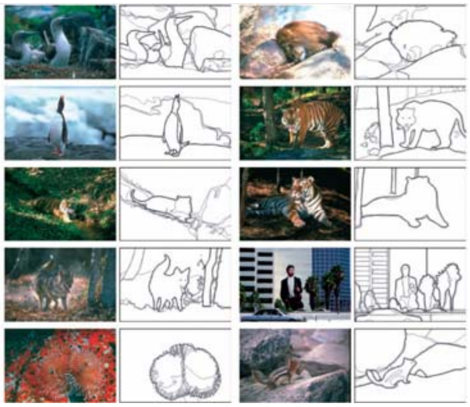

[D. Martin et al. PAMI 2004] Human-marked segment boundaries

What features are responsible for perceived edges?

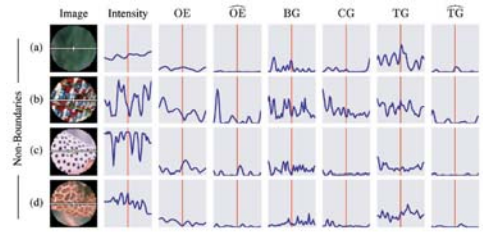

Feature profiles (oriented energy, brightness, color, and texture gradients) along the patch’s horizontal diameter

[D. Martin et al. PAMI 2004]

[D. Martin et al. PAMI 2004]

Summary

- Filters allow local image neighborhood to influence our description and features
  - Smoothing to reduce noise
  - Derivatives to locate contrast, gradient
- Convolution properties will influence the efficiency with which we can process images.
  - Associative
  - Filter separability
- Edge detection processes the image gradient to find curves, or chains of edgels.