Lecture 6 Lines and Corners







Administrative

A1 due Fri, April 18!!!

- You can use up to 2 late days







Administrative

- Recitation this Friday
- Geometric transformations







So far: Sobel Filter

Step 1: Calculate the gradient magnitude at every pixel location.

Step 2: Threshold the values to generate a binary image

Raymond Yu



Lecture 6 - 4

So far: challenges multiple disconnected edges



Raymond Yu

Lecture 6 - 5

So far: Canny edge detector Use Sobel filters to find line estimates



Raymond Yu

Lecture 6 - 6

So far: Non-maximum suppression







So far: Hysteresis thresholding Strong and weak edges



Source: S. Seitz

Raymond Yu

Lecture 6 - 8

So far: The Hough transform

- So: one point (x_i, y_i) gives a line in (a, b) space.
- Another point (x_i, y_i) will give rise to another line in (a,b)-space.
- Iterate over pairs of points, to vote for buckets of intersection in (a,b)-space



Raymond Yu

Lecture 6 - 9

So far: RANSAC

• Sample seed points, calculate line, count # of inliers, repeat



Raymond Yu



Today's agenda

- RANSAC
- Local Invariant Features
- Harris Corner Detector





Today's agenda

- RANSAC
- Local Invariant Features
- Harris Corner Detector





Why is Hough transform inefficient?

- It's not feasible to check all pairs of points to calculate possible lines. For example, Hough Transform algorithm runs in O(N²).
- Voting is a general technique where we let the each point vote for all models that are compatible with it.
 - Iterate through features, cast votes for parameters.
 - Filter parameters that receive a lot of votes.
- **Problem:** Noisy points will cast votes too, *but* typically their votes should be inconsistent with the majority of "good" edge points.

Raymond Yu



Difficulty of voting for lines



- Noisy edge pixels cast inconsistent votes:
 - Can we identify false edge pixels without iterating over all pairs like we do in Hough transforms?



- Canny can predict false positive edge points:
 - Can we eliminate them without needing to compare this pixel with every other edge pixel?

Raymond Yu

Lecture 6 - 14

Intuition





Raymond Yu

Intuition





Raymond Yu

RANSAC [Fischler & Bolles 1981]

- RANdom SAmple Consensus
- **Approach**: we want to avoid the impact of noisy outliers, so let's look for "inliers", and use only those.
- Intuition: if an outlier is chosen to compute the parameters (a,b) of a line, then the resulting line won't have much support from rest of the points.





- Task: Estimate the best line
 - Let's randomly select a subset of points and calculate a line



Raymond Yu

Lecture 6 - 18

- Task: Estimate the best line
 - Let's select only 2 points as an example

 \bigcirc



Lecture 6 - 19

Sample two points

- Task: Estimate the best line
 - Calculate the line parameters







- Task: Estimate the best line
 - Edges can be noisy. To account for this, let's say that the line is somewhere between the dashed lines





- Task: Estimate the best line
 - Calculate the number of points that lie within the dashed lines







How do we calculate the inliers? We use the distance from the point to the line



April 15, 2025

Lecture 6 - 23

Raymond Yu

- Task: Estimate the best line
 - Repeat with two other randomly selected points'







Raymond Yu

Lecture 6 - 25

RANSAC loop:

Repeat for *k* iterations:

1. Randomly select a *seed* subset of points on which to perform a model estimate (e.g., a group of edge points)





<u>April 15, 2025</u>

RANSAC loop:

Repeat for *k* iterations:

- 1. Randomly select a *seed* subset of points on which to perform a model estimate (e.g., a group of edge points)
- 2. Compute parameters (a, b) from seed group





RANSAC loop:

Repeat for *k* iterations:

- 1. Randomly select a *seed* subset of points on which to perform a model estimate (e.g., a group of edge points)
- 2. Compute parameters (a, b) from seed group
- 3. Find inliers for these parameters



RANSAC loop:

Repeat for *k* iterations:

- 1. Randomly select a *seed* subset of points on which to perform a model estimate (e.g., a group of edge points)
- 2. Compute parameters (a, b) from seed group
- 3. Find inliers for these parameters
- 4. If the number of inliers is larger than the best so far, save these parameters and the inliers

Raymond Yu



RANSAC loop:

Repeat for *k* iterations:

- 1. Randomly select a *seed* subset of points on which to perform a model estimate (e.g., a group of edge points)
- 2. Compute parameters (a, b) from seed group
- 3. Find inliers for these parameters
- 4. If the number of inliers is larger than the best so far, save these parameters and the inliers

If number of inliers in the best line is < m, return no line

Raymond Yu

Lecture 6 - 30

RANSAC loop:

Repeat for *k* iterations:

- 1. Randomly select a *seed* subset of points on which to perform a model estimate (e.g., a group of edge points)
- 2. Compute parameters (a, b) from seed group
- 3. Find inliers for these parameters
- 4. If the number of inliers is larger than the best so far, save these parameters and the inliers

If number of inliers in the best line is < m, return no line

Else re-calculate the final parameters with all the inliers

Raymond Yu

Lecture 6 - 31

Final step: Refining the parameters

- The best parameters were computed using a seed set of *n* points.
- We use these points to find the inliers.
- We can improve the parameters by estimating over all inliers (e.g. with standard least-squares minimization).
- But this may change the inliers, so repeat this last step until there is no change in inliers.



Raymond Yu

Lecture 6 - 32

How do you calculate the line from many points?

 (x_i, y_i) is a set of points we are going to use to estimate (a, b)

Least squares method:

 $b = y_i - ax_i$

$$a = \frac{\left(\sum_{i} x_{i} - \bar{x}\right)\left(\sum_{i} y_{i} - \bar{y}\right)}{\left(\sum_{i} x_{i} - \bar{x}\right)^{2}}$$

April 15, 2025

Raymond Yu

RANSAC loop:

Repeat for *k* iterations:

- 1. Randomly select a *seed* subset of points on which to perform a model estimate (e.g., a group of edge points)
- 2. Compute parameters (a, b) from seed group
- 3. Find inliers for these parameters
- 4. If the number of inliers is larger than the best so far, save these parameters and the inliers

If number of inliers in the best line is < m, return no line

Else re-calculate the final parameters with all the inliers

Raymond Yu

Lecture 6 - 34

The hyperparameters

- 1. How many points to sample in the seed set?
 - a. We used 2 in the example above





The hyperparameters

- 1. How many points to sample in the seed set?
 - a. We used 2 in the example above
- 2. How many times should we repeat?
 - a. More repetitions increase computation but increase chances of finding best line




The hyperparameters

- 1. How many points to sample in the seed set?
 - a. We used 2 in the example above
- 2. How many times should we repeat?
 - a. More repetitions increase computation but increase chances of finding best line
- 3. The threshold for the dashed lines
 - a. Larger the gap between dashed lines, the more false positive inliers
 - b. Smaller the gap, the more false negatives outliers



The hyperparameters

- 1. How many points to sample in the seed set?
 - a. We used 2 in the example above
- 2. How many times should we repeat?
 - a. More repetitions increase computation but increase chances of finding best line
- 3. The threshold for the dashed lines
 - a. Larger the gap between dashed lines, the more false positive inliers
 - b. Smaller the gap, the more false negatives outliers
- 4. The minimum number of inliers to confidently claim there is a line
 - a. Smaller the number, the more false negative lines
 - b. Larger the number, the fewer lines we will find

Raymond Yu

Lecture 6 - 38

RANSAC: Computed *k* (p=0.99)

Sample	Proportion of outliers						
n	5%	10%	20%	25%	30%	40%	50%
2	2	3	5	6	7	11	17
3	3	4	7	9	11	19	35
4	3	5	9	13	17	34	72
5	4	6	12	17	26	57	146
6	4	7	16	24	37	97	293
7	4	8	20	33	54	163	588
8	5	9	26	44	78	272	1177

Raymond Yu

Lecture 6 - 39

RANSAC: How many iterations "k"?

- How many samples are needed?
 - Suppose *w* is fraction of inliers (points from line).
 - n points needed to define hypothesis (2 for lines)
 - k samples chosen.
- Prob. that a single sample of n points is correct: wⁿ
- Prob. that a single sample of n points fails: $1 w^n$
- Prob. that all k samples fail is: $(1 w^n)^k$
- Prob. that at least one of the k samples is correct: $1 (1 w^n)^k$

 \Rightarrow Choose k high enough to keep this below desired failure rate.

Raymond Yu

Lecture 6 - 40

RANSAC: Pros and Cons

• <u>Pros</u>:

- General method suited for a wide range of parameter fitting problems
- Easy to implement and easy to calculate its failure rate

• <u>Cons</u>:

- Only handles a moderate percentage of outliers without cost blowing up
- Many real problems have high rate of outliers (but sometimes selective choice of random subsets can help)
- A voting strategy, The Hough transform, can handle high percentage of outliers

Raymond Yu

Lecture 6 - 41

Today's agenda

- RANSAC
- Local Invariant Features
- Harris Corner Detector







Image matching: a challenging problem



Q1. Will cross-correlation work?

Q2. Can we use match the lines?

Raymond Yu



Lecture 6 - 43

Q. How would you build a system that can detect this movie in the pile?



Raymond Yu



Challenge: Perspective / viewpoint changes



by <u>Diva Sian</u>



by <u>swashford</u>

April 15, 2025

Raymond Yu

Challenge: partial observability



by <u>Diva Sian</u>



by <u>scgbt</u>

Raymond Yu



Challenge even for us





NASA Mars Rover images

Raymond Yu



Answer Below (Look for tiny colored squares)





April 15, 2025

NASA Mars Rover images with SIFT feature matches (Figure by Noah Snavely)

Raymond Yu

Intuition behind how to match images

- Find matching patches
- Check to make sure enough patches



Raymond Yu

Lecture 6 - 49

Intuition behind how to match images

- Find matching patches
- Check to make sure enough patches

What do we need?

- We need to identify patches
- We need to learn to a way to describe each patch
- We need an algorithm to match the description between two patches

Raymond Yu



Motivation for using local features

- Matching large patches have major challenges (mentioned in previous slides)
- Instead, let's describe and match only local image patches
- Smaller, local patches are more likely to find an object even if it is partially occluded (covered)





• Intra-category variations





Raymond Yu

Lecture 6 - 51

General approach for search



1. Find a set of distinctive key-points

2. Define a region/patch around each keypoint

3. Normalize the region content

4. Compute a local descriptor from the normalized region

5. Match local descriptors

Raymond Yu

Lecture 6 - 52

Common Requirements

- Problem 1: How should we choose the key-points?
 - We want to detect the same points independently in both images





No chance to match if the key-points aren't the same

We need a repeatable detector!

Raymond Yu

Lecture 6 - 53

Common Requirements

- Problem 1: How should we choose the key-points?
 Detect the same point independently in both images
- Problem 2: How should we describe each patch?
 For each point correctly recognize the corresponding one



We need a reliable and distinctive descriptor!

Raymond Yu

Lecture 6 - 54

Descriptions should be invariant to rotation and translation



Raymond Yu

Lecture 6 - 55

Descriptions should be invariant to photometric transformations



- Often modeled as a linear transformation:
 - Scaling + Offset

Raymond Yu

Slide credit: Tinne Tuytelaars



Levels of geometric transformations





April 15, 2025

Lecture 6 - 57

Raymond Yu

Requirements for Local Features

- Patch selection needs to be repeatable and accurate
 Invariant to translation, rotation, scale changes
 Robust to out-of-plane (≈affine) transformations
 Robust to lighting variations, noise, blur, quantization
- Locality: Features are local, therefore robust to occlusion and clutter.
- Quantity: We need a sufficient number of regions to cover the object.
- **Distinctiveness**: The regions should contain "unique" structure.
- Efficiency: Close to real-time performance.

Raymond Yu



What are good patches?

Q. Is this a good patch for image matching?



Raymond Yu

Lecture 6 - 59

What are good patches?

Q. What about this one?



Raymond Yu

Lecture 6 - 60

What are good patches?

Q. Let's try another one?



Raymond Yu

Lecture 6 - 61

Many existing feature detectors available

- Hessian & Harris
- Laplacian, DoG

- [Beaudet '78], [Harris '88]
 - [Lindeberg '98], [Lowe '99]
- Harris-/Hessian-Laplace [Mikolajczyk & Schmid '01]
- Harris-/Hessian-Affine
- EBR and IBR
- MSER
- Salient Regions
- Neural networks

[Mikolajczyk & Schmid '04] [Tuytelaars & Van Gool '04] [Matas '02] [Kadir & Brady '01] [Krichevsky '12]

• Those detectors have become a basic building block for many applications in Computer Vision.

Raymond Yu



Today's agenda

- Local Invariant Features
- Harris Corner Detector







Keypoint Localization



• Goals:

• Repeatable detection

- Precise localization
- Interesting content

intuition ⇒ Look for 2D signal changes (LSI systems strike again)

Raymond Yu

Lecture 6 - 64

Finding Corners



How do we find corners using LSI systems?

The image gradient around a corner has two or more dominant directions

Corners are **repeatable** and **distinctive**

C.Harris and M.Stephens. <u>"A Combined Corner and Edge Detector."</u> *Proceedings of the 4th Alvey Vision Conference*, 1988.

Raymond Yu



Corners are distinctive key-points

- We should easily recognize the corner point by looking through a small image patch (*locality*)
- Shifting the window in any direction should give a large change in intensity (good localization)







"flat" region: no change in all directions

"edge": no change along the edge direction

"corner": significant change in all directions

Slide credit: Alyosha Efros

April 15, 2025

Raymond Yu

Flat patches have small image gradients



Raymond Yu

Lecture 6 - 67

Edges have high gradient in one direction



Raymond Yu

Lecture 6 - 68

Corners versus edges



April 15, 2025

Lecture 6 - 69

Raymond Yu

Generalizing to corners in any direction











Harris Detector Formulation

- Find patches that result in large change of pixel values when shifted in *any direction*.
- When we shift by [*u*, *v*], the intensity change at the center pixel is:

[u, v] I(x + u, y + v)I(x, y)

"corner": significant change in all directions

Raymond Yu

- Measure change as intensity difference:
 (I(x + u, y + v) I(x, y))
- That's for a single point, but we have to accumulate over the patch or "small window" around that point...

Lecture 6 - 71

Harris Detector Formulation

• When we shift by [u, v], the change in intensity for the "small window" is:



Raymond Yu

Lecture 6 - 72

Slide cApril ck15; 182025
Change in intensity function

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

We can rewrite the shifted intensity using Taylor's expansion:

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

Substituting it back into E(u, v):

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

Raymond Yu

Lecture 6 - 73

 $E(u,v) = \sum w(x,y)[I_x u + I_y v]^2$ $_{x,y}$





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





$$\begin{split} E(u,v) &= \sum_{x,y} w(x,y) [I_x u + I_y v]^2 \\ &= \sum_{x,y} w(x,y) (I_x^2 u^2 + 2I_x I_y u v + I_y^2 v^2) \\ &= (\sum_{x,y} w I_x^2) u^2 + 2(\sum_{x,y} w I_x I_y) u v + (\sum_{x,y} w I_y^2) v^2 \end{split}$$

Raymond Yu

Lecture 6 - 76

$$\begin{split} E(u,v) &= \sum_{x,y} w(x,y) [I_x u + I_y v]^2 \\ &= \sum_{x,y} w(x,y) (I_x^2 u^2 + 2I_x I_y u v + I_y^2 v^2) \\ &= (\sum_{x,y} w I_x^2) u^2 + 2 (\sum_{x,y} w I_x I_y) u v + (\sum_{x,y} w I_y^2) v^2 \\ &= \begin{bmatrix} u & v \end{bmatrix} \sum_{x,y} w(x,y) \begin{bmatrix} I_x^2 & I_x I_y \\ I_x I_y & I_y^2 \end{bmatrix} \begin{bmatrix} u \\ v \end{bmatrix} \end{split}$$

Raymond Yu

Lecture 6 - 77

$$\begin{split} E(u,v) &= \sum_{x,y} w(x,y) [I_x u + I_y v]^2 \\ &= \sum_{x,y} w(x,y) (I_x^2 u^2 + 2I_x I_y u v + I_y^2 v^2) \\ &= (\sum_{x,y} w I_x^2) u^2 + 2(\sum_{x,y} w I_x I_y) u v + (\sum_{x,y} w I_y^2) v^2 \\ &= [u \quad v] \sum_{x,y} w(x,y) \begin{bmatrix} I_x^2 & I_x I_y \\ I_x I_y & I_y^2 \end{bmatrix} \begin{bmatrix} u \\ v \end{bmatrix} \\ &= [u \quad v] M \begin{bmatrix} u \\ v \end{bmatrix} \end{split}$$

April 15, 2025

where:

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

Raymond Yu

Simplifying M for a second:

Assuming w(x, y) = 1

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

where:

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

Raymond Yu

Lecture 6 - 79

Change in intensity in a patch

• So, using Taylor's expansion, the change in intensity in an image patch:

$$E(u,v) \approx \begin{bmatrix} u & v \end{bmatrix} M \begin{bmatrix} u \\ v \end{bmatrix}$$

where M is a 2×2 matrix computed from image derivatives:



Raymond Yu

Lecture 6 - 80

$$E(u,v) = egin{bmatrix} u & v \end{bmatrix} M egin{bmatrix} u \\ v \end{bmatrix}$$

Does anyone know what this part of the equation is?

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

Raymond Yu

Lecture 6 - 81

Harris Detector Formulation



where M is a 2×2 matrix computed from image derivatives:



Raymond Yu

Lecture 6 - 82

It's the equation of an ellipse

$$5u^{2} - 4uv + 5v^{2} = 1$$
$$\begin{bmatrix} u & v \end{bmatrix} M \begin{bmatrix} u \\ v \end{bmatrix}$$
$$M = \begin{bmatrix} 5 & -2 \\ -2 & 5 \end{bmatrix}$$



Raymond Yu

Lecture 6 - 83



If only
$$\sum I_x^2 \longrightarrow$$
 Large ,

Q. What is the matrix M going to look like?

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

Raymond Yu

Lecture 6 - 84



If only
$$\sum I_x^2 \longrightarrow$$
 Large ,

Q. What is the matrix M going $M = \begin{bmatrix} M \\ M \end{bmatrix}$

$$M = \begin{bmatrix} \sum_{x,y} I_x^2 & \sum_{x,y} I_x I_y \\ \sum_{x,y} I_x I_y & \sum_{x,y} I_y^2 \end{bmatrix}$$
$$= \begin{bmatrix} & \text{Large} & \text{Small} & ' \\ & & \text{Small} & \text{Small} \end{bmatrix}$$

April 15, 2025

Raymond Yu

Lecture 6 - 85

M



If only $\sum I_x^2 \longrightarrow$ Large , Q. what kind of ellipse would you expect to see?

$$M = \begin{bmatrix} \sum_{x,y} I_x^2 & \sum_{x,y} I_x I_y \\ \sum_{x,y} I_x I_y & \sum_{x,y} I_y^2 \end{bmatrix}$$
$$M = \begin{bmatrix} & \text{Large} & \text{Small} & ' \\ & \text{Small} & \text{Small} \end{bmatrix}$$

April 15, 2025

Lecture 6 - 86

Raymond Yu



Raymond Yu

If only $\sum I_x^2 \longrightarrow$ Large , Q. what kind of ellipse would you expect to see? $M = \begin{bmatrix} \sum_{x,y} I_x^2 & \sum_{x,y} I_x I_y \\ \sum_{x,y} I_x I_y & \sum_{x,y} I_y^2 \end{bmatrix}$ $M = \begin{bmatrix} \text{Large} & \text{Small} & \text{Small} \end{bmatrix}$

April 15, 2025



If only $\sum I_y^2 \longrightarrow$ Large , Q. what kind of ellipse would you expect to see?

$$M = \begin{bmatrix} \sum_{x,y} I_x^2 & \sum_{x,y} I_x I_y \\ \sum_{x,y} I_x I_y & \sum_{x,y} I_y^2 \end{bmatrix}$$
$$M = \begin{bmatrix} & \mathsf{Small} & \mathsf{Small} & \prime \\ & & \mathsf{Small} & \mathsf{Large} \end{bmatrix}$$

April 15, 2025

Lecture 6 - 88

Raymond Yu



If only $\sum I_y^2 \longrightarrow$ Large , Q. what kind of ellipse would you expect to see?



Raymond Yu

$$M = \begin{bmatrix} \sum_{x,y} I_x^2 & \sum_{x,y} I_x I_y \\ \sum_{x,y} I_x I_y & \sum_{x,y} I_y^2 \end{bmatrix}$$
$$M = \begin{bmatrix} & \text{Small} & \text{Small} & \prime \\ & & \text{Small} & \text{Large} \end{bmatrix}$$

April 15, 2025





Q. What is the matrix M going $M = \begin{bmatrix} \sum_{x,y} I_x^2 & \sum_{x,y} I_x I_y \\ \sum_{x,y} I_x I_y & \sum_{x,y} I_y^2 \end{bmatrix}$

$$M = \begin{bmatrix} ??? & ??? \\ ??? & ??? \\ ??? & ??? \end{bmatrix}$$

April 15, 2025

Raymond Yu





Q. What is the matrix M going $M = \begin{bmatrix} \sum_{x,y} I_x^2 & \sum_{x,y} I_x I_y \\ \sum_{x,y} I_x I_y & \sum_{x,y} I_y^2 \end{bmatrix}$

$$M = \begin{bmatrix} Large & small \\ small & Large \end{bmatrix}$$

April 15, 2025

Lecture 6 - 91

Raymond Yu





Q. What is the ellipse going to look like?

Raymond Yu



$$M = \begin{bmatrix} \sum_{x,y} I_x^2 & \sum_{x,y} I_x I_y \\ \sum_{x,y} I_x I_y & \sum_{x,y} I_y^2 \end{bmatrix}$$
$$= \begin{bmatrix} Large & small & \prime \\ small & Large \end{bmatrix}$$

April 15, 2025

Lecture 6 - 92

M

But what about these ones?



Q. What would the matrix and ellipses look like?

$$5u^{2} - 4uv + 5v^{2} = 1$$

$$\begin{bmatrix} u & v \end{bmatrix} M \begin{bmatrix} u \\ v \end{bmatrix}$$

$$M = \begin{bmatrix} \sum_{x,y} I_{x}^{2} & \sum_{x,y} I_{x}I_{y} \\ \sum_{x,y} I_{x}I_{y} & \sum_{x,y} I_{y}^{2} \end{bmatrix}$$

$$M = \begin{bmatrix} 5 & -2 \\ -2 & 5 \end{bmatrix}$$

April 15, 2025

Raymond Yu

What Does This Matrix Reveal?

- First, let's consider an axis-aligned corner.
- In that case, the dominant gradient directions align with the x or the y axis

•
$$M = \begin{bmatrix} \sum I_x^2 & \sum I_x I_y \\ \sum I_x I_y & \sum I_y^2 \end{bmatrix} = \begin{bmatrix} \lambda_1 & 0 \\ 0 & \lambda_2 \end{bmatrix}$$

Raymond Yu

- This means: if either λ is close to 0, then this is not a corner, so look for image windows where both lambdas are large.
- What if we have a corner that is not aligned with the image axes?



<u>April 15, 2025</u>

• Since $M = \begin{bmatrix} I_x^2 & I_x I_y \\ I_x I_y & I_y^2 \end{bmatrix}$ is symmetric, we can re-rewrite $M = R^{-1} \begin{bmatrix} \lambda_1 & 0 \\ 0 & \lambda_2 \end{bmatrix} R$ (Eigenvalue decomposition)

 We can think of M as an ellipse with its axis lengths determined by the eigenvalues λ₁ and λ₂; and its orientation determined by R



Raymond Yu

Lecture 6 - 95

Interpreting the Eigenvalues

• Classification of image points using eigenvalues of *M*:



But calculating eigenvalues is expensive. **Solution: Corner Response Function** $\theta = \det(M) - \alpha \operatorname{trace}(M)^2 = \lambda_1 \lambda_2 - \alpha (\lambda_1 + \lambda_2)^2$

 Λ_{γ} Edge" $\theta < 0$ "Corner" $\theta > 0$ "Flat" "Edge" region $\theta < 0$ λ_{i}

- Fast approximation

 Avoid computing the eigenvalues
 α: constant
 - (0.04 to 0.06)

Raymond Yu

Lecture 6 - 97

Slide credit April 15, 2025

Window Function *w*(*x*,*y*)

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

- Option 1: uniform window
 - Sum over square window

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



• Problem: not rotation invariant

1 in window, 0 outside

- Option 2: Smooth with Gaussian
 - Gaussian already performs weighted sum

$$M = g(\sigma) * \begin{bmatrix} I_x^2 & I_x I_y \\ I_x I_y & I_y^2 \end{bmatrix}$$

• Result is rotation invariant
an

Raymond Yu

Lecture 6 - 98

Summary: Harris Detector [Harris88]

• Compute second moment matrix (autocorrelation matrix)

 $M(\sigma_{I},\sigma_{D}) = g(\sigma_{I}) * \begin{bmatrix} I_{x}^{2}(\sigma_{D}) & I_{x}I_{y}(\sigma_{D}) \\ I_{x}I_{y}(\sigma_{D}) & I_{y}^{2}(\sigma_{D}) \end{bmatrix}$

 σ_D : for Gaussian in the derivative calculation σ_I : for Gaussian in the windowing function



4. Cornerness function - two strong eigenvalues

 $\theta = \det[M(\sigma_{I}, \sigma_{D})] - \alpha[\operatorname{trace}(M(\sigma_{I}, \sigma_{D}))]^{2}$ = $g(I_{x}^{2})g(I_{y}^{2}) - [g(I_{x}I_{y})]^{2} - \alpha[g(I_{x}^{2}) + g(I_{y}^{2})]^{2}$

5. Perform non-maximum suppression





Lecture 6 - 99

• Input Image



Raymond Yu

Lecture 6 - 100

- Input Image
- Compute corner response function θ



Raymond Yu

Lecture 6^{Slide}101^{ted from Daya Frolova}Aprilina5, 2025

- Input Image
- Compute corner response function θ
- Take only the local maxima of θ,
 where θ > threshold



April 15, 2025

Raymond Yu

- Input Image
- Compute corner response function θ
- Take only the local maxima of θ,
 where θ > threshold



Raymond Yu

Lecture 6 - 103

Harris Detector – Responses [Harris88]



Effect: A very precise corner detector.



Raymond Yu



Harris Detector – Responses [Harris88]



Raymond Yu



Harris Detector – Responses [Harris88]



Raymond Yu

• Results are great for finding correspondences matches between images



Summary

- Local Invariant Features
- Harris Corner Detector







Harris Detector: Properties

• Translation invariance?




Harris Detector: Properties

- Translation invariance
- Rotation invariance?



Ellipse rotates but its shape (i.e. eigenvalues) remains the same

It is invariant to image rotation

Raymond Yu

Lecture 6 - 109

April 15, 2025

Harris Detector: Properties

- Translation invariance
- Rotation invariance
- Scale invariance?



Raymond Yu

Lecture 6 - 110

Slide crediApren 45, 2025

Next time

Detectors and Descriptors

Raymond Yu

Lecture 6 - 111

