Another kind of feature: blob

- “Peak” or “Valley” in intensity

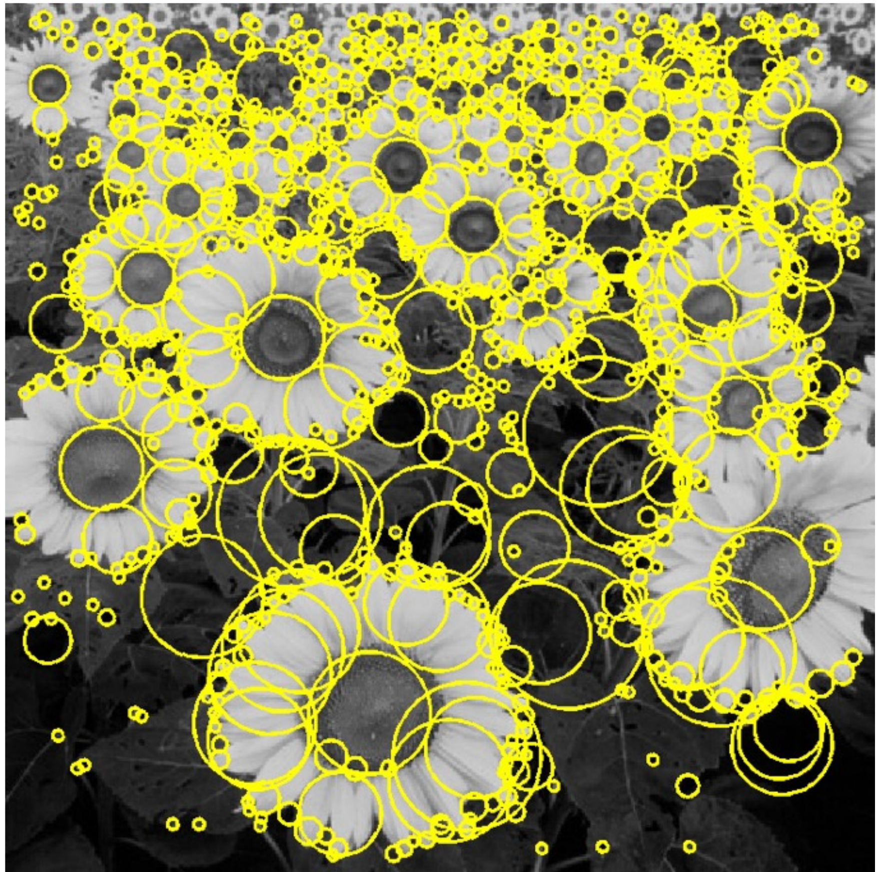

Detecting peaks and valleys

- The *Laplacian of Gaussian* (*LoG*)

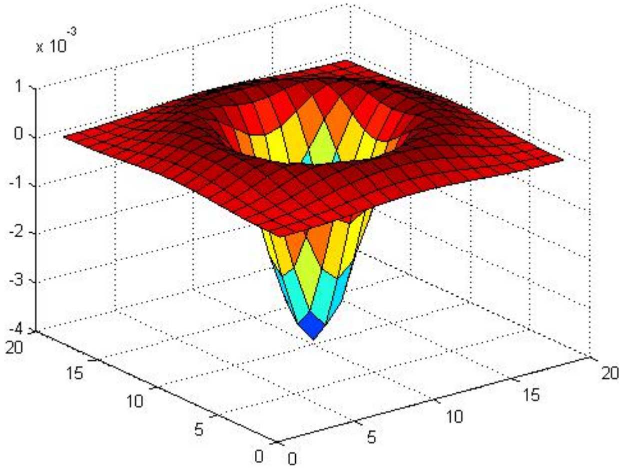

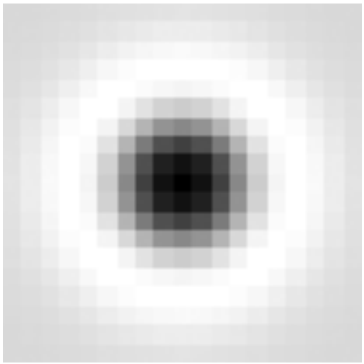

$$\nabla^2 g = \frac{\partial^2 g}{\partial x^2} + \frac{\partial^2 g}{\partial y^2}$$

(very similar to a Difference of Gaussians (DoG) – i.e. a Gaussian minus a slightly smaller Gaussian)

Scale selection

- At what scale does the Laplacian achieve a maximum response for a binary circle of radius r?

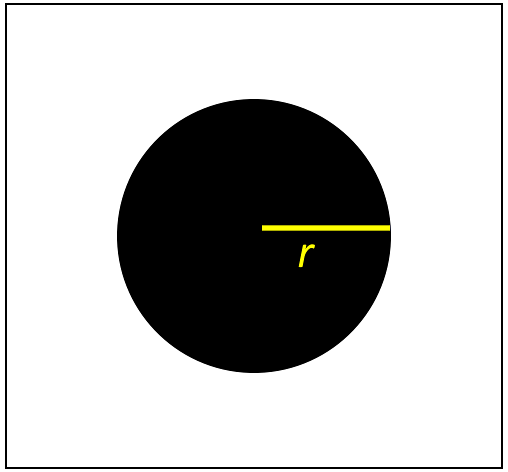

image

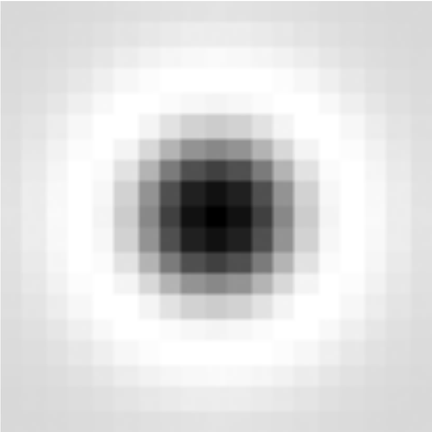

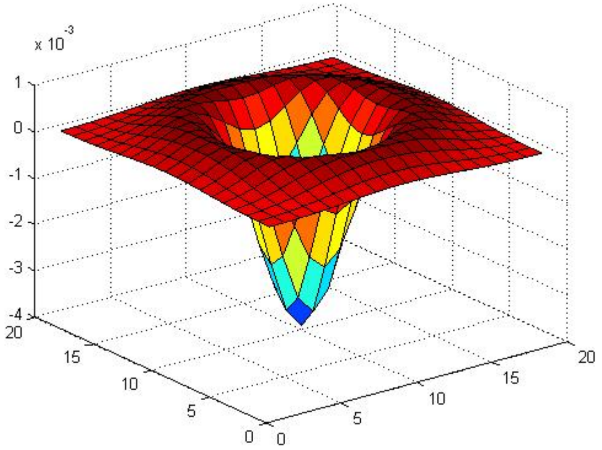

Laplacian

Laplacian of Gaussian

- “Blob” detector

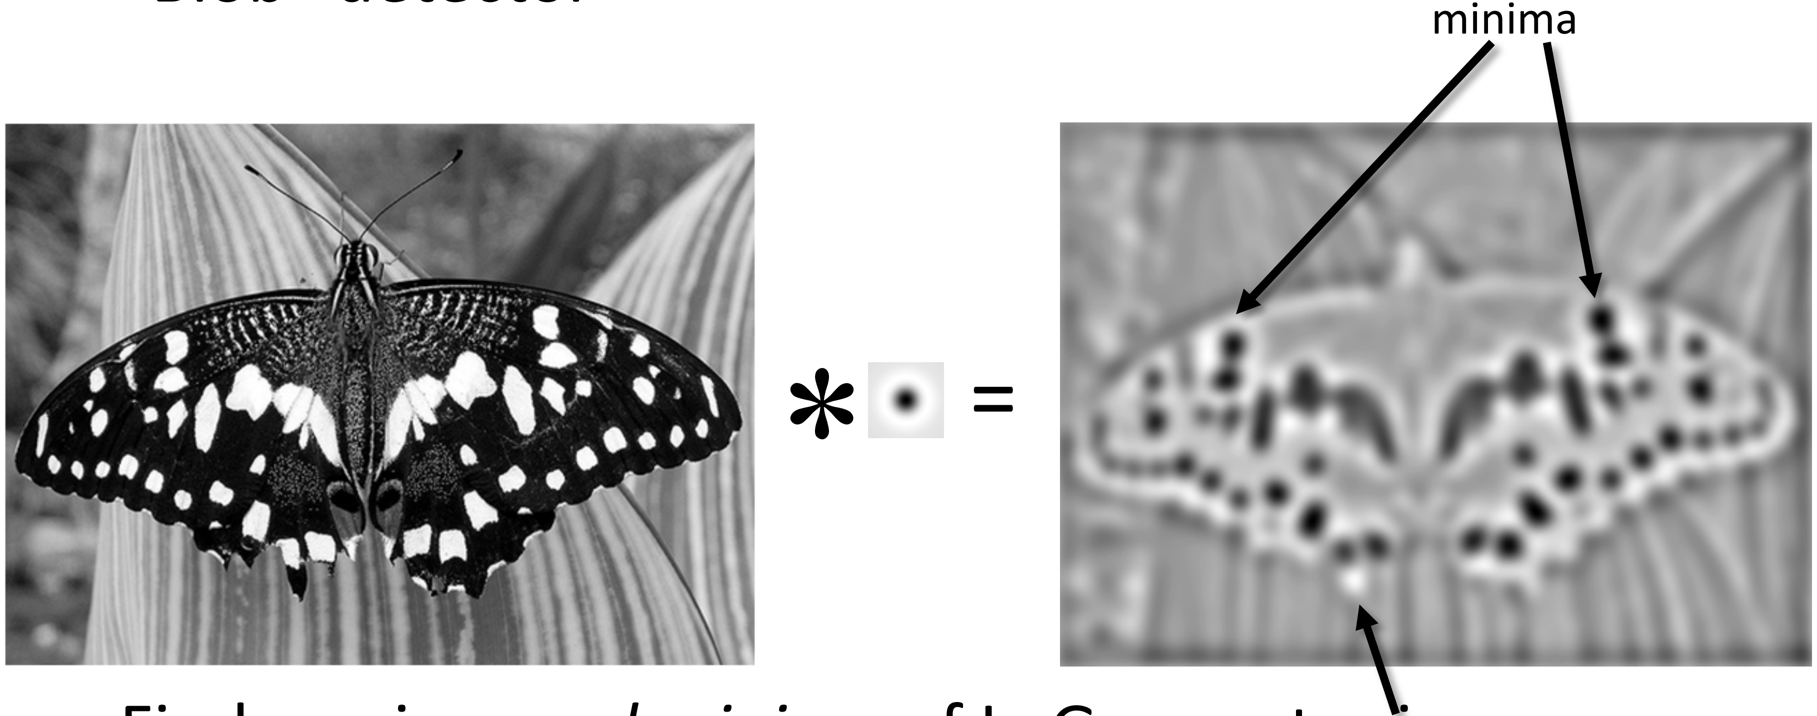

- Find maxima *and minima* of LoG operator in space and scale

Characteristic scale

- The scale that produces peak of Laplacian response

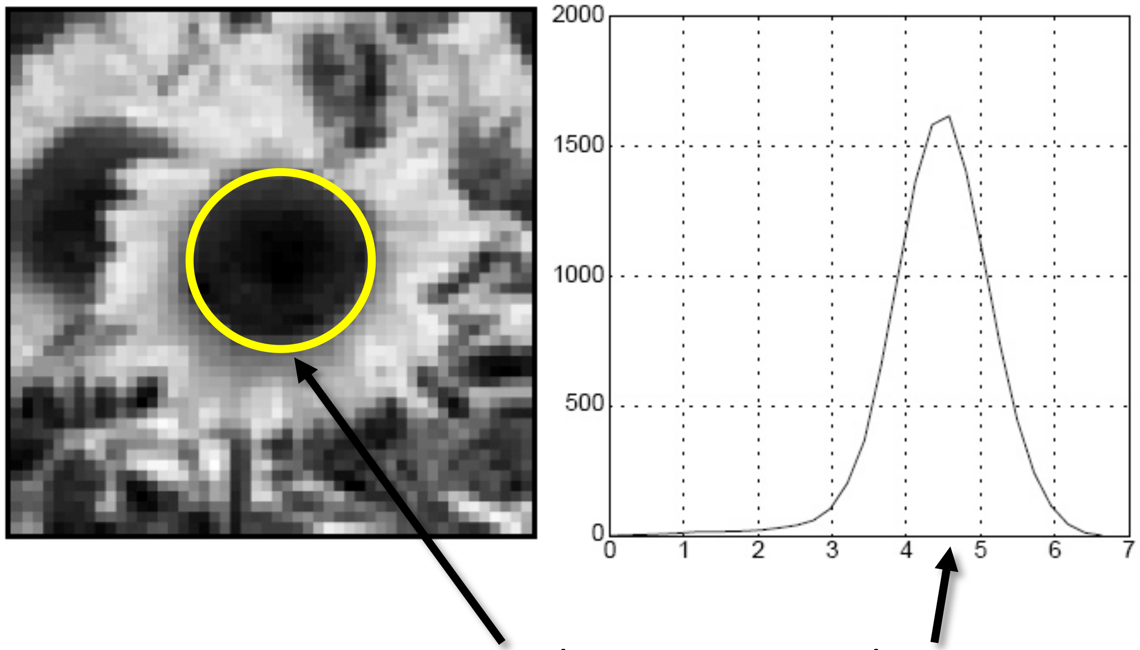

characteristic scale

T. Lindeberg (1998). "Feature detection with automatic scale selection." *International Journal of Computer Vision* **30** (2): pp 77--116.

Find local maxima in position-scale space

$$L_{xx}(\sigma) + L_{yy}(\sigma)$$

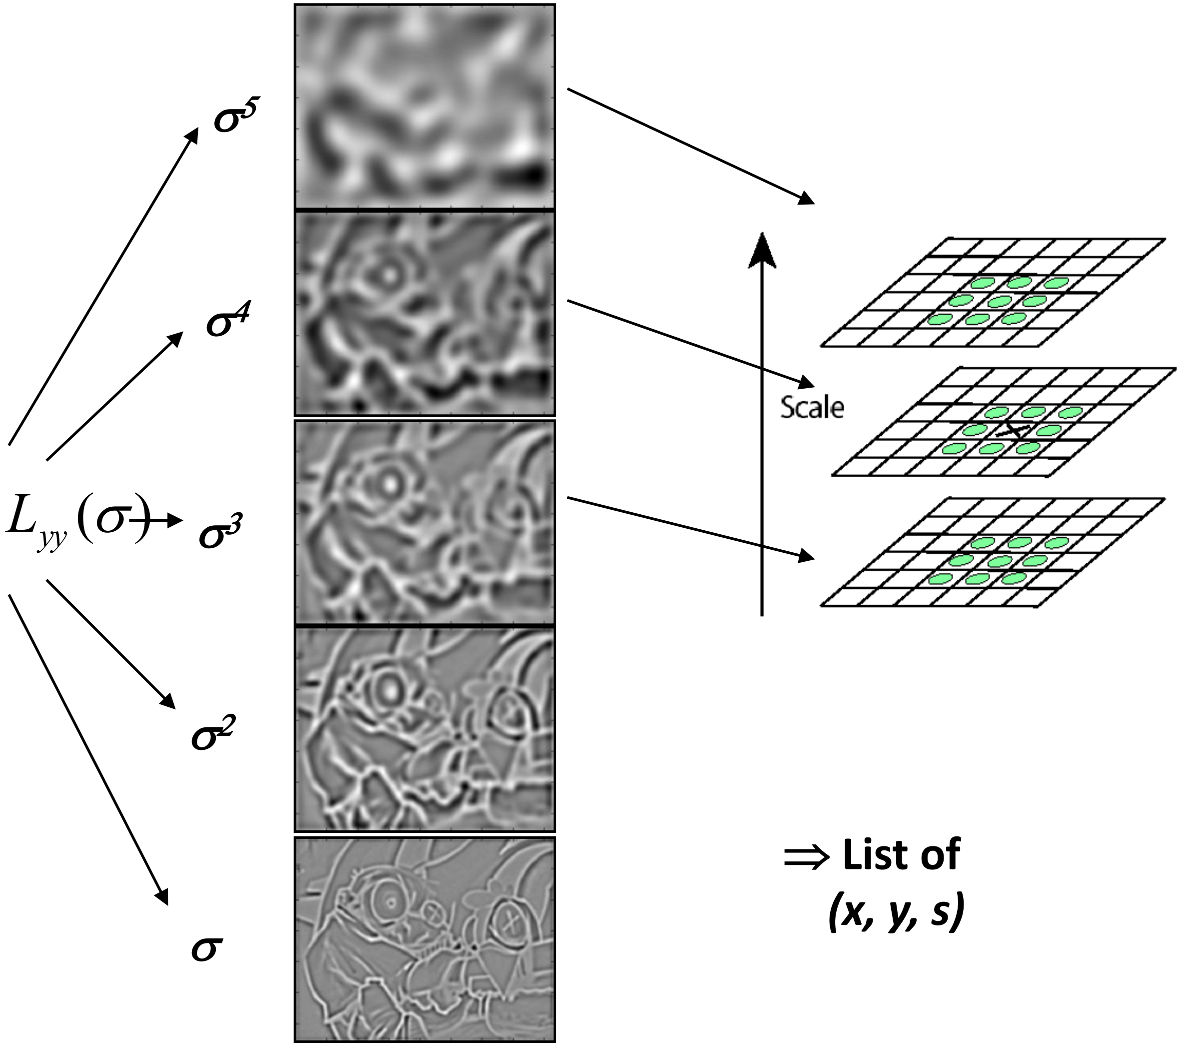

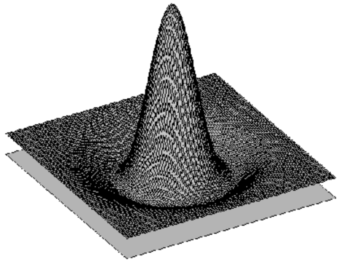

⇒ List of
(x, y, s)

Scale-space blob detector: Example

Scale-space blob detector: Example

sigma = 11.9912

Scale-space blob detector: Example

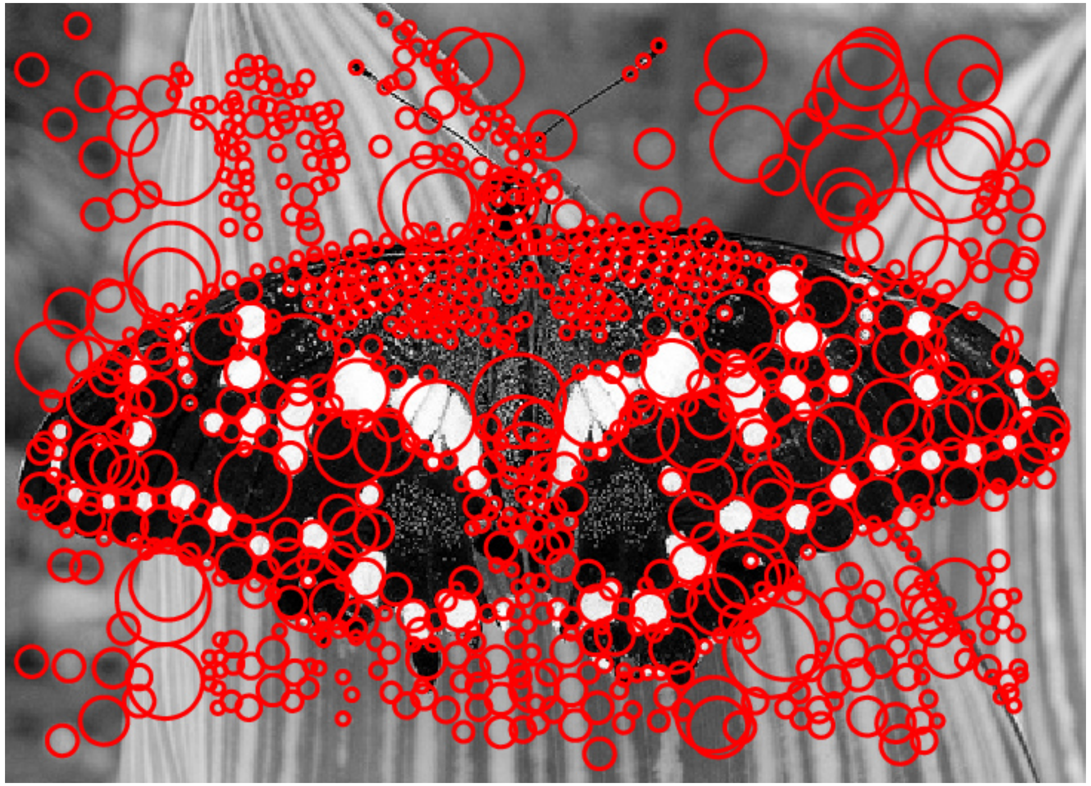

Matching feature points

We know how to detect good points
Next question: **How to match them?**

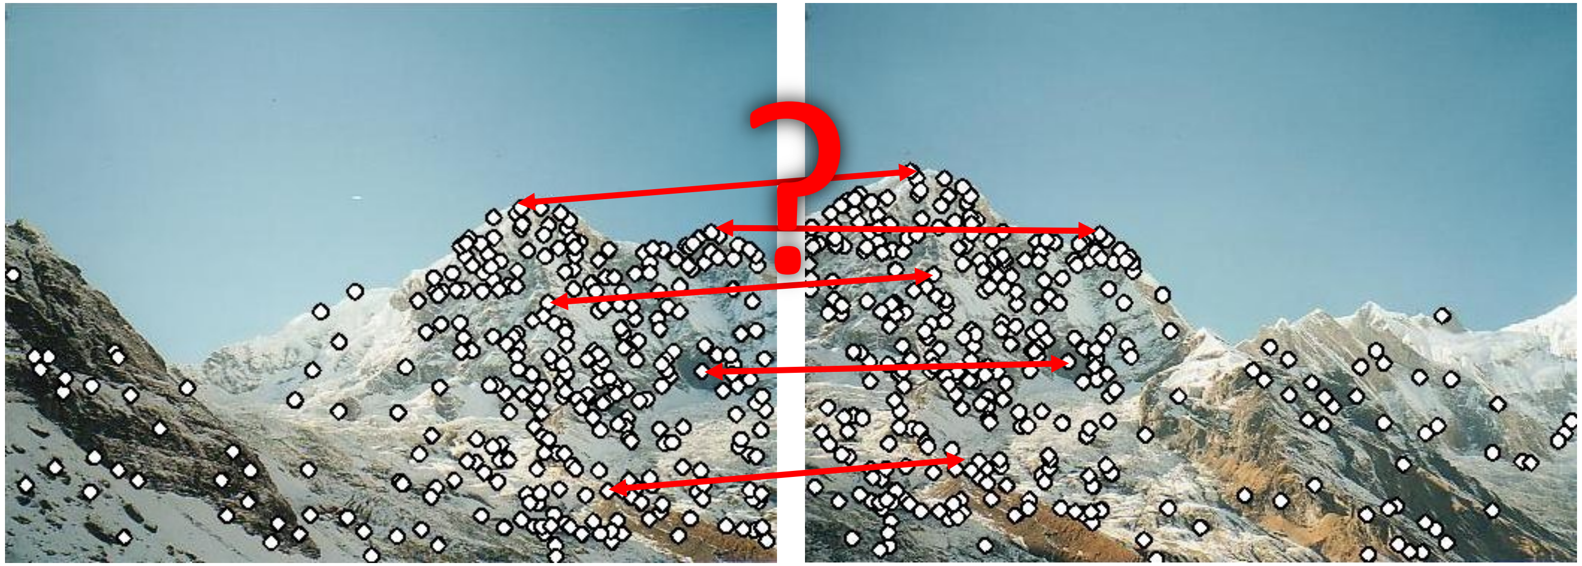

Two interrelated questions:

1. How do we *describe* each feature point?
2. How do we *match* descriptions?

Feature descriptor

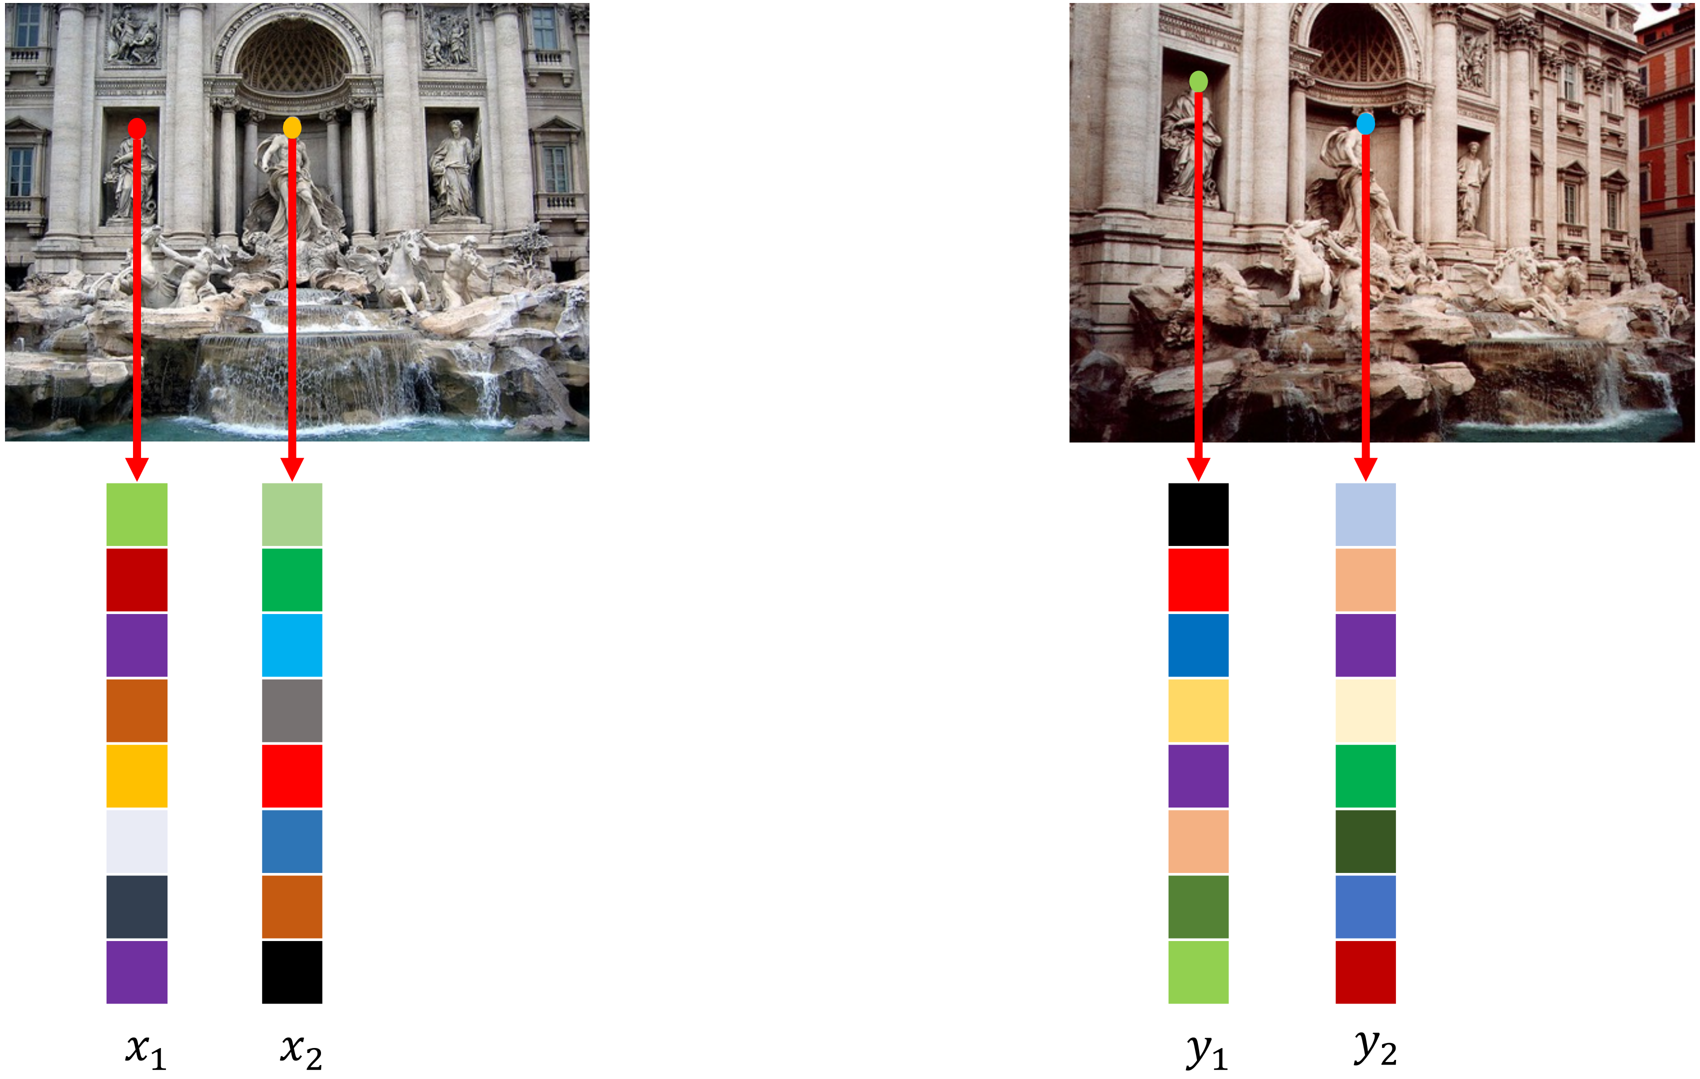

Feature matching

- Measure the distance between (or similarity between) every pair of descriptors

	y_1	y_2
x_1	$d(x_1, y_1)$	$d(x_1, y_2)$
x_2	$d(x_2, y_1)$	$d(x_2, y_2)$

Invariance vs. discriminability

- Invariance:
  - Distance between descriptors of corresponding points should be small even if image is transformed

- Discriminability:
  - Descriptor for a point should be highly unique for each point (far away from other points in the image)

Invariance

- Most feature descriptors are designed to be invariant to
  - Translation, 2D rotation, scale

- They can usually also handle
  - Limited 3D rotations (SIFT works up to about 60 degrees)
  - Limited affine transformations (some are fully affine invariant)
  - Limited illumination/contrast changes

Simple baseline descriptors

Design an invariant feature descriptor

- Simplest descriptor: a single 0
  - What’s this invariant to?
  - Is this discriminative?
- Next simplest descriptor: a single pixel
  - What’s this invariant to?
  - Is this discriminative?

The aperture problem revisited

- If we describe the corner by its color alone then it won't be discriminative
- Will it be invariant?

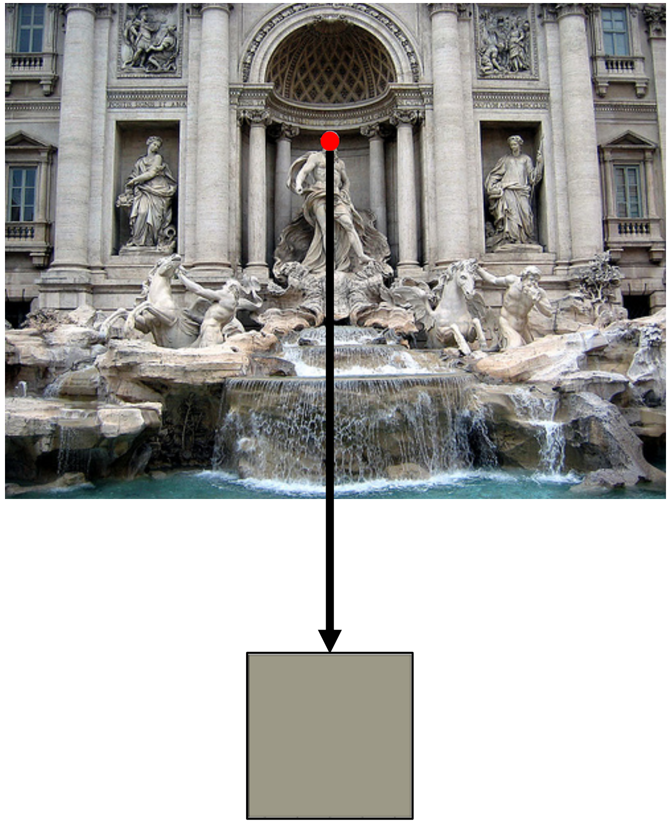

The aperture problem

- We can increase the aperture and look at a whole patch
- How does this affect discriminability? Invariance?

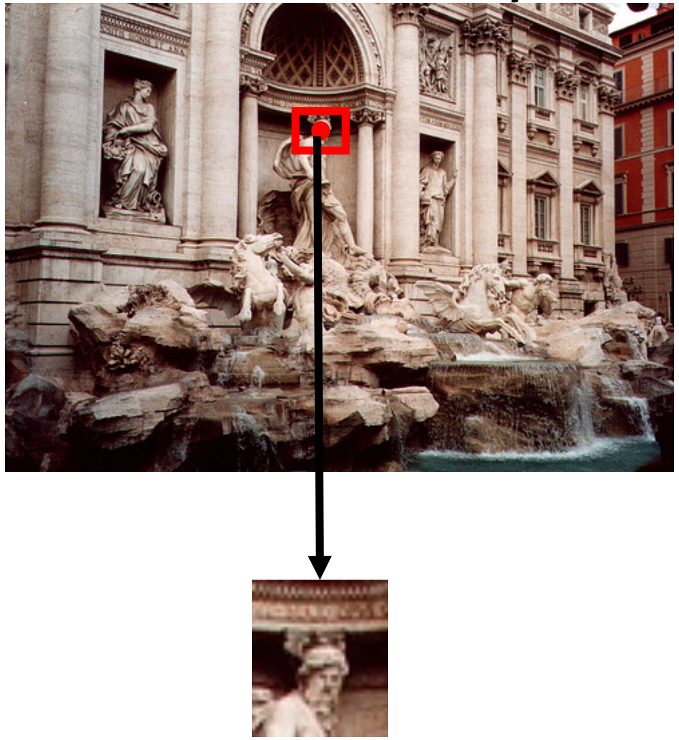

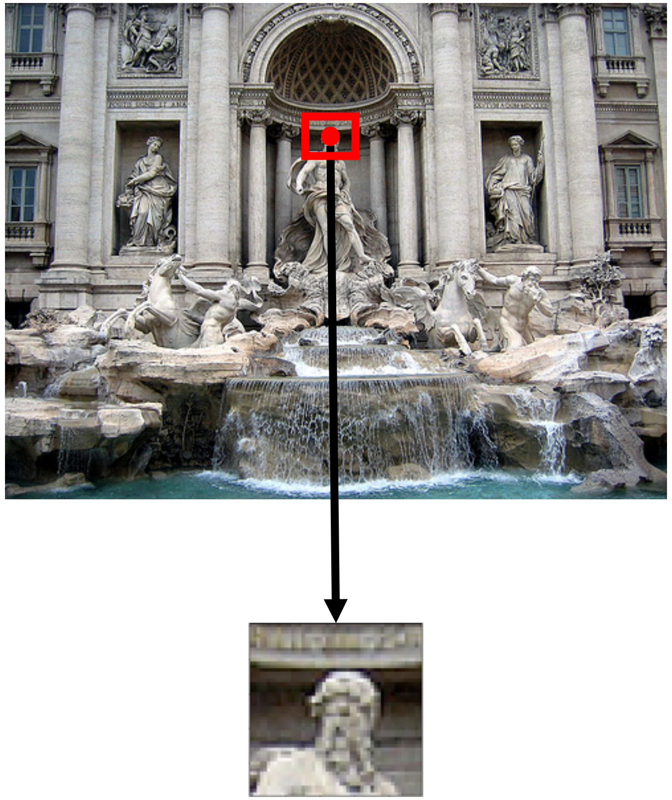

SSD

- Use as descriptor the whole patch
- Match descriptors using euclidean distance
- $d(x, y) = \ ||x \ - y||^2$



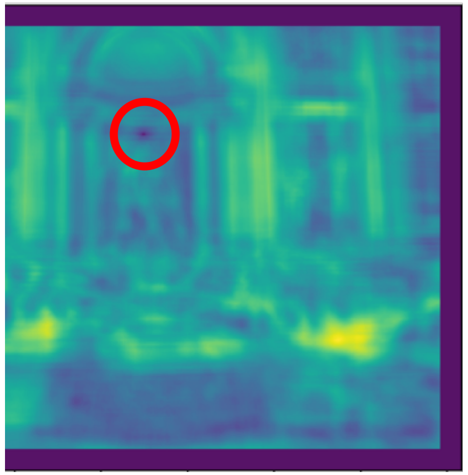

SSD

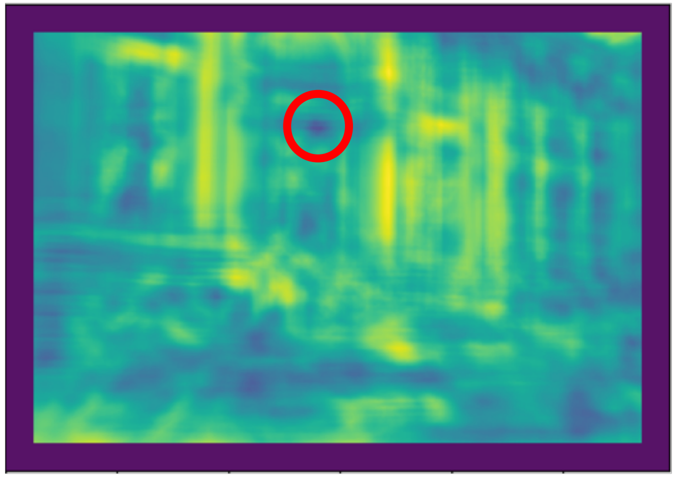

SSD

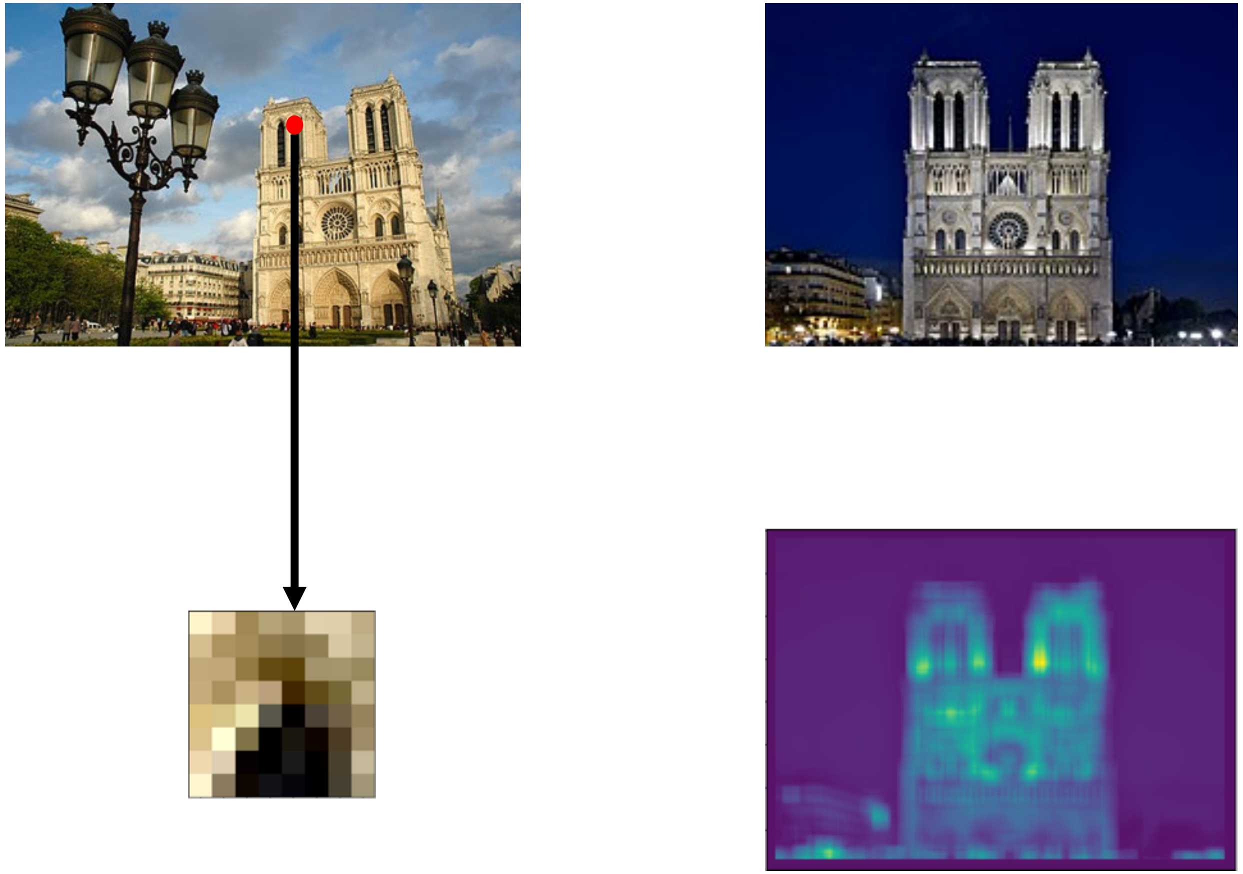

NCC - Normalized Cross Correlation

- Lighting and color change pixel intensities
- Example: increase brightness / contrast
- $I' = \alpha I + \beta$
- Subtract patch mean: invariance to β
- Divide by norm of vector: invariance to α
- $x' = x - < x >$
- $x'' = \frac{x'}{||x'||}$
- *similarity* = $x'' \cdot y''$

NCC - Normalized cross correlation

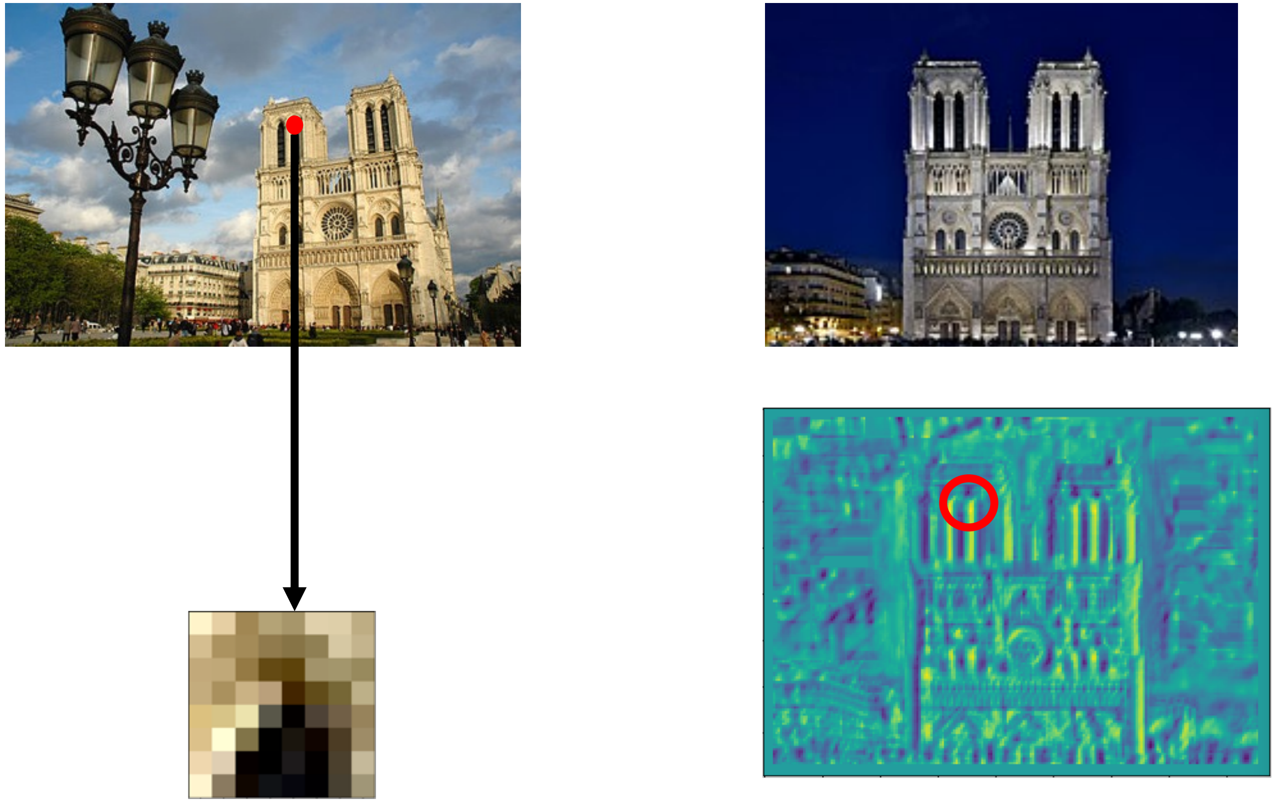

Basic correspondence

- Image patch as descriptor, NCC as similarity
- Invariant to?
  - Photometric transformations?
  - Translation?
  - Rotation?

Rotation invariance for feature descriptors

- Find dominant orientation of the image patch
  - This is given by $\mathbf{x}_{\mathbf{max}}$, the eigenvector of **M** corresponding to λ_{max} (the *larger* eigenvalue)
  - Rotate the patch according to this angle
  - Figure: line represents x_{max}, box represents patch we take as feature descriptor

Figure by Matthew Brown

Multiscale Oriented PatcheS descriptor

Take 40x40 square window around detected feature

- Scale to 1/5 size (using prefiltering)
- Rotate to horizontal
- Sample 8x8 square window centered at feature
- Intensity normalize the window by subtracting the mean, dividing by the standard deviation in the window

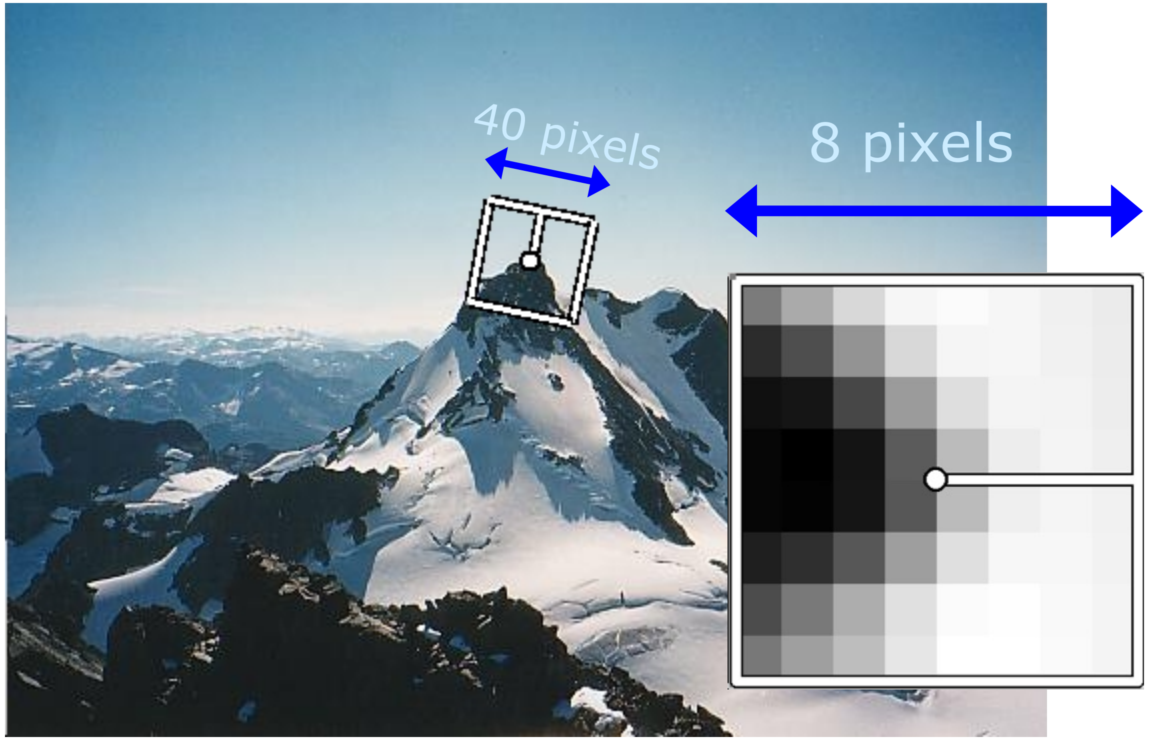

Adapted from slide by Matthew Brown

Scale invariance for feature descriptors

- Recall that corner detector searches over scales for maximum response: record scale which gives maximum response
- Use a patch of the same scale, then resize it to a fixed size

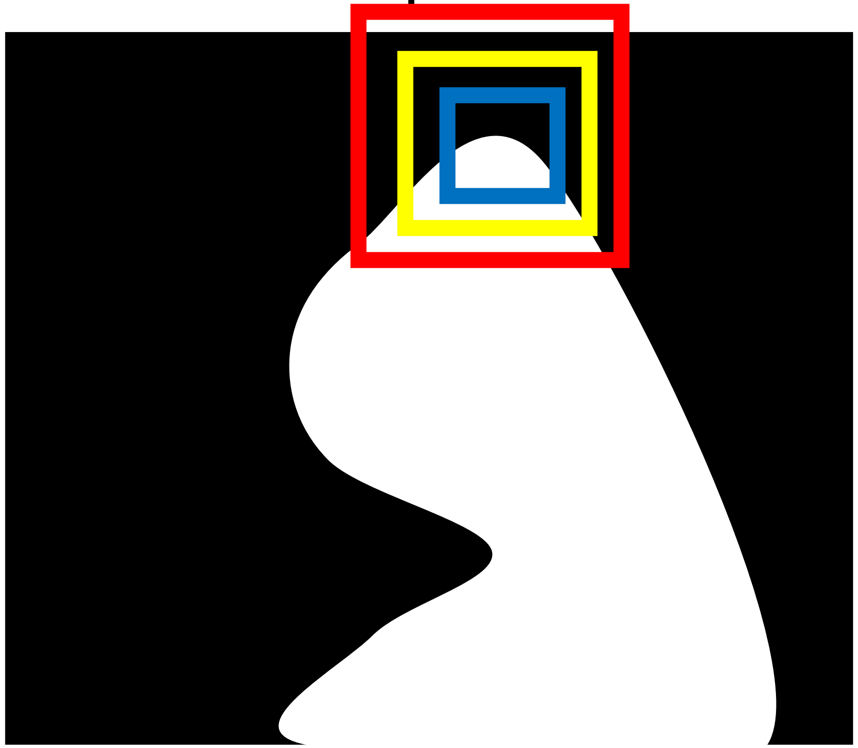

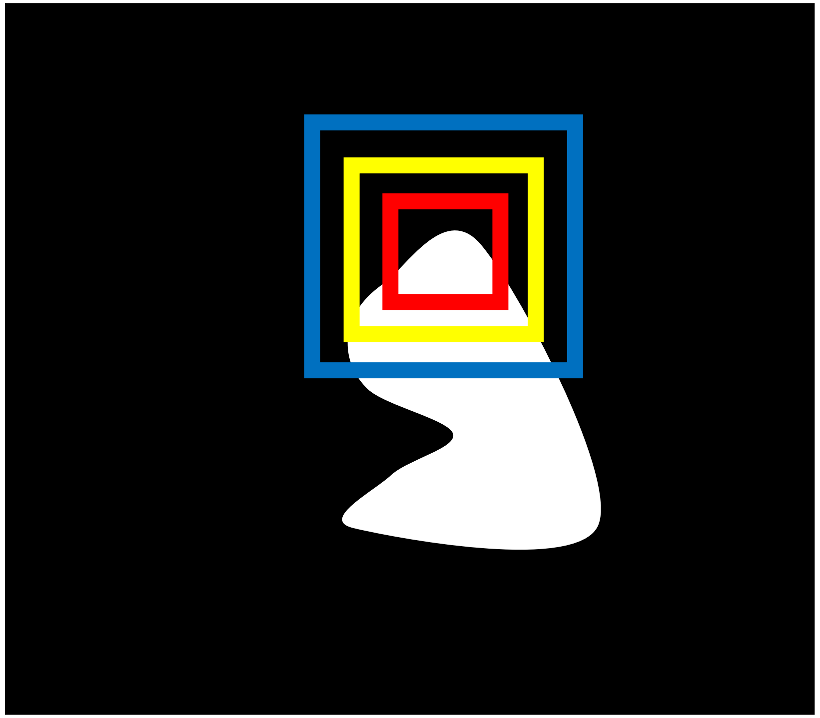

Multiscale Oriented PatcheS descriptor

Take 40x40 square window around detected feature *at appropriate scale*

- Scale to 1/5 size (using prefiltering)
- Rotate to horizontal
- Sample 8x8 square window centered at feature
- Intensity normalize the window by subtracting the mean, dividing by the standard deviation in the window

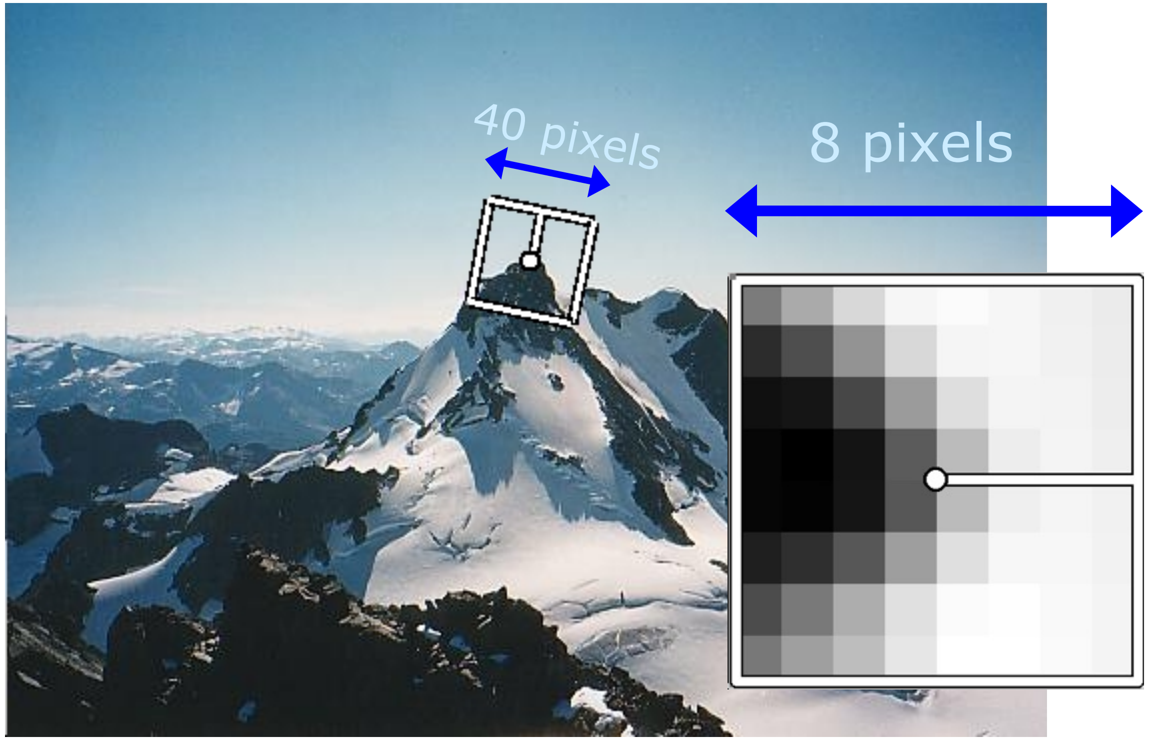

Adapted from slide by Matthew Brown

Detections at multiple scales

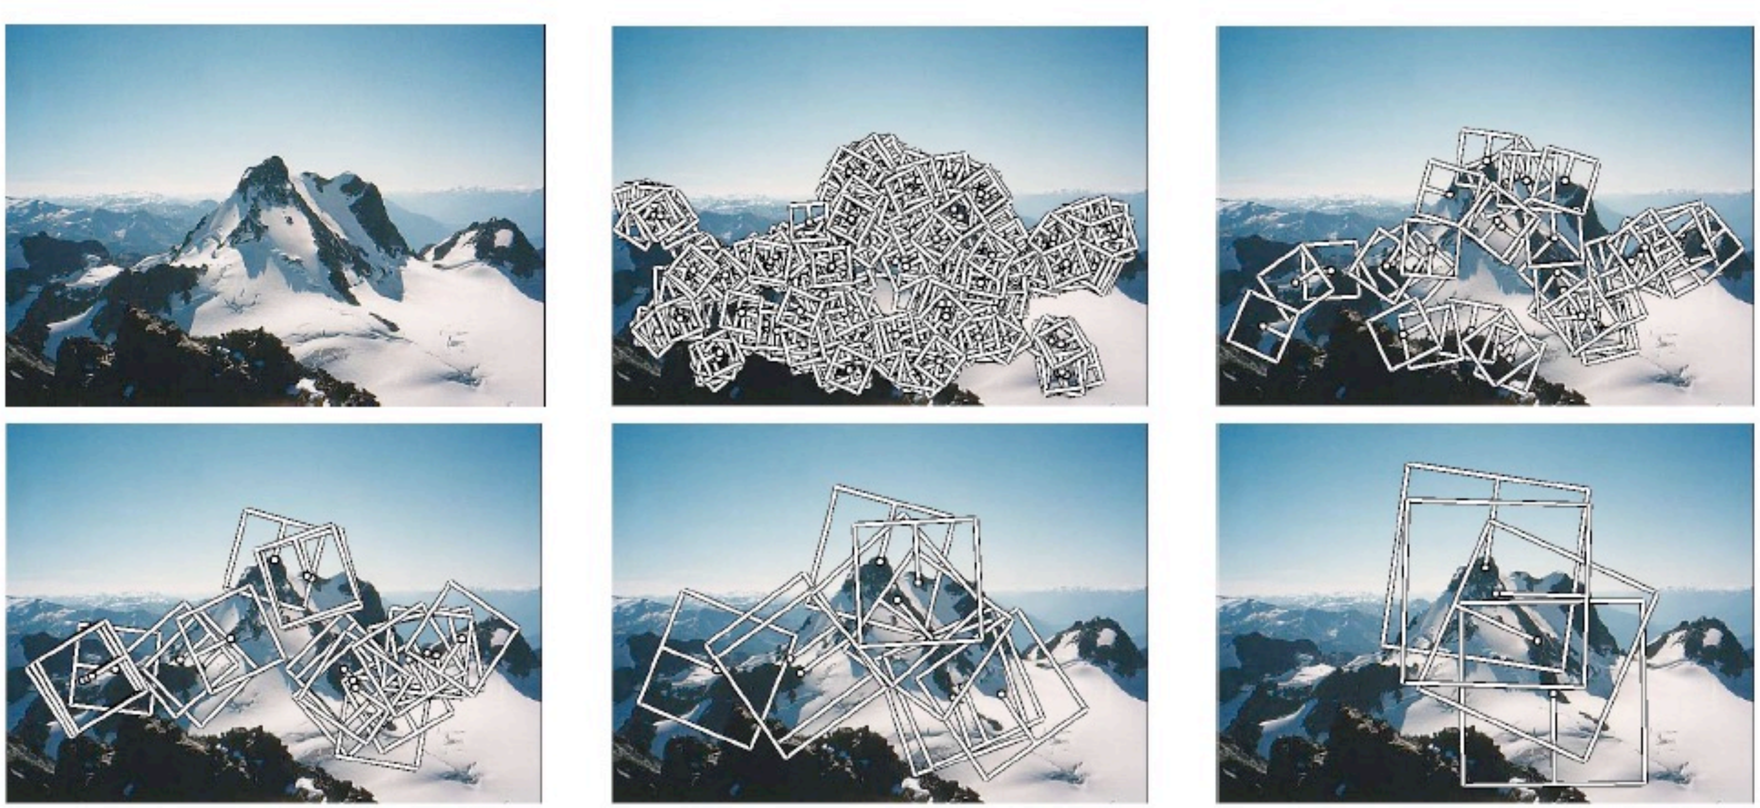

Figure 1. Multi-scale Oriented Patches (MOPS) extracted at five pyramid levels from one of the Matier images. The boxes show the feature orientation and the region from which the descriptor vector is sampled.

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