Image gradients and edges

April 9th, 2020

Yong Jae Lee
UC Davis

Review

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

original image g

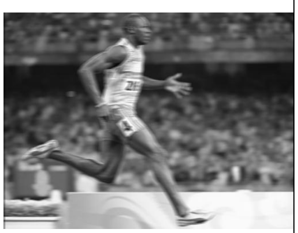

filtered

2

Slide credit: Kristen Grauman

Review

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

original image g

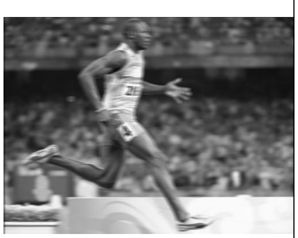

filtered

3

Slide credit: Kristen Grauman

Last time

- Image formation
- 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

4

Separability

- In some cases, filter is separable, and we can factor into two steps:
  - Convolve all rows
  - Convolve all columns

5

Slide credit: Kristen Grauman

Separability

- In some cases, filter is separable, and we can factor into two steps: e.g.,

g

1	2	1

h

2	3	3
3	5	5
4	4	6

f

What is the computational complexity advantage for a separable filter of size k x k, in terms of number of operations per output pixel?

f * (g * h) = (f * g) * h

Slide credit: Kristen Grauman

Effect of smoothing filters

5x5

Additive Gaussian noise

Salt and pepper noise

7

Slide credit: Kristen Grauman

Median filter

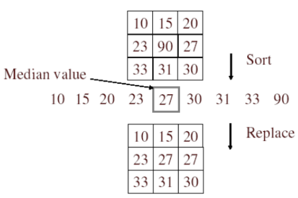

- No new pixel values introduced
- Removes spikes: good for impulse, salt & pepper noise
- Non-linear filter

8

Slide credit: Kristen Grauman

Median filter

Salt and pepper noise

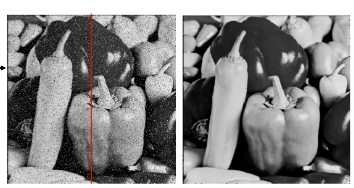

Median filtered

Matlab: output im = medfilt2(im, [h w]);

9

Slide credit: Martial Hebert

Median filter

- Median filter is edge preserving

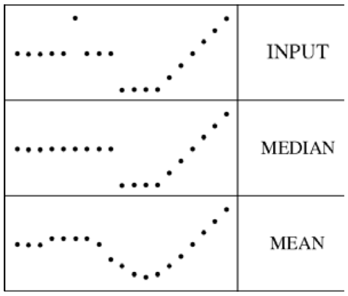

10

Slide credit: Kristen Grauman

Review

Median filter f:

Is f(a+b) = f(a)+f(b)?

Example:
a = [10 20 30 40 50]
b = [55 20 30 40 50]

Is f linear?

11

Slide credit: Devi Parikh

Filtering application: Hybrid Images

12

Slide credit: Kristen Grauman

Aude Oliva & Antonio Torralba & Philippe G Schyns, SIGGRAPH 2006

Application: Hybrid Images

A. Oliva, A. Torralba, P.G. Schyns, “Hybrid Images,” SIGGRAPH 2006

unit

13

Slide credit: Kristen Grauman

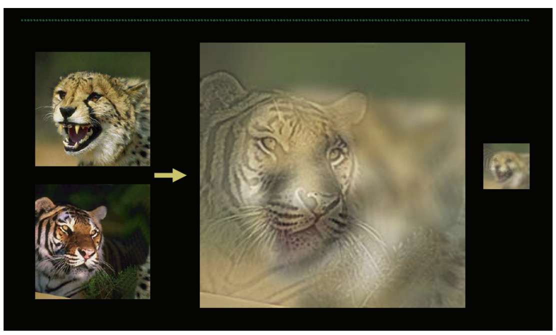

14

Aude Oliva & Antonio Torralba & Philippe G Schyns, SIGGRAPH 2006

Changing expression

SIGGRAPH2006

Sad ⟷ Surprised

Aude Oliva & Antonio Torralba & Philippe G Schyns, SIGGRAPH 2006

Summary

- Image formation
- 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

16

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, increase contrast, etc)
  - Extract information (texture, edges, interest points, etc)
  - Detect patterns (template matching)

17

Slide credit: Kristen Grauman, Adapted from Derek Hoiem

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, increase contrast, etc)
  - Extract information (texture, edges, interest points, etc)
  - Detect patterns (template matching)

18

Slide credit: Kristen Grauman, 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?**

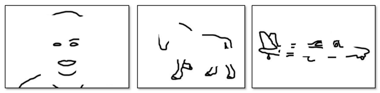

Figure from J. Shotton et al., PAMI 2007

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

19

Slide credit: Kristen Grauman

What causes an edge?

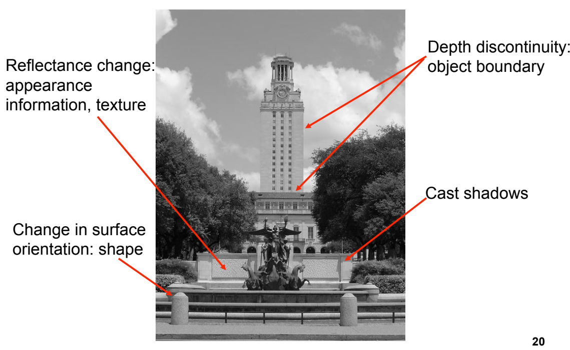

20

Edges/gradients and invariance

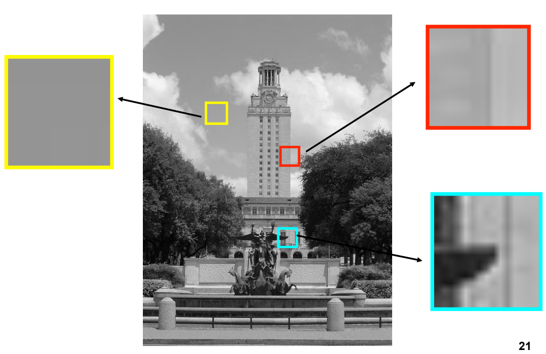

21

Derivatives and edges

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

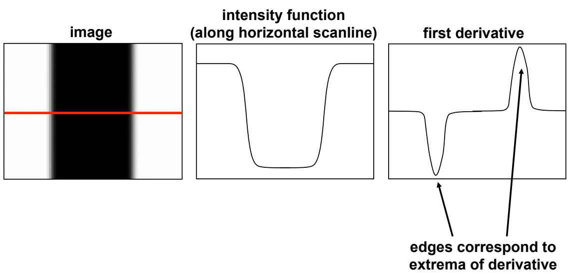

22

Slide credit: Svetlana 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?

23

Slide credit: Kristen Grauman

Partial derivatives of an image

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

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

-1	1

-1	?	1	
1	or	-1	

Which shows changes with respect to x?

(showing filters for correlation)

24

Slide credit: Kristen Grauman

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

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;
```

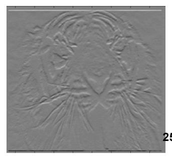

25

Slide credit: Kristen Grauman

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

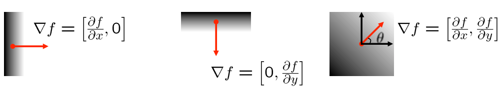

The **gradient direction** (orientation of edge normal) is given by:

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

The **edge strength** is given by the gradient magnitude

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

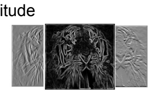

Slide credit: Steve Seitz

26

Effects of noise

Consider a single row or column of the image

- Plotting intensity as a function of position gives a signal

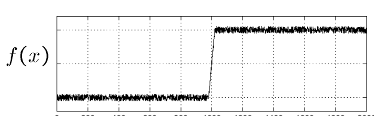

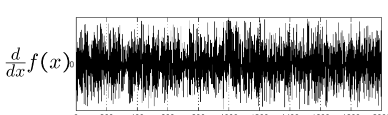

Where is the edge?

27

Slide credit: Steve Seitz

Effects of noise

- Difference filters respond strongly to noise
  - Image noise results in pixels that look very different from their neighbors
  - Generally, the larger the noise the stronger the response
- What can we do about it?

28

Source: D. Forsyth

Solution: smooth first

f

h

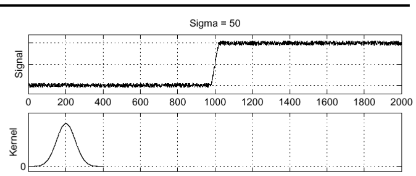

Where is the edge? Look for peaks in $\frac{\partial}{\partial x}(h \star f)$

29

Slide credit: Kristen Grauman

Derivative theorem of convolution

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

Differentiation property of convolution.

f

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

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

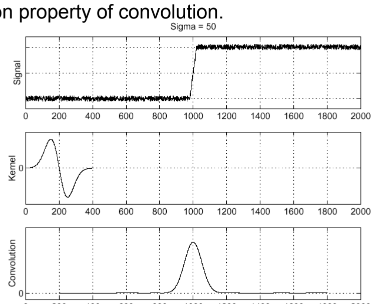

30

Slide credit: Steve Seitz

Derivative of Gaussian filters

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

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

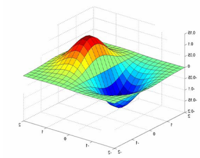

31

Slide credit: Kristen Grauman

Derivative of Gaussian filters

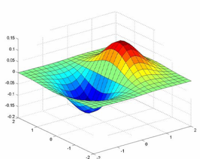

***x*-direction**

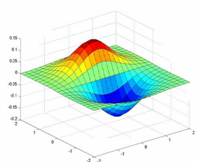

***y*-direction**

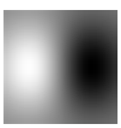

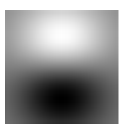

32

Slide credit: Svetlana Lazebnik

Laplacian of Gaussian

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

f

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

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

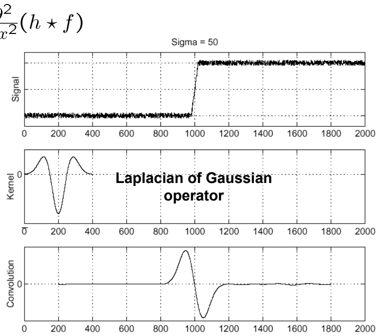

Where is the edge? Zero-crossings of bottom graph

33

Slide credit: Steve Seitz

2D edge detection filters

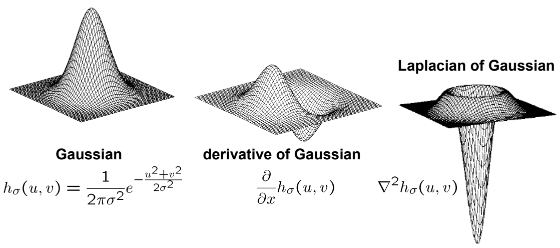

$h_\sigma(u,v) = \frac{1}{2\pi\sigma^2} e^{-\frac{u^2+v^2}{2\sigma^2}}$ $\quad \frac{\partial}{\partial x} h_\sigma(u,v)$ $\quad \nabla^2 h_\sigma(u,v)$

- ∇^2 is the Laplacian operator:

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

34

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.

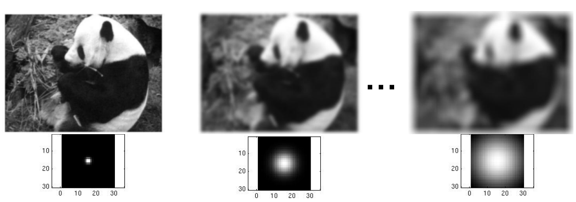

35

Slide credit: Kristen Grauman

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

36

So, what scale to choose?

It depends what we're looking for.

37

Slide credit: 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

38

Slide credit: Kristen Grauman

Seam carving: main idea

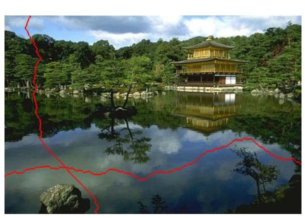

[Shai & Avidan, SIGGRAPH 2007]

39

Slide credit: Kristen Grauman

Seam carving: main idea

Content-aware resizing

Traditional resizing

[Shai & Avidan, SIGGRAPH 2007]

40

Slide credit: Kristen Grauman

Seam carving: main idea

video

41

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.

42

Slide credit: Kristen Grauman

Seam carving: main idea

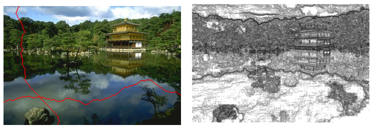

$$Energy(f) = \sqrt{\left(\frac{\partial f}{\partial x}\right)^2 + \left(\frac{\partial f}{\partial y}\right)^2}$$

- 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.

43

Slide credit: Kristen Grauman

Seam carving: algorithm

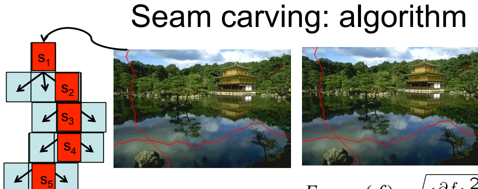

$$Energy(f) = \sqrt{\left(\frac{\partial f}{\partial x}\right)^2 + \left(\frac{\partial f}{\partial y}\right)^2}$$

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. 44

Slide credit: Kristen Grauman

How to identify the minimum cost seam?

- How many possible seams are there?
  - height *h*, width *w*
- First, consider a **greedy** approach:

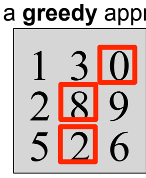

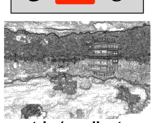

Energy matrix (gradient magnitude)

45

Slide credit: Adapted from Kristen Grauman

Seam carving: algorithm

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

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

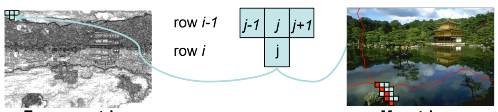

Energy matrix (gradient magnitude)

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**.

46

Slide credit: Kristen Grauman

Example

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

1	3	0
2	8	9
5	2	6

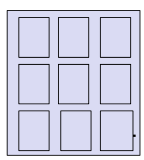

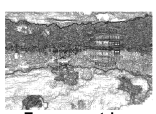

Energy matrix (gradient magnitude)

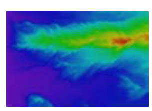

M matrix (for vertical seams)

47

Slide credit: Kristen Grauman

Example

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

1	3	0
2	8	9
5	2	6

1	3	0
3	8	9
8	5	14

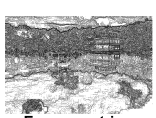

Energy matrix (gradient magnitude)

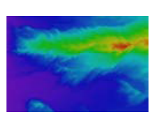

M matrix (for vertical seams)

48

Slide credit: Kristen Grauman

Real image example

Original Image

Energy Map

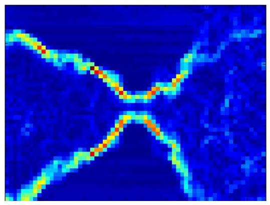

Blue = low energy
Red = high energy

Slide credit: Kristen Grauman

49

Real image example

Slide credit: Kristen Grauman

50

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

51

Slide credit: Kristen Grauman

Example seam carving results

(a) Original input

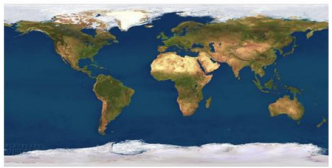

(b) Content-aware resizing

(c) Image from 'imresize'

52

53

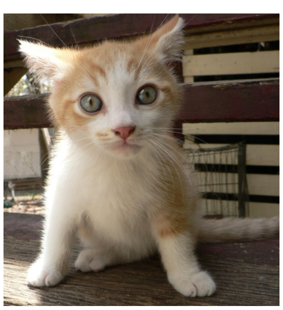

54

Original image

Conventional resize

Seam carving result

55

Original image

Conventional resize

Seam carving result

56

Original image (599 by 799)

Conventional resize (399 by 599)

Seam carving (399 by 599)

57

Removal of a marked object

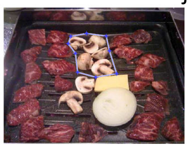

(a) Selected an area.

(b) Object is removed.

(c) Selected an area.

(d) Object is removed.

58

Removal of a marked object

59

Removal of a marked object

Figure A

Figure C

60

“Failure cases” with seam carving

61

“Failure cases” with seam carving

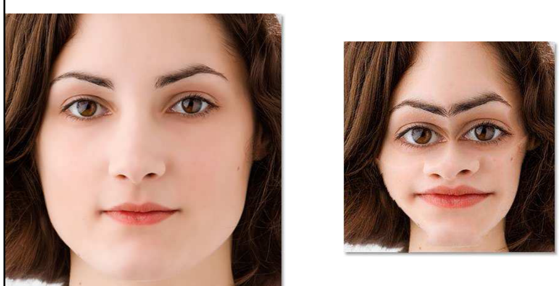

62

Questions?

See you Tuesday!

63