

Patch Descriptors 1

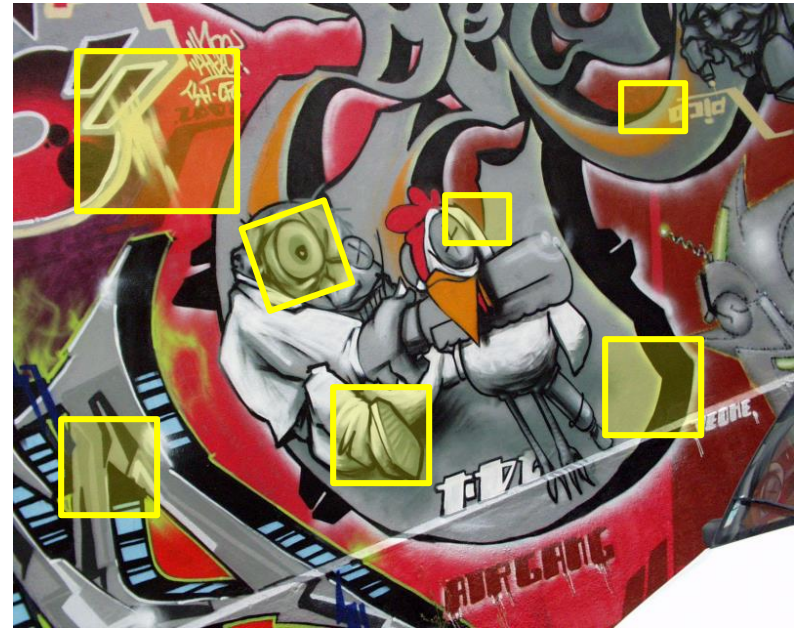
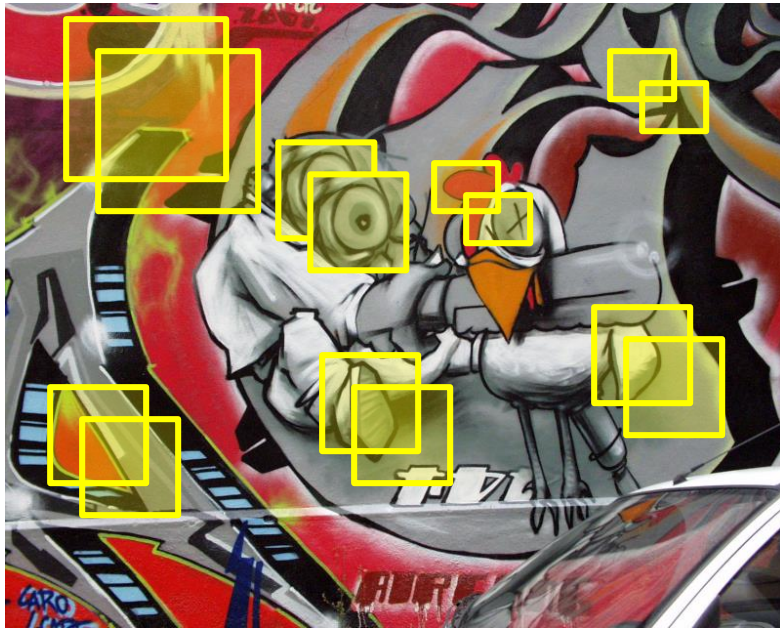
ECE P 596

Linda Shapiro

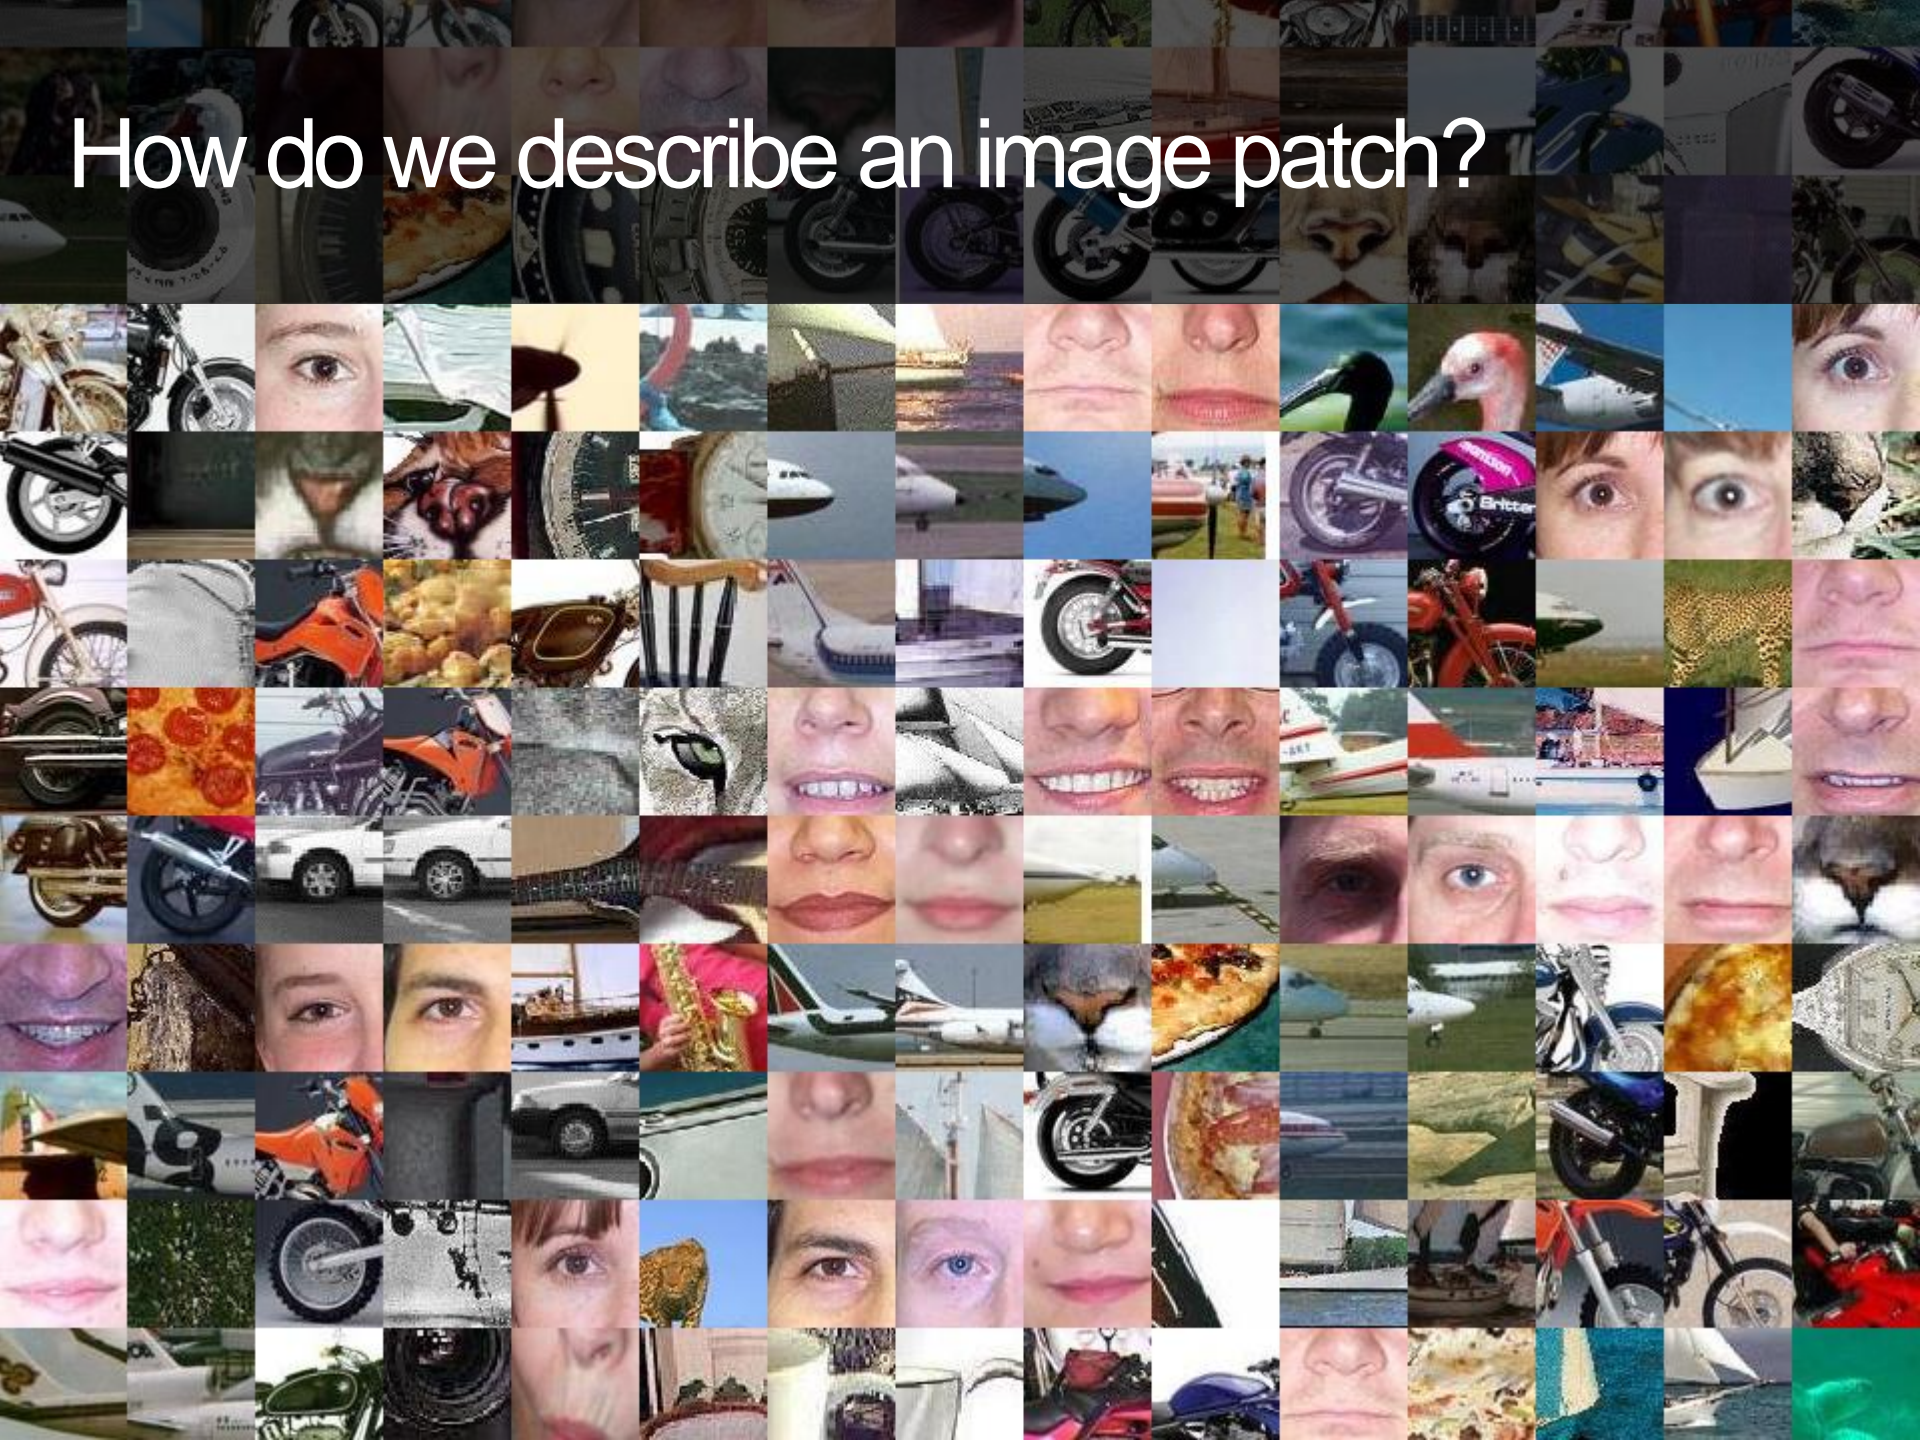
How can we find corresponding points?



How can we find correspondences?

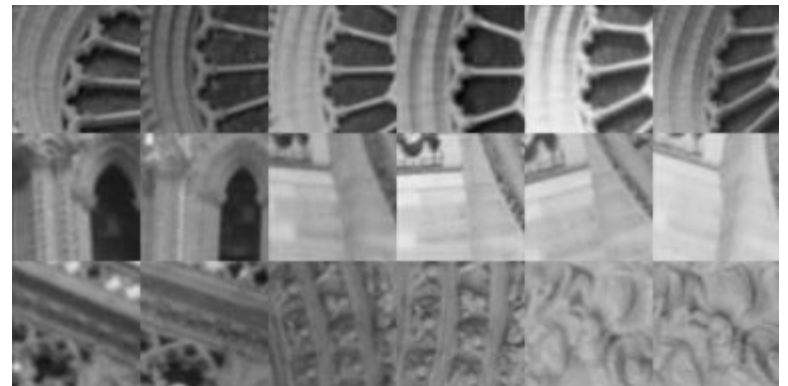


How do we describe an image patch?

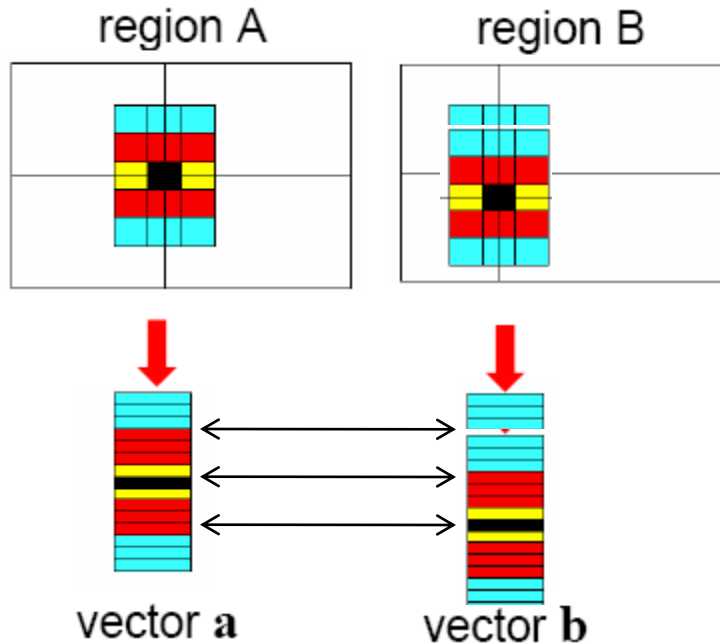


How do we describe an image patch?

Patches with similar content should have similar descriptors.



Raw patches as local descriptors



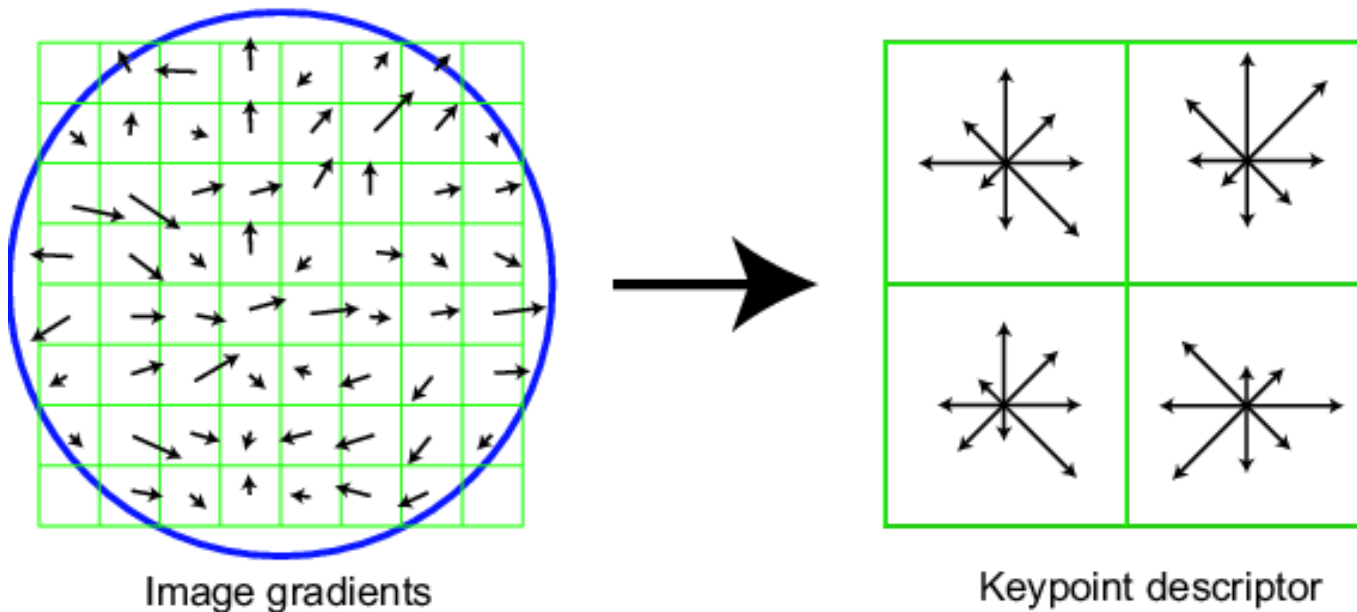
The simplest way to describe the neighborhood around an interest point is to write down the list of intensities to form a **feature vector**.

But this is very sensitive to even small shifts, rotations.

SIFT descriptor

Full version

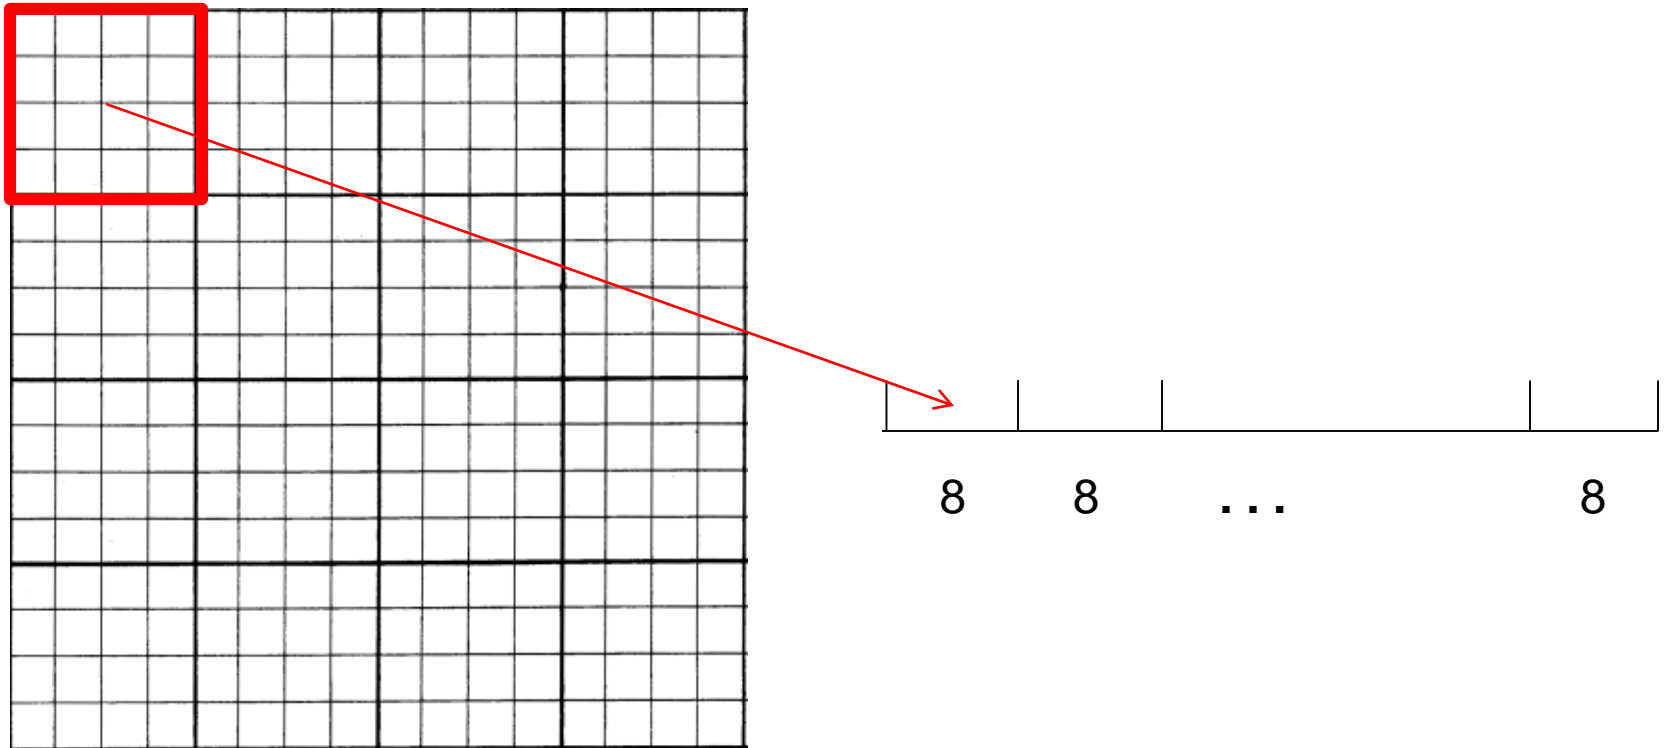
- Divide the 16x16 window into a 4x4 grid of cells (2x2 case shown below)
- Compute an **orientation histogram** for each cell
- 16 cells * 8 orientations = 128 dimensional descriptor



SIFT descriptor

Full version

- Divide the **16x16 window** into a 4x4 grid of cells
- Compute an **orientation histogram** for each cell
- 16 cells * 8 orientations = **128 dimensional descriptor**



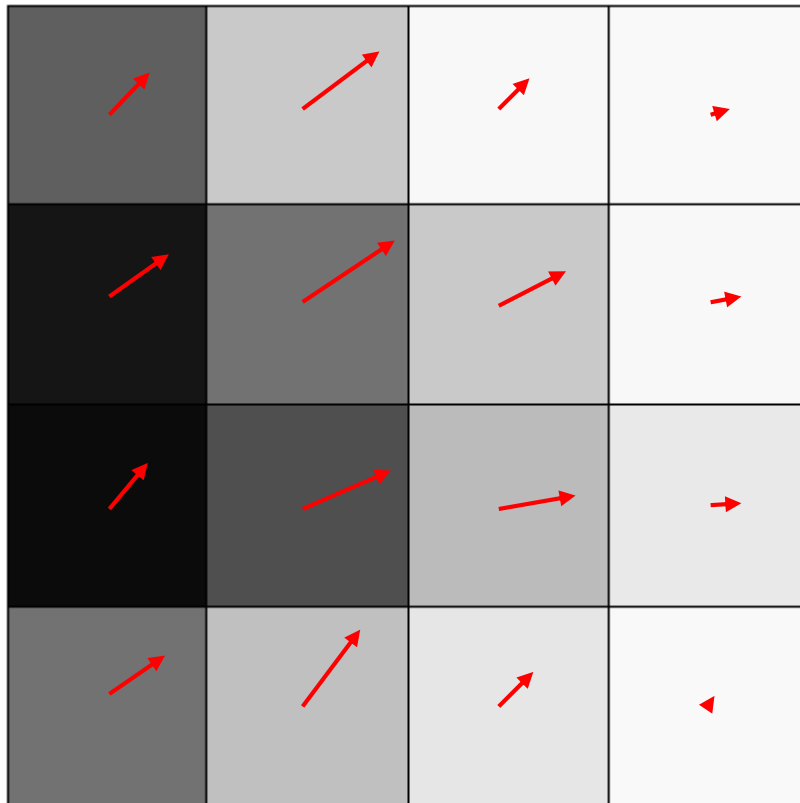
Numeric Example

0.37	0.79	0.97	0.98
0.08	0.45	0.79	0.97
0.04	0.31	0.73	0.91
0.45	0.75	0.90	0.98

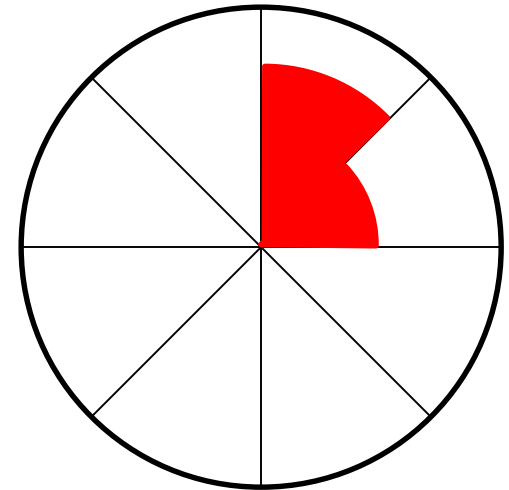


$$\text{magnitude}(x, y) = \sqrt{(L(x+1, y) - L(x-1, y))^2 + (L(x, y+1) - L(x, y-1))^2}$$

$$\theta(x, y) = \text{atan}\left(\frac{L(x, y+1) - L(x, y-1)}{L(x+1, y) - L(x-1, y)}\right)$$



Orientations in each of the 16 pixels of the cell



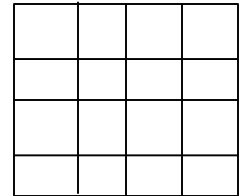
The orientations all ended up in two bins: 11 in one bin, 5 in the other. (rough count)

5 11 0 0 0 0 0 0

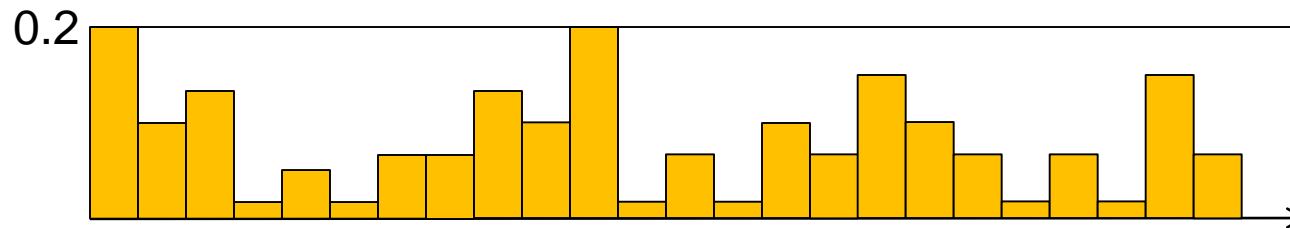
SIFT descriptor

Full version

- Start with a 16x16 window (256 pixels)
- Divide the 16x16 window into a 4x4 grid of cells (16 cells)
- Compute an orientation histogram for each cell
- 16 cells * 8 orientations = 128 dimensional descriptor
- Threshold normalize the descriptor:



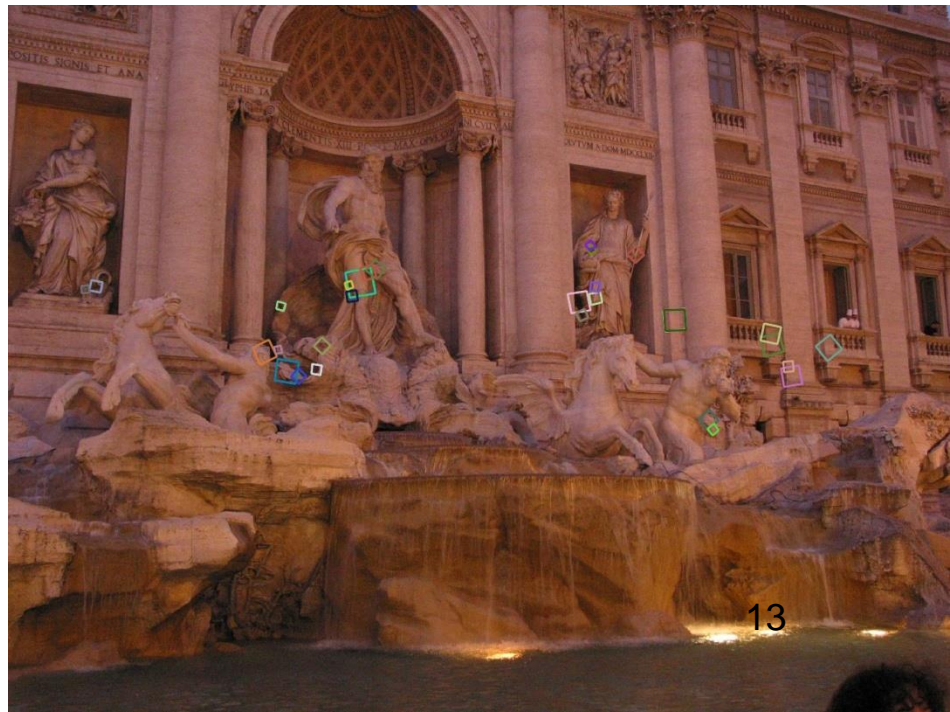
$$\sum_i d_i^2 = 1 \quad \text{such that: } d_i < 0.2$$



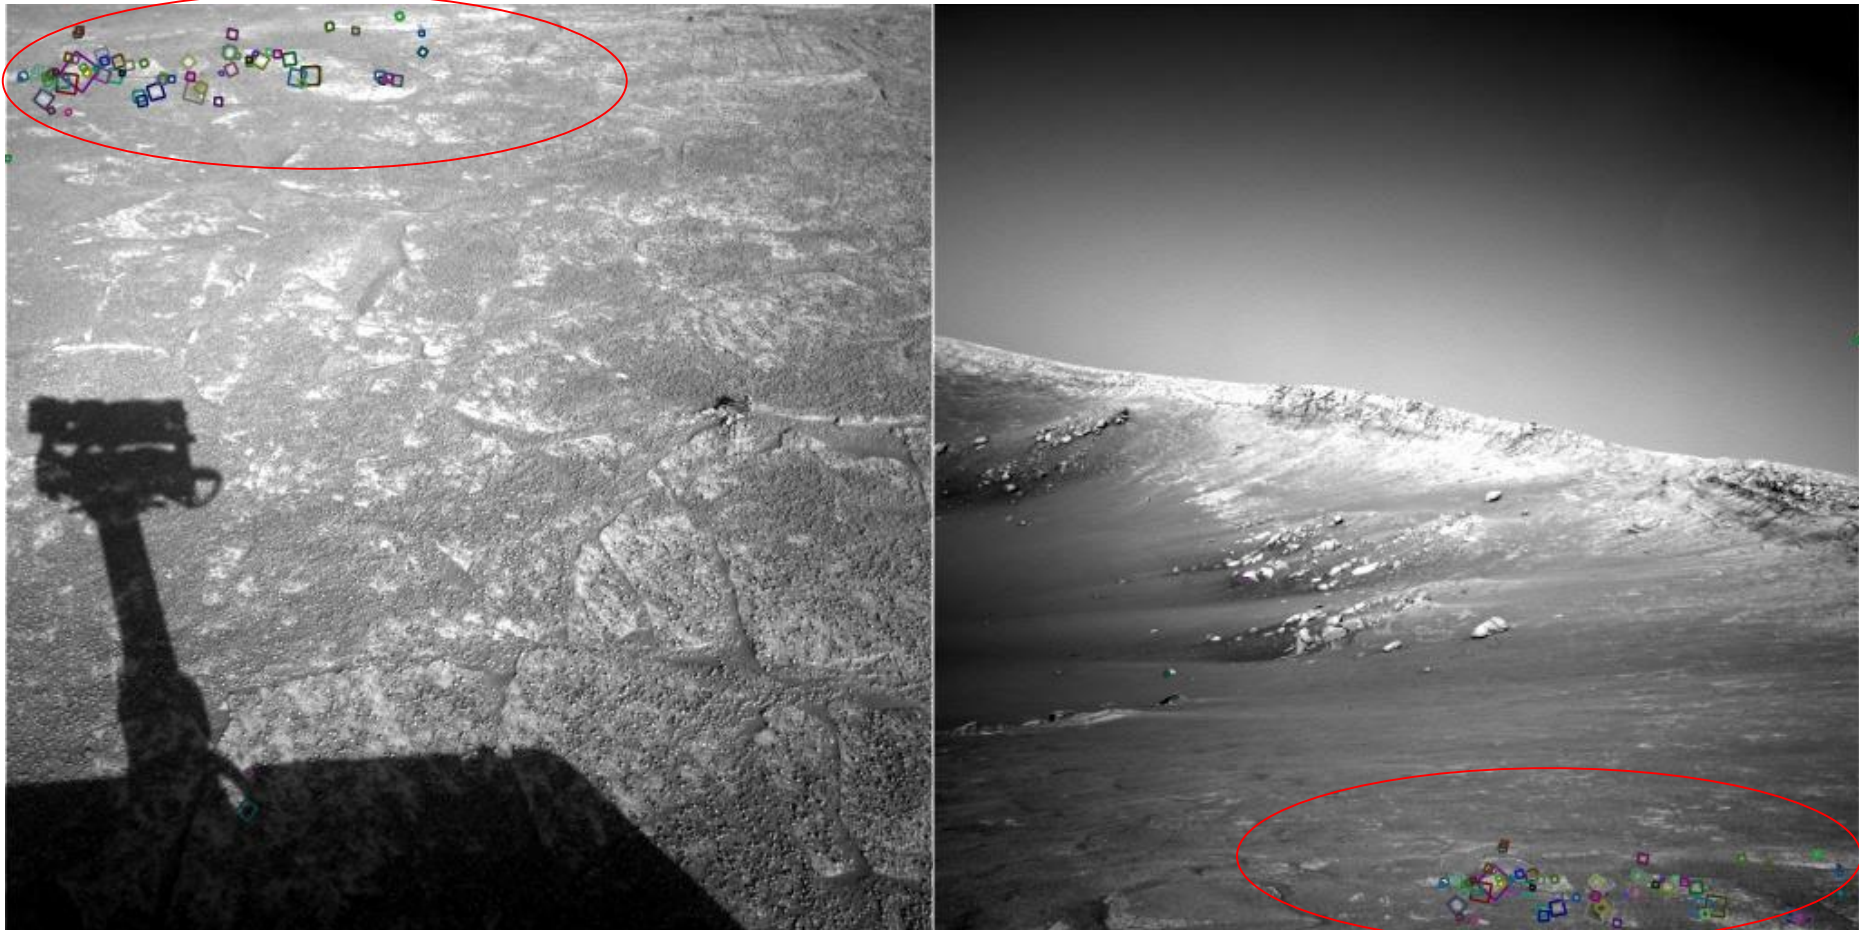
Properties of SIFT

Extraordinarily robust matching technique

- Can handle changes in viewpoint
 - Up to about 30 degree out of plane rotation
- Can handle significant changes in illumination
 - Sometimes even day vs. night (below)
- Fast and efficient—can run in real time
- Various code available
 - <http://www.cs.ubc.ca/~lowe/keypoints/>



Example



NASA Mars Rover images
with SIFT feature matches
Figure by Noah Snaveley

Example: Object Recognition



SIFT is extremely powerful for object instance recognition, especially for well-textured objects

Example: Google Goggle

Google Goggles in Action

Click the icons below to see the different ways Google Goggles can be used.



[Landmark](#)



[Book](#)



[Contact Info.](#)



[Artwork](#)



[Places](#)



[Wine](#)



[Logo](#)



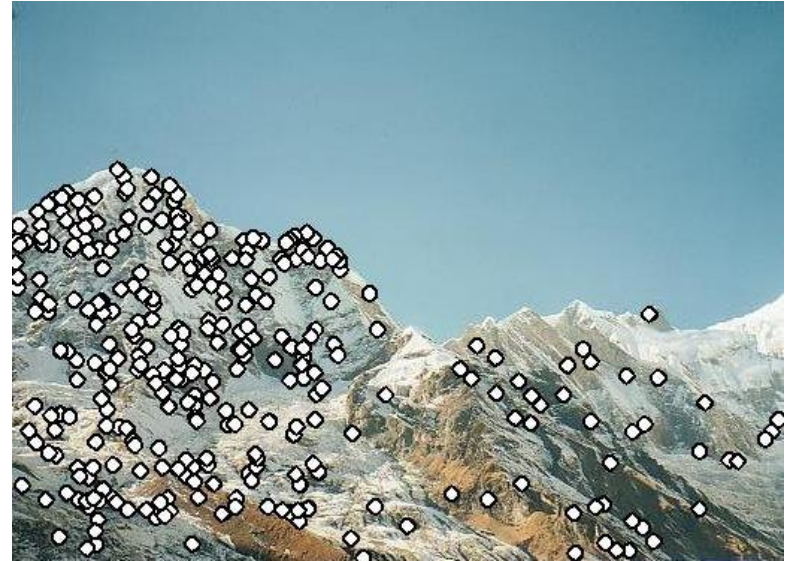
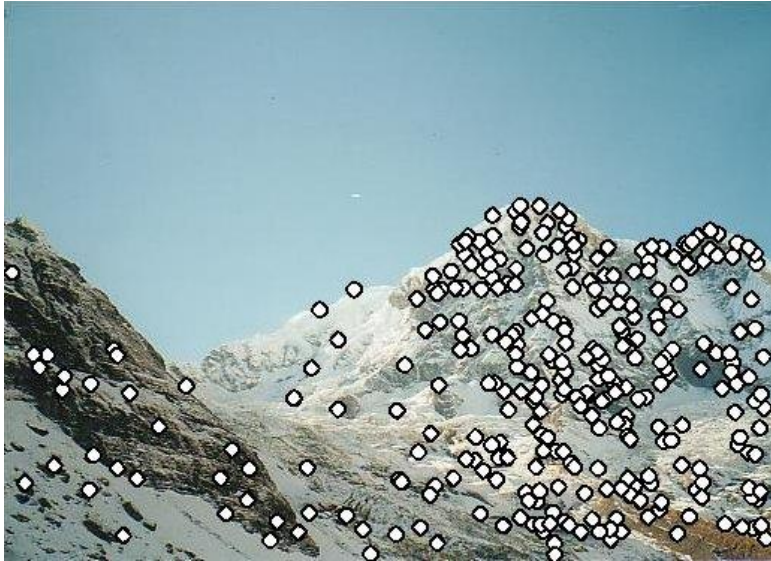
panorama?

- We need to match (align) images



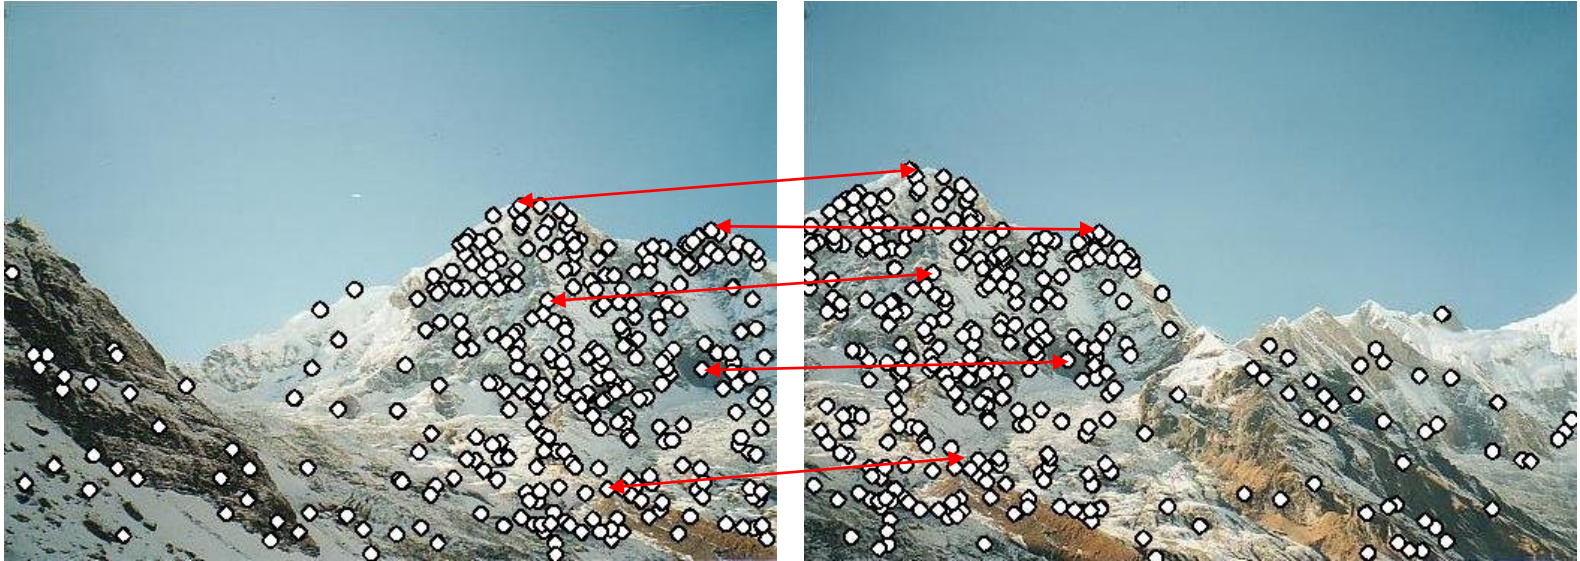
Matching with Features

- Detect feature points in both images



Matching with Features

- Detect feature points in both images
- Find corresponding pairs

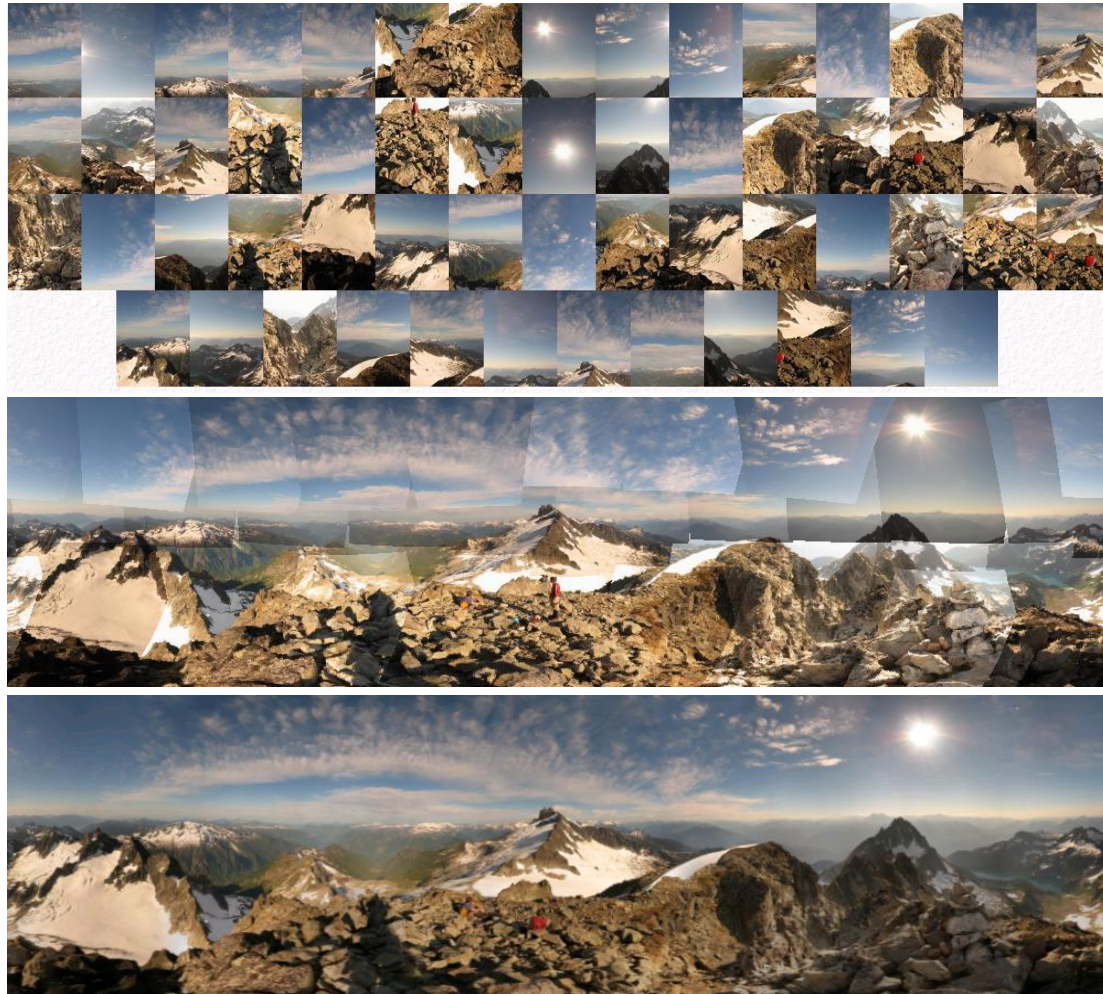


Matching with Features

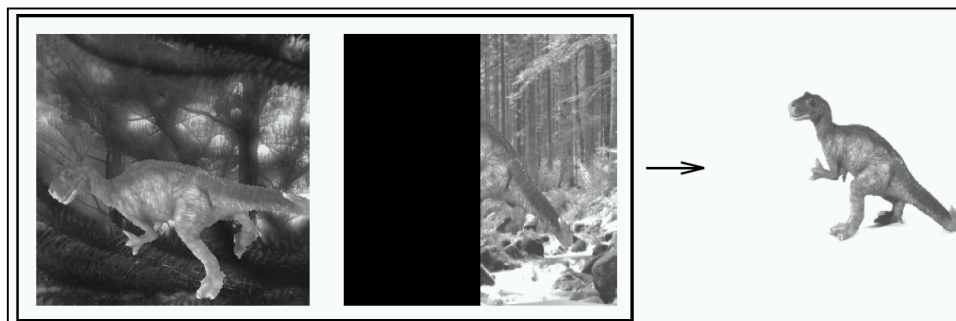
- Detect feature points in both images
- Find corresponding pairs
- Use these matching pairs to align images - the required mapping is called a **homography**.



Automatic mosaicing



Recognition of specific objects, scenes



Schmid and Mohr 1997



Sivic and Zisserman, 2003



Rothganger et al. 2003



Lowe 2002

Example: 3D Reconstructions

- Photosynth (also called Photo Tourism) developed at UW by Noah Snavely, Steve Seitz, Rick Szeliski and others

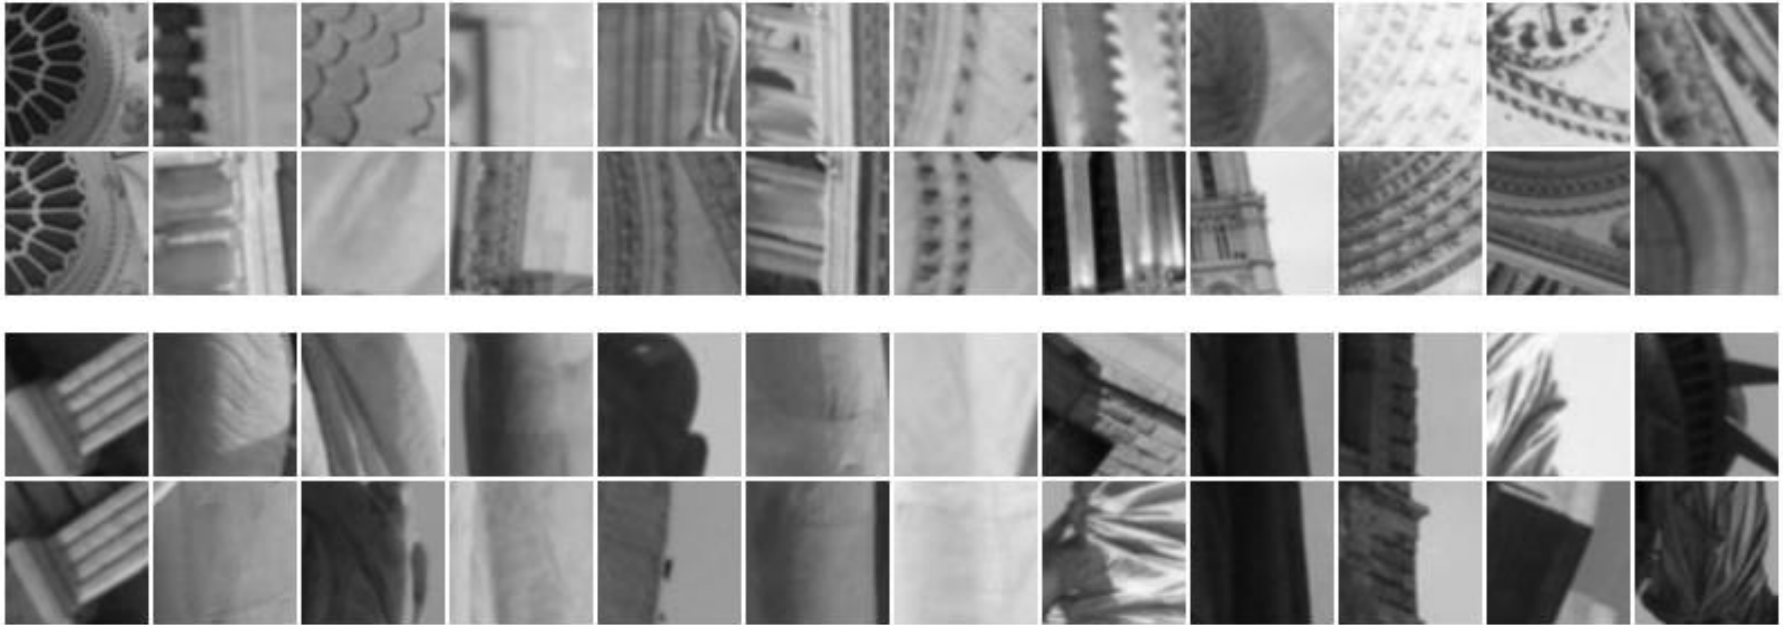
<http://www.youtube.com/watch?v=p16frKJLVi0>

- Building Rome in a day, developed at UW by Sameer Agarwal, Noah Snavely, Steve Seitz and others

http://www.youtube.com/watch?v=kxtQqYLRaSQ&feature=player_embedded

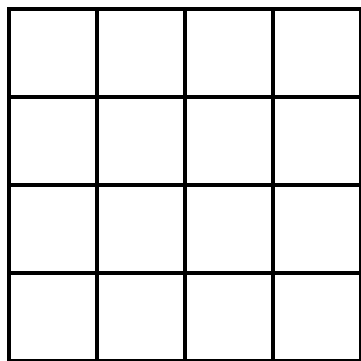
When does the SIFT descriptor fail?

Patches SIFT thought were the same but aren't:

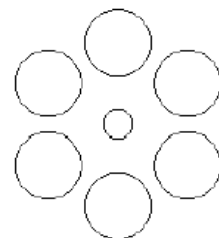


Other methods: Daisy

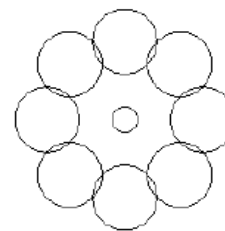
Circular gradient binning



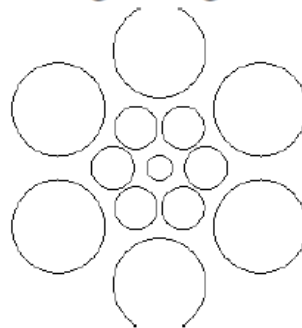
SIFT



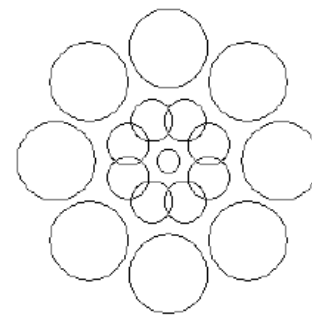
1 Ring 6 Segments



1 Ring 8 Segments



2 Rings 6 Segments



2 Rings 8 Segments

Daisy

Other methods: SURF

For **computational efficiency** only compute gradient histogram with 4 bins:

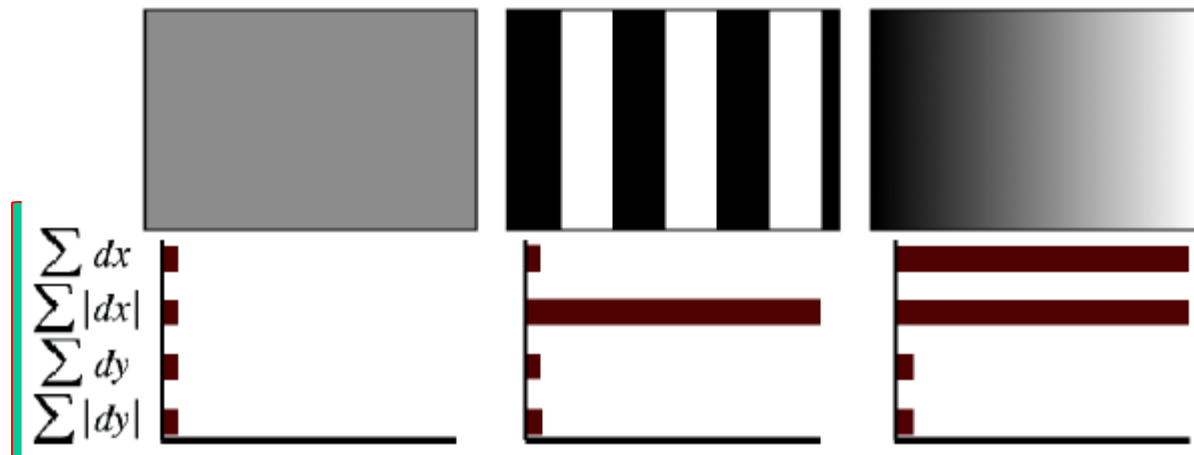


Fig. 3. The descriptor entries of a sub-region represent the nature of the underlying intensity pattern. Left: In case of a homogeneous region, all values are relatively low. Middle: In presence of frequencies in x direction, the value of $\sum |d_x|$ is high, but all others remain low. If the intensity is gradually increasing in x direction, both values $\sum d_x$ and $\sum |d_x|$ are high.

SURF: Speeded Up Robust Features

Herbert Bay, Tinne Tuytelaars, and Luc Van Gool, ECCV 2006

Other methods: BRIEF

Randomly sample pair of pixels a and b .
1 if $a > b$, else 0. Store binary vector.

011000111000...

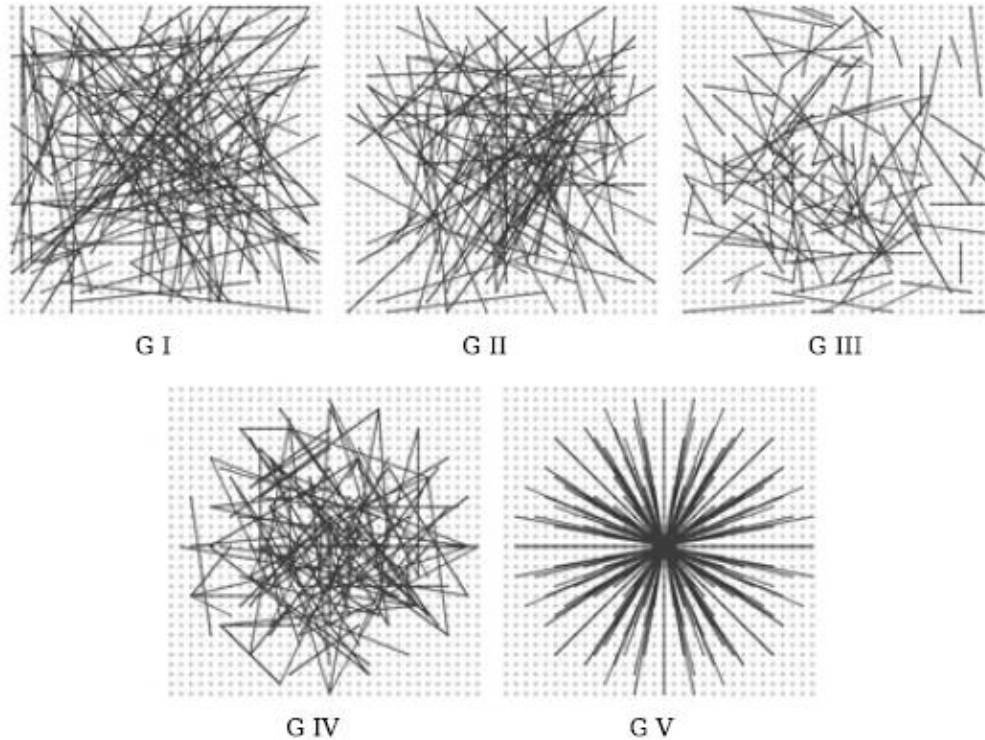
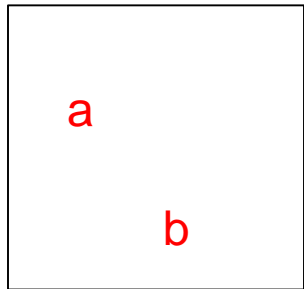


Fig. 2. Different approaches to choosing the test locations. All except the rightmost one are selected by random sampling. Showing 128 tests in every image.

BRIEF: binary robust independent elementary features,
Calonder, V Lepetit, C Strecha, ECCV 2010

Descriptors and Matching

- The SIFT descriptor and the various variants are used to **describe** an image patch, so that we can match two image patches.
- In addition to the descriptors, we need a **distance measure** to calculate how different the two patches are?



?



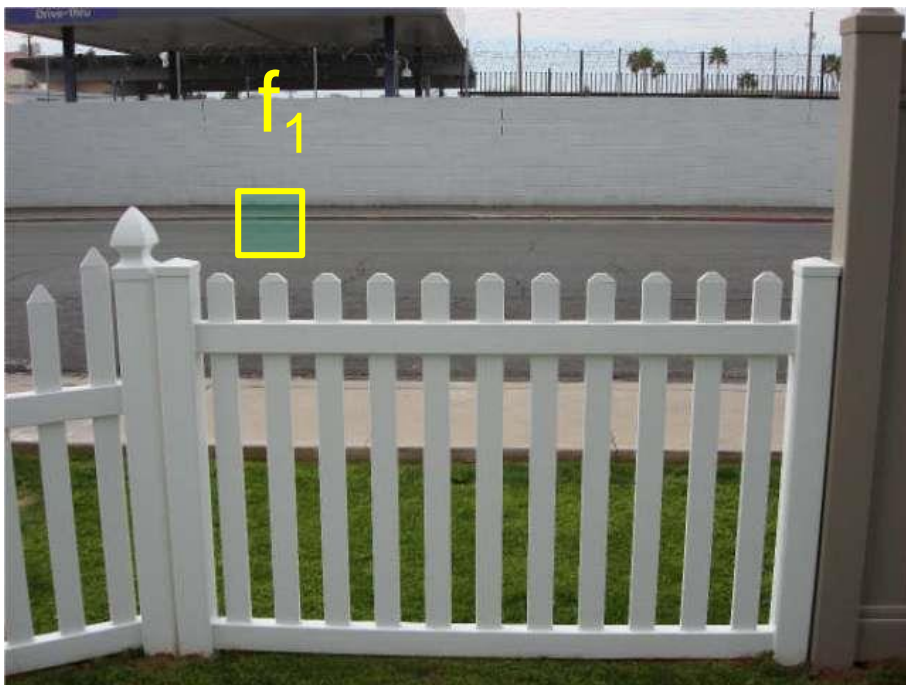
Feature distance

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

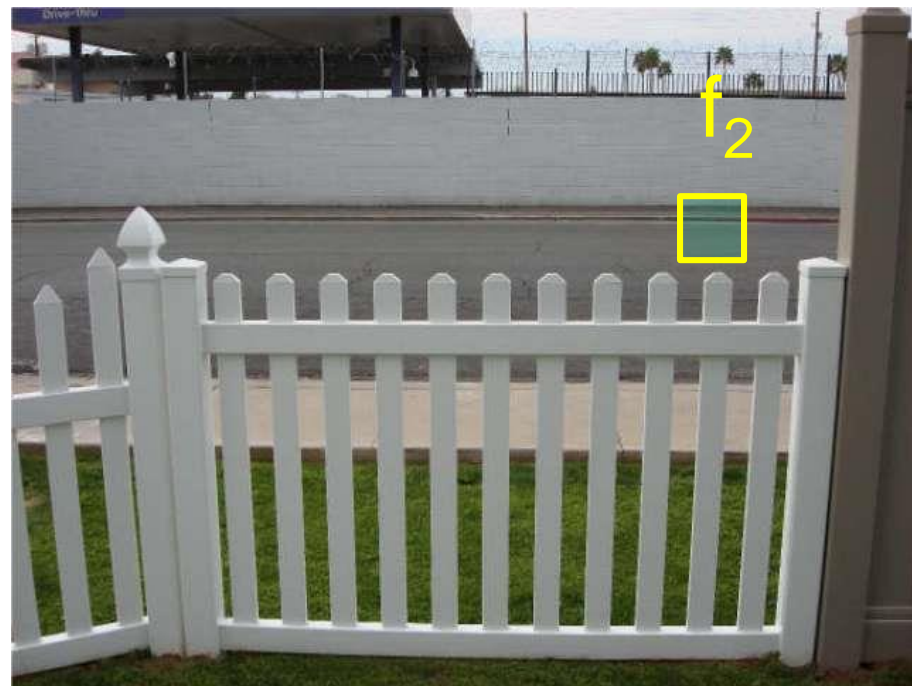
- Simple approach is $\text{SSD}(f_1, f_2)$
 - sum of square differences between entries of the two descriptors

$$\sum_i (f_{1i} - f_{2i})^2$$

- But it can give good scores to very ambiguous (bad) matches



I_1

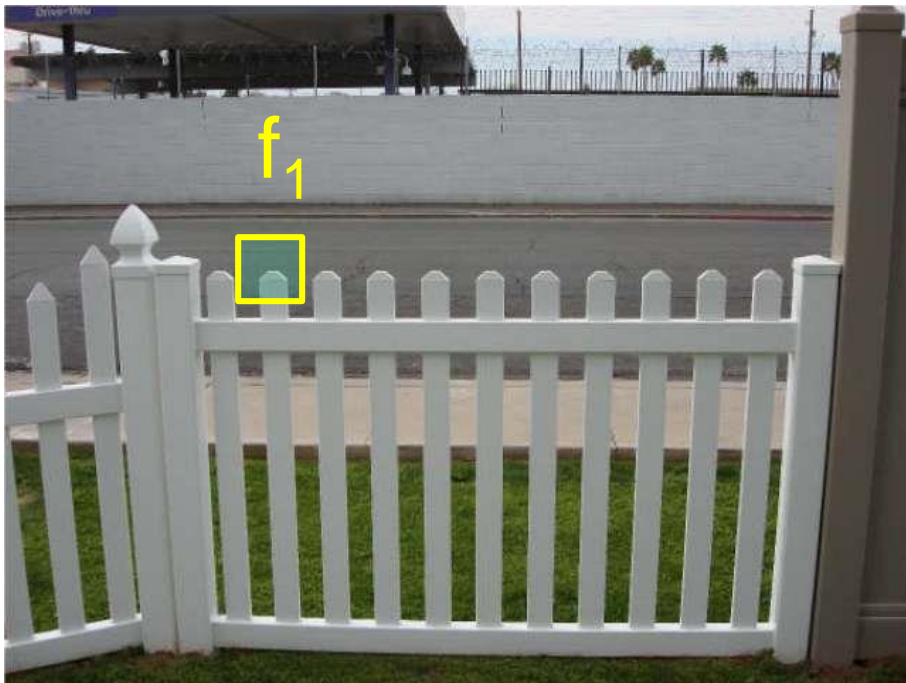


I_2

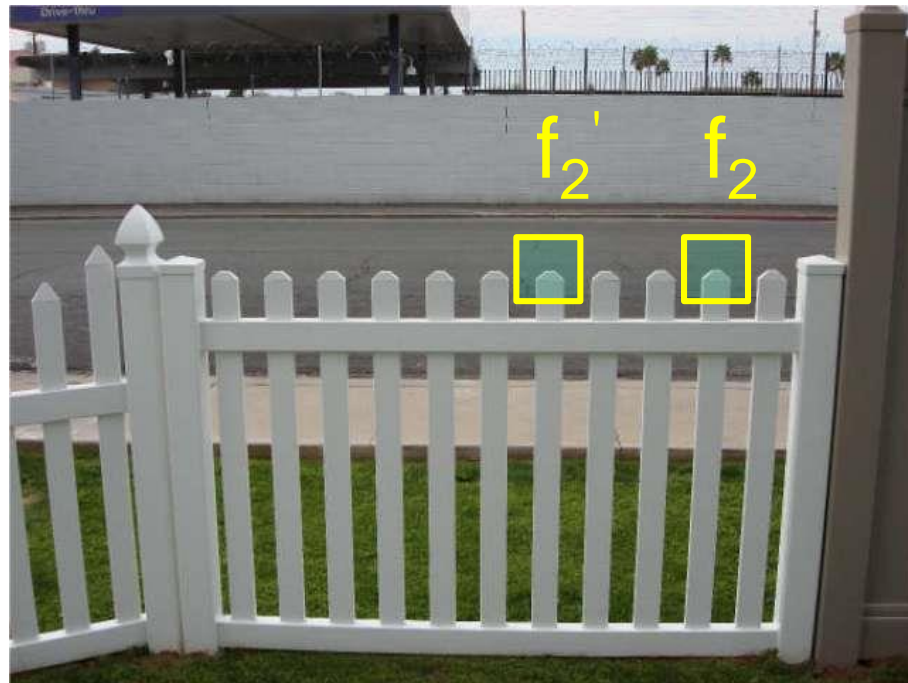
Feature distance in practice

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

- Better approach: ratio distance = $\text{SSD}(f_1, f_2) / \text{SSD}(f_1, f_2')$
 - 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 (~ 1) for ambiguous matches WHY?

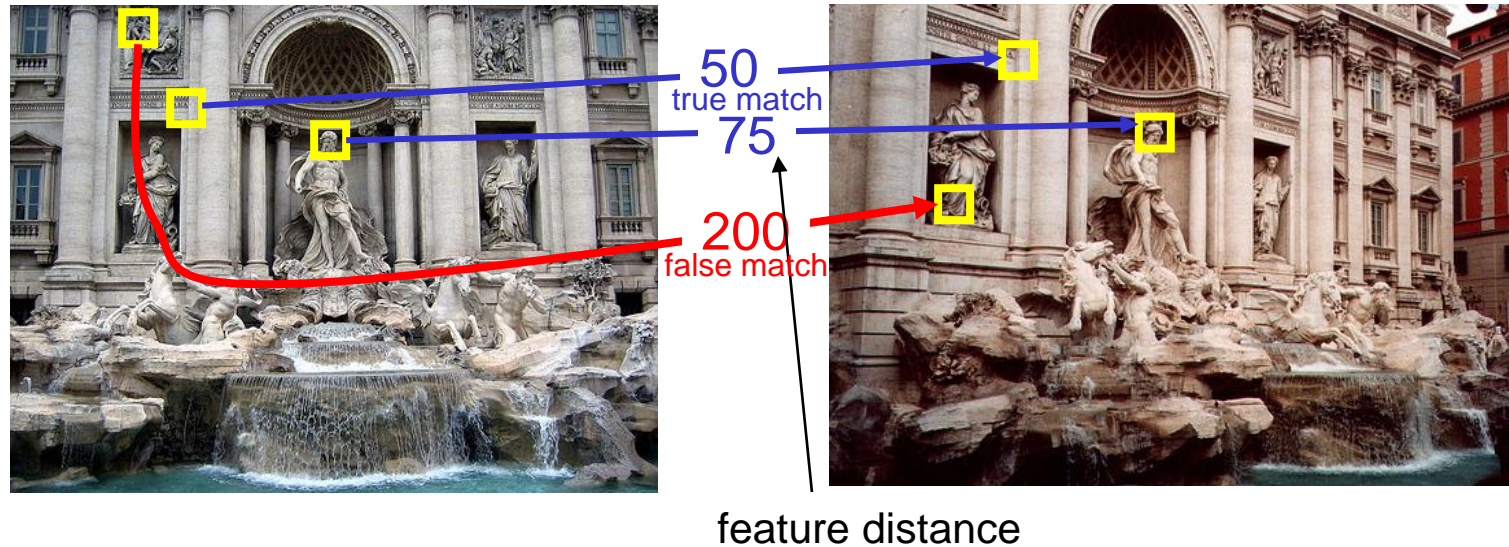


I_1



I_2

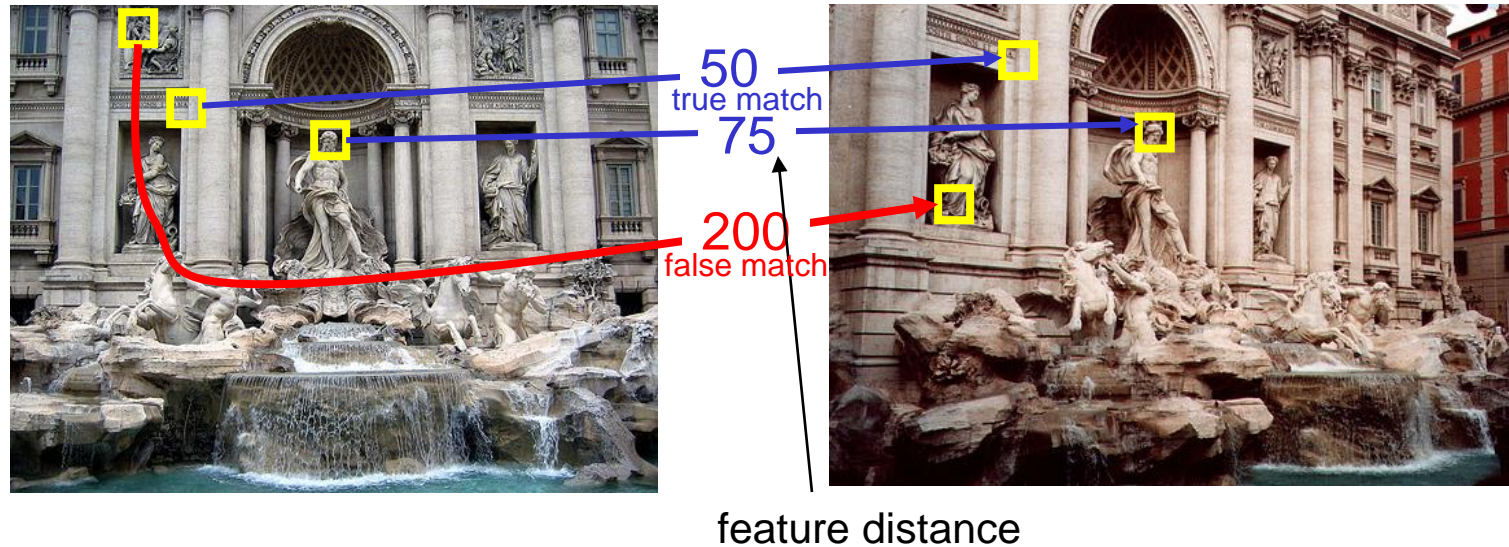
Eliminating more bad matches



Throw out features with distance $>$ threshold

- How to choose the threshold?

True/false positives



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?

Finale

- Describing images or image patches is very important for matching and recognition
- The SIFT descriptor was invented in 1999 and is still very heavily used.
- Other descriptors are also available, some much simpler, but less powerful.