


# CAP 4453

## Robot Vision

Dr. Gonzalo Vaca-Castaño  
[gonzalo.vacacastano@ucf.edu](mailto:gonzalo.vacacastano@ucf.edu)






# Administrative details

- Homework 1 doubts

- Should we convert the default BGR images to RGB before displaying them? In part 4 we're instructed to "convert it back to BGR to visualize it," does that mean we should keep all the images as BGR instead of converting to RGB to display them as they were originally?
- For parts 3, 4, and 5, should we be modifying/plotting the red-green swapped image we created in part 2, or using the original image (without those channels swapped)?



# Questions?



# Robot Vision

## 4. Image Filtering II

# Credits

- Some slides comes directly from:
  - Yogesh S Rawat (UCF)
  - Noah Snavely (Cornell)
  - Ioannis (Yannis) Gkioulekas (CMU)
  - Mubarak Shah (UCF)
  - S. Seitz
  - James Tompkin
  - Ulas Bagci
  - L. Lazebnik

# Outline

- Image as a function
- Extracting useful information from Images
  - ~~Histogram~~
  - ~~Filtering (linear)~~
  - ~~Smoothing/Removing noise~~
  - ~~Convolution/Correlation~~
  - Image Derivatives/Gradient
  - **Edges**


# Edge Detection

- Identify sudden changes in an image
  - Semantic and shape information
  - Marks the border of an object
  - More compact than pixels





# Images as functions...




- Edges look like steep cliffs



# Characterizing edges

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





# Detecting edges

How would you go about detecting edges in an image (i.e., discontinuities in a function)?

✓ You take derivatives: derivatives are large at discontinuities.

How do you differentiate a discrete image (or any other discrete signal)?



# Detecting edges

How would you go about detecting edges in an image (i.e., discontinuities in a function)?

- ✓ You take derivatives: derivatives are large at discontinuities.

How do you differentiate a discrete image (or any other discrete signal)?

- ✓ You use finite differences.



# Finite differences

High-school reminder: definition of a derivative using forward difference

$$f'(x) = \lim_{h \rightarrow 0} \frac{f(x+h) - f(x)}{h}$$



# Finite differences

High-school reminder: definition of a derivative using forward difference

$$f'(x) = \lim_{h \rightarrow 0} \frac{f(x+h) - f(x)}{h}$$

Alternative: use central difference

$$f'(x) = \lim_{h \rightarrow 0} \frac{f(x+0.5h) - f(x-0.5h)}{h}$$

For discrete signals: Remove limit and set  $h = 2$

$$f'(x) = \frac{f(x+1) - f(x-1)}{2}$$

What convolution kernel does this correspond to?



# Finite differences

High-school reminder: definition of a derivative using forward difference

$$f'(x) = \lim_{h \rightarrow 0} \frac{f(x+h) - f(x)}{h}$$

Alternative: use central difference

$$f'(x) = \lim_{h \rightarrow 0} \frac{f(x+0.5h) - f(x-0.5h)}{h}$$

For discrete signals: Remove limit and set  $h = 2$

$$f'(x) = \frac{f(x+1) - f(x-1)}{2}$$

-1	0	1
----	---	---

1	0	-1
---	---	----

# Example 1D signal

How do we compute the derivative of a discrete signal?

10	20	10	200	210	250	250
----	----	----	-----	-----	-----	-----



$$f'(x) = \frac{f(x+1) - f(x-1)}{2} = \frac{210 - 10}{2} = 100$$

-1	0	1
----	---	---

1D derivative filter



# The Sobel filter

1	0	-1
2	0	-2
1	0	-1

Sobel filter

=

1
2
1

What filter  
is this?

\*

1	0	-1
---	---	----

1D derivative  
filter

# The Sobel filter

1	0	-1
2	0	-2
1	0	-1

Sobel filter

=

1
2
1

Blurring

\*

1	0	-1
---	---	----

1D derivative  
filter

In a 2D image, does this filter responses along horizontal or vertical lines?

# The Sobel filter

1	0	-1
2	0	-2
1	0	-1

Sobel filter

=

1
2
1

Blurring

\*

1	0	-1
---	---	----

1D derivative  
filter

Does this filter return large responses on vertical or horizontal lines?

# The Sobel filter

Horizontal Sober filter:

$$\begin{bmatrix} 1 & 0 & -1 \\ 2 & 0 & -2 \\ 1 & 0 & -1 \end{bmatrix} = \begin{bmatrix} 1 \\ 2 \\ 1 \end{bmatrix} * \begin{bmatrix} 1 & 0 & -1 \end{bmatrix}$$

What does the vertical Sobel filter look like?

# The Sobel filter

Horizontal Sober filter:

$$\begin{bmatrix} 1 & 0 & -1 \\ 2 & 0 & -2 \\ 1 & 0 & -1 \end{bmatrix} = \begin{bmatrix} 1 \\ 2 \\ 1 \end{bmatrix} * \begin{bmatrix} 1 & 0 & -1 \end{bmatrix}$$


Vertical Sobel filter:

$$\begin{bmatrix} 1 & 2 & 1 \\ 0 & 0 & 0 \\ -1 & -2 & -1 \end{bmatrix} = \begin{bmatrix} 1 \\ 0 \\ -1 \end{bmatrix} * \begin{bmatrix} 1 & 2 & 1 \end{bmatrix}$$


# Sobel filter example



original



which Sobel filter?




which Sobel filter?




# Sobel filter example



original




horizontal Sobel filter



vertical Sobel filter




# Sobel filter example



original



horizontal Sobel filter



vertical Sobel filter

# Several derivative filters

Sobel

1	0	-1
2	0	-2
1	0	-1

1	2	1
0	0	0
-1	-2	-1

Scharr

3	0	-3
10	0	-10
3	0	-3

3	10	3
0	0	0
-3	-10	-3

Prewitt

1	0	-1
1	0	-1
1	0	-1

1	1	1
0	0	0
-1	-1	-1

Roberts

0	1
-1	0

1	0
0	-1

- How are the other filters derived and how do they relate to the Sobel filter?
- How would you derive a derivative filter that is larger than 3x3?

# Computing image gradients

1. Select your favorite derivative filters.

$$S_x = \begin{array}{|c|c|c|} \hline 1 & 0 & -1 \\ \hline 2 & 0 & -2 \\ \hline 1 & 0 & -1 \\ \hline \end{array}$$

$$S_y = \begin{array}{|c|c|c|} \hline 1 & 2 & 1 \\ \hline 0 & 0 & 0 \\ \hline -1 & -2 & -1 \\ \hline \end{array}$$

# Computing image gradients

1. Select your favorite derivative filters.

$$\mathbf{S}_x = \begin{array}{|c|c|c|} \hline 1 & 0 & -1 \\ \hline 2 & 0 & -2 \\ \hline 1 & 0 & -1 \\ \hline \end{array}$$

$$\mathbf{S}_y = \begin{array}{|c|c|c|} \hline 1 & 2 & 1 \\ \hline 0 & 0 & 0 \\ \hline -1 & -2 & -1 \\ \hline \end{array}$$

2. Convolve with the image to compute derivatives.

$$\frac{\partial f}{\partial x} = \mathbf{S}_x \otimes f$$

$$\frac{\partial f}{\partial y} = \mathbf{S}_y \otimes f$$

# Computing image gradients

1. Select your favorite derivative filters.

$$\mathbf{S}_x = \begin{bmatrix} 1 & 0 & -1 \\ 2 & 0 & -2 \\ 1 & 0 & -1 \end{bmatrix}$$

$$\mathbf{S}_y = \begin{bmatrix} 1 & 2 & 1 \\ 0 & 0 & 0 \\ -1 & -2 & -1 \end{bmatrix}$$

2. Convolve with the image to compute derivatives.

$$\frac{\partial f}{\partial x} = \mathbf{S}_x \otimes f$$

$$\frac{\partial f}{\partial y} = \mathbf{S}_y \otimes f$$

3. Form the image gradient, and compute its direction and amplitude.

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

gradient

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


direction

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


amplitude

# Image Gradient


Gradient in x only



Gradient in y only



Gradient in both x and y



## Gradient direction

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

*How does the gradient direction relate to the edge?*

## Gradient magnitude


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

*What does a large magnitude look like in the image?*




# Image gradient example


original



vertical  
derivative



gradient  
amplitude



horizontal  
derivative




How does the gradient direction relate to these edges?



# How do you find the edge of this signal?




intensity plot




# How do you find the edge of this signal?

intensity plot



Using a derivative filter:


derivative plot



What's the problem here?

# Differentiation is very sensitive to noise

When using derivative filters, it is critical to blur first!




How much  
should we blur?


# Derivative of Gaussian (DoG) filter

Derivative theorem of convolution:  $\frac{\partial}{\partial x}(h \star f) = (\frac{\partial}{\partial x}h) \star f$


input



derivative of  
Gaussian



output (same  
as before)



- How many operations did we save?
- Any other advantages beyond efficiency?

# Laplace filter

A.K.A. Laplacian, Laplacian of Gaussian (LoG), Marr filter, Mexican Hat Function



# Laplace filter

Basically a second derivative filter.

- We can use finite differences to derive it, as with first derivative filter.

first-order  
finite difference

$$f'(x) = \lim_{h \rightarrow 0} \frac{f(x + 0.5h) - f(x - 0.5h)}{h}$$



1D derivative filter

1	0	-1
---	---	----

second-order  
finite difference

$$f''(x) = \lim_{h \rightarrow 0} \frac{f(x + h) - 2f(x) + f(x - h)}{h^2}$$



Laplace filter

?

# Laplace filter

Basically a second derivative filter.

- We can use finite differences to derive it, as with first derivative filter.

first-order  
finite difference

$$f'(x) = \lim_{h \rightarrow 0} \frac{f(x + 0.5h) - f(x - 0.5h)}{h}$$



1D derivative filter

1	0	-1
---	---	----

second-order  
finite difference

$$f''(x) = \lim_{h \rightarrow 0} \frac{f(x + h) - 2f(x) + f(x - h)}{h^2}$$




Laplace filter

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




# Laplacian of Gaussian (LoG) filter

As with derivative, we can combine Laplace filtering with Gaussian filtering



# Laplacian of Gaussian (LoG) filter

As with derivative, we can combine Laplace filtering with Gaussian filtering



“zero crossings” at edges

# Laplace and LoG filtering examples



Laplacian of Gaussian filtering



Laplace filtering



# Laplacian of Gaussian vs Derivative of Gaussian




Laplacian of Gaussian filtering




Derivative of Gaussian filtering

# Laplacian of Gaussian vs Derivative of Gaussian




Laplacian of Gaussian filtering



Derivative of Gaussian filtering


Zero crossings are more accurate at localizing edges (but not very convenient).

# 2D Gaussian filters




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

Gaussian



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

Derivative of Gaussian



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

Mexican hat →

Laplacian of Gaussian



# References

Basic reading:

- Szeliski textbook, Section 3.2



# Questions?