


CS5670: Computer Vision

Noah Snavely

Lecture 6: Feature matching



Reading


- Szeliski: 4.1

Announcements


- Project 1 artifact voting online shortly
- Project 2 to be released soon
- Quiz at the beginning of class today

Local features: main components


1) Detection: Identify the interest points





2) Description: Extract vector feature descriptor surrounding each interest point.



3) Matching: Determine correspondence between descriptors in two views





SIFT Example



868 SIFT features

Which features match?



Feature matching


Given a feature in I_1 , how to find the best match in I_2 ?

1. Define distance function that compares two descriptors
2. Test all the features in I_2 , find the one with min distance


Feature distance

How to define the difference between two features f_1, f_2 ?

- Simple approach: L_2 distance, $\|f_1 - f_2\|$ (aka SSD)
- can give good scores to ambiguous (incorrect) matches



I_1




I_2


Feature distance

How to define the difference between two features f_1, f_2 ?

- Better approach: ratio distance = $\|f_1 - f_2\| / \|f_1 - f_2'\|$
 - f_2 is best SSD match to f_1 in I_2
 - f_2' is 2nd best SSD match to f_1 in I_2
 - gives large values for ambiguous matches



I_1




I_2

Feature distance


- Does the SSD vs “ratio distance” change the best match to a given feature in image 1?

Feature matching example



51 matches (thresholded by ratio score)


Feature matching example



58 matches (thresholded by ratio score)

Evaluating the results


How can we measure the performance of a feature matcher?



feature distance

True/false positives

How can we measure the performance of a feature matcher?



The distance threshold affects performance


- True positives = # of detected matches that are correct
 - Suppose we want to maximize these—how to choose threshold?
- False positives = # of detected matches that are incorrect
 - Suppose we want to minimize these—how to choose threshold?

Evaluating the results

How can we measure the performance of a feature matcher?

$$\frac{\text{\# true positives}}{\text{\# matching features (positives)}}$$

“recall”

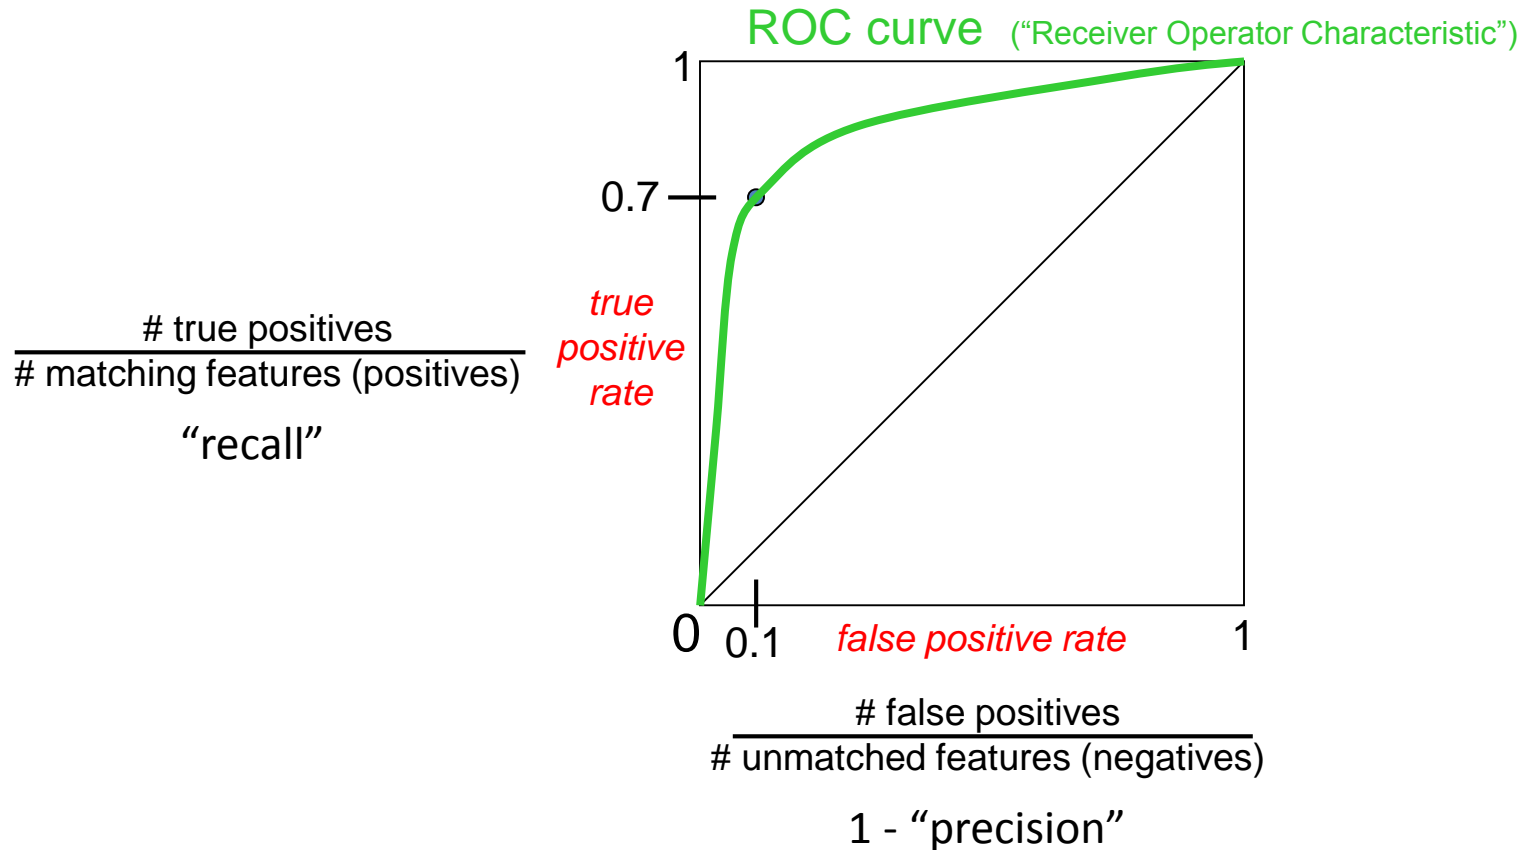


$$\frac{\text{\# false positives}}{\text{\# unmatched features (negatives)}}$$

1 - “precision”

Evaluating the results

How can we measure the performance of a feature matcher?



Available at a web site near you...

- For most local feature detectors, executables are available online:
 - <http://www.robots.ox.ac.uk/~vgg/research/affine>
 - <http://www.cs.ubc.ca/~lowe/keypoints/>
 - <http://www.vision.ee.ethz.ch/~surf>

Questions?