

Detectors and Descriptors

EECS 442 – Prof. David Fouhey
Winter 2019, University of Michigan

http://web.eecs.umich.edu/~fouhey/teaching/EECS442_W19/

Administrivia

- Today/Thursday: Detecting edges and corners in images and describing them
- Discussion Section: reviewing projection and convolution.
- If you have accommodation needs or need an alternate midterm, please email me **this week.**
- Please do not schedule stuff for the class that is listed as the midterm going forward.

General Hints for Vision Debugging

- Visualize
- **Visualize**
- **Visualize**
- Break it into bite-sized chunks; verify each
- Test the smallest version possible
- Use a debugger (pdb or use the jupyter cells)
- Try explaining your code verbally

Goal


How big is this image as a vector?
 $389 \times 600 = 233,400$ dimensions (**big**)



Applications To Have In Mind




Part of the same photo?



Same computer from another angle?

Applications To Have In Mind

Building a 3D Reconstruction Out Of Images



Applications To Have In Mind

Stitching photos taken at different angles



One Familiar Example

Given two images: how do you align them?



One (Hopefully Familiar) Solution

```
for y in range(-ySearch,ySearch+1):  
    for x in range(-xSearch,xSearch+1):  
        #Touches all HxW pixels!  
        check_alignment_with_images()
```

One Motivating Example

Given these images: how do you align them?



These aren't off by a small 2D translation but instead by a 3D rotation + translation of the camera.

One (Hopefully Familiar) Solution

```
for y in yRange:  
    for x in xRange:  
        for z in zRange:  
            for xRot in xRotVals:  
                for yRot in yRotVals:  
                    for zRot in zRotVals:  
                        #touches all HxW pixels!  
                        check_alignment_with_images()
```

This code should make you really unhappy

Note: this actually isn't even the full number of parameters; it's actually 8 for loops.


An Alternate Approach

Given these images: how would you align them?



An Alternate Approach

Finding and Matching




1: find corners+features


2: match based on local image data

What Now?

Given pairs
p1, p2 of
correspondence,
how do I align?




Consider translation-
only case from HW1.



An Alternate Approach

Solving for a Transformation



3: Solve for transformation T (e.g. such that $\mathbf{p1} \equiv T \mathbf{p2}$) that fits the matches well

Note the homogeneous coordinates, you'll see them again.

Slide Credit: S. Lazebnik, original figure: M. Brown, D. Lowe

An Alternate Approach


Blend Them Together



Key insight: we don't work with full image. We work with only parts of the image.

Today

Finding edges (part 1) and corners (part 2) in images.

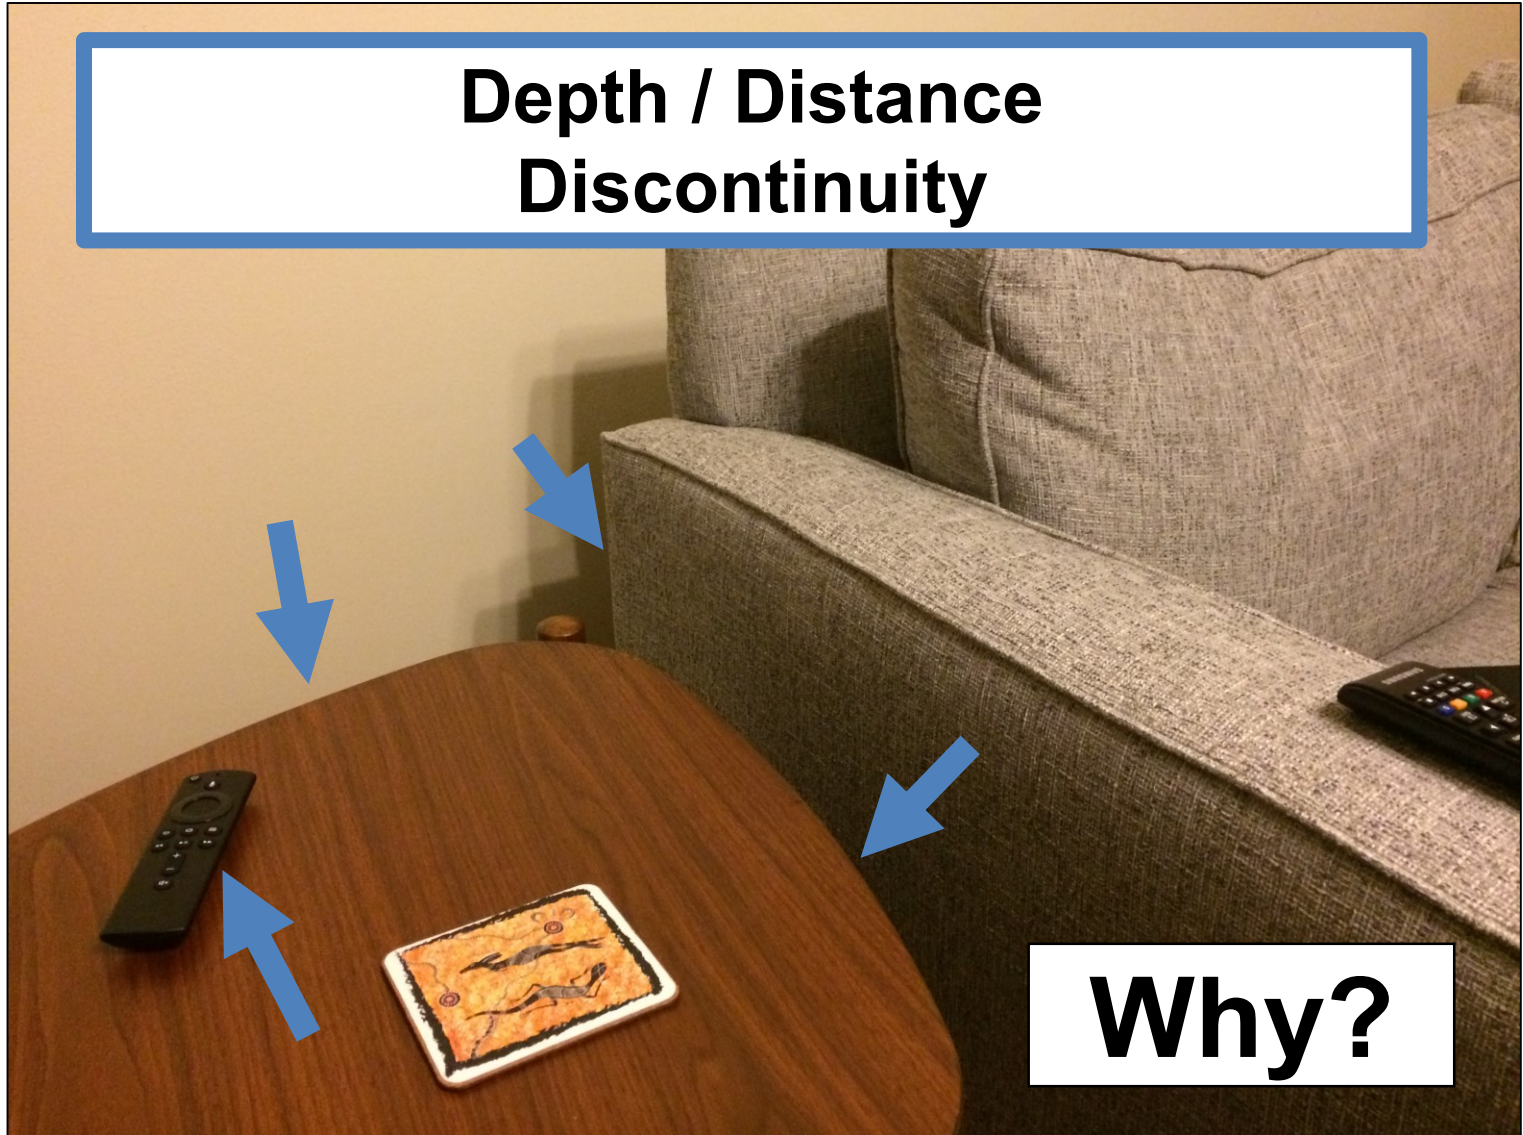


Where do Edges Come From?



Where do Edges Come From?


**Depth / Distance
Discontinuity**



Why?

Where do Edges Come From?


**Surface Normal / Orientation
Discontinuity**



Why?

Where do Edges Come From?

**Surface Color / Reflectance
Properties Discontinuity**



Where do Edges Come From?

**Illumination
Discontinuity**



Last Time



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

I_x

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

^T

I_y



Last Time

$$(\sqrt{x^2 + y^2})^{1/2}$$



Why Does This Work?

Image is function $f(x,y)$

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

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

-1	1
----	---

Another one:
$$\frac{\partial f(x, y)}{\partial x} \approx \frac{f(x + 1, y) - f(x - 1, y)}{2}$$

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

Other Differentiation Operations

Horizontal

Vertical

Prewitt

$$\begin{bmatrix} -1 & 0 & 1 \\ -1 & 0 & 1 \\ -1 & 0 & 1 \end{bmatrix}$$

$$\begin{bmatrix} 1 & 1 & 1 \\ 0 & 0 & 0 \\ -1 & -1 & -1 \end{bmatrix}$$

Sobel

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

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

Why might people use these compared to $[-1,0,1]$?

Review: Gradients

Gradient is just the collection of partial derivatives in each dimension/direction.

Points in direction that increases the most.

$$\mathbf{g} = \nabla f = \begin{bmatrix} \frac{\partial f}{\partial x_1} \\ \vdots \\ \frac{\partial f}{\partial x_n} \end{bmatrix} \quad \mathbf{g}(\mathbf{a}) = \begin{bmatrix} \frac{\partial f}{\partial x_1}(\mathbf{a}) \\ \vdots \\ \frac{\partial f}{\partial x_n}(\mathbf{a}) \end{bmatrix}$$

How much f changes in the direction of x_1 at point \mathbf{a} .

Images as Functions or Points


Key idea: can treat image as a point in $\mathbb{R}^{(H \times W)}$
or as a function of x, y .

$$\nabla I(x, y) = \begin{bmatrix} \frac{\partial I}{\partial x}(x, y) \\ \frac{\partial I}{\partial y}(x, y) \end{bmatrix}$$


How much the intensity
of the image changes
as you go horizontally
at (x, y)
(Often called I_x)

Image Gradient Direction


Some gradients



$$\nabla f = \left[\frac{\partial f}{\partial x}, 0 \right]$$



$$\nabla f = \left[0, \frac{\partial f}{\partial y} \right]$$




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

Image Gradient

Gradient: direction of maximum change.
What's the relationship to edge direction?

I_x



I_y





Image Gradient

$(I_x^2 + I_y^2)^{1/2}$: magnitude



Image Gradient



$\text{atan2}(I_y, I_x)$: orientation



I'm making the lightness equal to gradient magnitude

Image Gradient

$\text{atan2}(I_y, I_x)$: orientation




Now I'm showing *all* the gradients

Image Gradient



$\text{atan2}(I_y, I_x)$: orientation

Why is there structure at 1 and not at 2?



Noise

Consider a row of $f(x,y)$ (i.e., fix y)



Noise

Conv. image + per-pixel noise with

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

$$I_{i,j} = \text{True image} \quad \epsilon_{i,j} \sim N(0, \sigma^2)$$

$$D_{i,j} = (I_{i,j+1} + \epsilon_{i,j+1}) - (I_{i,j-1} + \epsilon_{i,j-1})$$

$$D_{i,j} = \underbrace{(I_{i,j+1} - I_{i,j-1})}_{\text{True difference}} + \underbrace{\epsilon_{i,j+1} - \epsilon_{i,j-1}}_{\text{Sum of 2 Gaussians}}$$



True
difference

Sum of 2
Gaussians

$$\epsilon_{i,j} - \epsilon_{k,l} \sim N(0, 2\sigma^2) \rightarrow \text{Variance doubles!}$$

Noise





Consider a row of $f(x,y)$ (i.e., make y constant)



How can we use the last class to fix this?


Handling Noise

Sigma = 50




Noise in 2D



Noisy Input



I_x via $[-1,01]$



Zoom




Noise + Smoothing





Smoothed Input



I_x via $[-1,01]$



Zoom




Let's Make It One Pass (1D)


$$\frac{d}{dx}(f * g) = f * \frac{d}{dx}g$$

Sigma = 50


f



$\frac{d}{dx}g$







$f * \frac{d}{dx}g$



Let's Make It One Pass (2D)

Gaussian Derivative Filter




Which one finds the X direction?

Applying the Gaussian Derivative

1 pixel

3 pixels



7 pixels



Removes noise, but blurs edge

Compared with the Past

Gaussian
Derivative



Sobel
Filter


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

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

Why would anybody use the bottom filter?


Filters We've Seen

Smoothing



Gaussian

Derivative



Deriv. of gauss

Example

Goal

Remove noise

Find edges

Only +?

Yes

No

Sums to

1

0


Why sum to 1 or 0, intuitively?

Problems



Image



human segmentation



gradient magnitude



Still an active area of research

Corners


9300 Harris Corners Pkwy, Charlotte, NC



Desirables


- Repeatability: should find same things even with distortion
- Saliency: each feature should be distinctive
- Compactness: shouldn't just be all the pixels
- Locality: should only depend on local image data

Example



Can you find the correspondences?


Example Matches




Look for the colored squares

Basic Idea


Should see where we are based on small window, or any shift \rightarrow big intensity change.



“flat” region:
no change in
all directions



“edge”:
no change
along the edge
direction




“corner”:
significant
change in all
directions

Formalizing Corner Detection

Sum of squared differences between image and image shifted u, v pixels over.

$$E(u, v) = \sum_{(x,y) \in W} (I[x + u, y + v] - I[x, y])^2$$




Formalizing Corner Detection

Sum of squared differences between image and image shifted u, v pixels over.


$$E(u, v) = \sum_{(x,y) \in W} (I[x + u, y + v] - I[x, y])^2$$

What's the value of $E(0,0)$?

Image $I(x,y)$




Plot of $E(u,v)$



Formalizing Corner Detection

Can compute $E[u,v]$ for any window and u,v .
But we'd like an simpler function of u,v .



Aside: Taylor Series for Images

Recall Taylor Series:

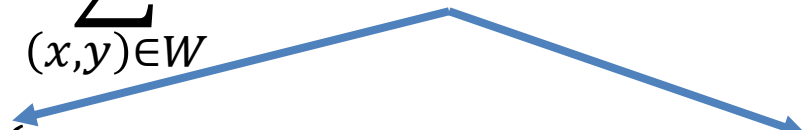
$$f(x + d) \approx f(x) + \frac{\partial f}{\partial x} d$$

Do the same with images, treating them as
function of x, y

$$I(x + u, y + v) \approx I(x, y) + I_x u + I_y v$$

Formalizing Corner Detection

Taylor series expansion for I at every single point in window

$$E(u, v) = \sum_{(x,y) \in W} (I[x+u, y+v] - I[x, y])^2$$

$$\approx \sum_{(x,y) \in W} (I[x, y] + I_x[x, y]u + I_y[x, y]v - I[x, y])^2$$

Cancel

$$= \sum_{(x,y) \in W} (I_x[x, y]u + I_y[x, y]v)^2$$

Expand

$$= \sum_{(x,y) \in W} I_x^2 u^2 + 2I_x I_y uv + I_y^2 v^2$$

For brevity: $I_x = I_x$ at point (x,y) , $I_y = I_y$ at point (x,y)

Formalizing corner Detection

By linearizing image, we can approximate $E(u,v)$ with quadratic function of u and v

$$E(u, v) \approx \sum_{(x,y) \in W} (I_x^2 u^2 + 2I_x I_y uv + I_y^2 v^2)$$
$$= [u, v] \mathbf{M} [u, v]^T$$

$$\mathbf{M} = \begin{bmatrix} \sum_{x,y \in W} I_x^2 & \sum_{x,y \in W} I_x I_y \\ \sum_{x,y \in W} I_x I_y & \sum_{x,y \in W} I_y^2 \end{bmatrix}$$

M is called the second moment matrix

Intuitively what is M?


Pretend for now gradients are either vertical or horizontal at a pixel (so $I_x I_y = 0$)

Obviously Wrong!

$$M = \begin{bmatrix} \sum_{x,y \in W} I_x^2 & \sum_{x,y \in W} I_x I_y \\ \sum_{x,y \in W} I_x I_y & \sum_{x,y \in W} I_y^2 \end{bmatrix} = \begin{bmatrix} a & 0 \\ 0 & b \end{bmatrix}$$


If a,b are both small:

flat




If one is big, one is small:

edge



If a,b both big:

corner



Review: Quadratic Forms


Suppose have symmetric matrix \mathbf{M} , scalar a , vector $[u,v]$:

$$E([u, v]) = [u, v]\mathbf{M}[u, v]^T$$

Then the isocontour / slice-through of F , i.e.

$$E([u, v]) = a$$

is an ellipse.




Review: Quadratic Forms

We can look at the shape of this ellipse by decomposing M into a rotation + scaling

$$M = R^{-1} \begin{bmatrix} \lambda_1 & 0 \\ 0 & \lambda_2 \end{bmatrix} R$$

**What are
 λ_1 and λ_2 ?**



Interpreting The Matrix M

The second moment matrix tells us how quickly the image changes and in which directions.

Can compute at each pixel

$$\mathbf{M} = \begin{bmatrix} \sum_{x,y \in W} I_x^2 & \sum_{x,y \in W} I_x I_y \\ \sum_{x,y \in W} I_x I_y & \sum_{x,y \in W} I_y^2 \end{bmatrix} = \mathbf{R}^{-1} \begin{bmatrix} \lambda_1 & 0 \\ 0 & \lambda_2 \end{bmatrix} \mathbf{R}$$


Directions

Amounts


Visualizing M



Visualizing M




Interpreting Eigenvalues of M



Putting Together The Eigenvalues

$$R = \det(\mathbf{M}) - \alpha \text{trace}(\mathbf{M})^2 \\ = \lambda_1 \lambda_2 - \alpha (\lambda_1 + \lambda_2)^2$$

α : constant (0.04 to 0.06)



In Practice

1. Compute partial derivatives I_x , I_y per pixel
2. Compute \mathbf{M} at each pixel, using Gaussian weighting w

$$\mathbf{M} = \begin{bmatrix} \sum_{x,y \in W} w(x,y) I_x^2 & \sum_{x,y \in W} w(x,y) I_x I_y \\ \sum_{x,y \in W} w(x,y) I_x I_y & \sum_{x,y \in W} w(x,y) I_y^2 \end{bmatrix}$$

C.Harris and M.Stephens. [“A Combined Corner and Edge Detector.”](#)
Proceedings of the 4th Alvey Vision Conference: pages 147—151, 1988.

In Practice

1. Compute partial derivatives I_x , I_y per pixel
2. Compute \mathbf{M} at each pixel, using Gaussian weighting w
3. Compute response function R


$$\begin{aligned} R &= \det(\mathbf{M}) - \alpha \operatorname{trace}(\mathbf{M})^2 \\ &= \lambda_1 \lambda_2 - \alpha (\lambda_1 + \lambda_2)^2 \end{aligned}$$

C.Harris and M.Stephens. ["A Combined Corner and Edge Detector."](#)
Proceedings of the 4th Alvey Vision Conference: pages 147—151, 1988.

Computing R



Computing R



In Practice

1. Compute partial derivatives I_x , I_y per pixel
2. Compute \mathbf{M} at each pixel, using Gaussian weighting w
3. Compute response function R
4. Threshold R

C.Harris and M.Stephens. ["A Combined Corner and Edge Detector."](#)
Proceedings of the 4th Alvey Vision Conference: pages 147—151, 1988.

Thresholded R




In Practice

1. Compute partial derivatives I_x , I_y per pixel
2. Compute \mathbf{M} at each pixel, using Gaussian weighting w
3. Compute response function R
4. Threshold R
5. Take only local maxima (called non-maxima suppression)

C.Harris and M.Stephens. ["A Combined Corner and Edge Detector."](#)
Proceedings of the 4th Alvey Vision Conference: pages 147—151, 1988.

Thresholded, NMS R



Final Results



Desirable Properties

If our detectors are repeatable, they should be:

- **Invariant** to some things: image is transformed and corners remain the same
- **Covariant/equivariant** with some things: image is transformed and corners transform with it.

Recall Motivating Problem

Images may be different in lighting and geometry




Affine Intensity Change

$$I_{new} = aI_{old} + b$$


M only depends on derivatives, so b is irrelevant

But a scales derivatives and there's a threshold



Partially invariant to affine intensity changes


Image Translation



All done with convolution. Convolution is translation invariant.

Equivariant with translation


Image Rotation



Rotations just cause the corner rotation to change.
Eigenvalues remain the same.

Equivariant with rotation

Image Scaling



One pixel can become many pixels and vice-versa.

Not equivariant with scaling

Next time

- Fixing this scaling issue
- Describing the corners

Desirable Properties


Repeatable: we can find the same place even after photometric metric and geometric distortion.



Desirable Properties

Compactness: we don't just use all the pixels

Saliency: the place is distinctive



Desirable Properties

Locality: the feature doesn't depend on the whole image but instead some part

