Computational Photography

Prof. Feng Liu

Spring 2022

http://www.cs.pdx.edu/~fliu/courses/cs510/

04/26/2022

Last Time

- Panorama
  - Feature detection

Today

- Panorama
  - SIFT
  - Feature matching

With slides by C. Dyer, Y. Chuang, R. Szeliski, S. Seitz, M. Brown and V. Hlavac

Harris Detector: Summary

- Average intensity change in direction [*u,v*] can be expressed in bilinear form:

$$E(u,v) \cong [u,v] \; M \begin{bmatrix} u \\ v \end{bmatrix}$$

- Describe a point in terms of eigenvalues of **M**: *measure of corner response:*

$$R = \lambda_1 \lambda_2 - k\left(\lambda_1 + \lambda_2\right)^2$$

- A good (corner) point should have a *large intensity change* in *all directions*, i.e., *R* should be a large positive value

Credit: C. Dyer

Harris Detector: Example

Credit: C. Dyer

Harris Detector: Example

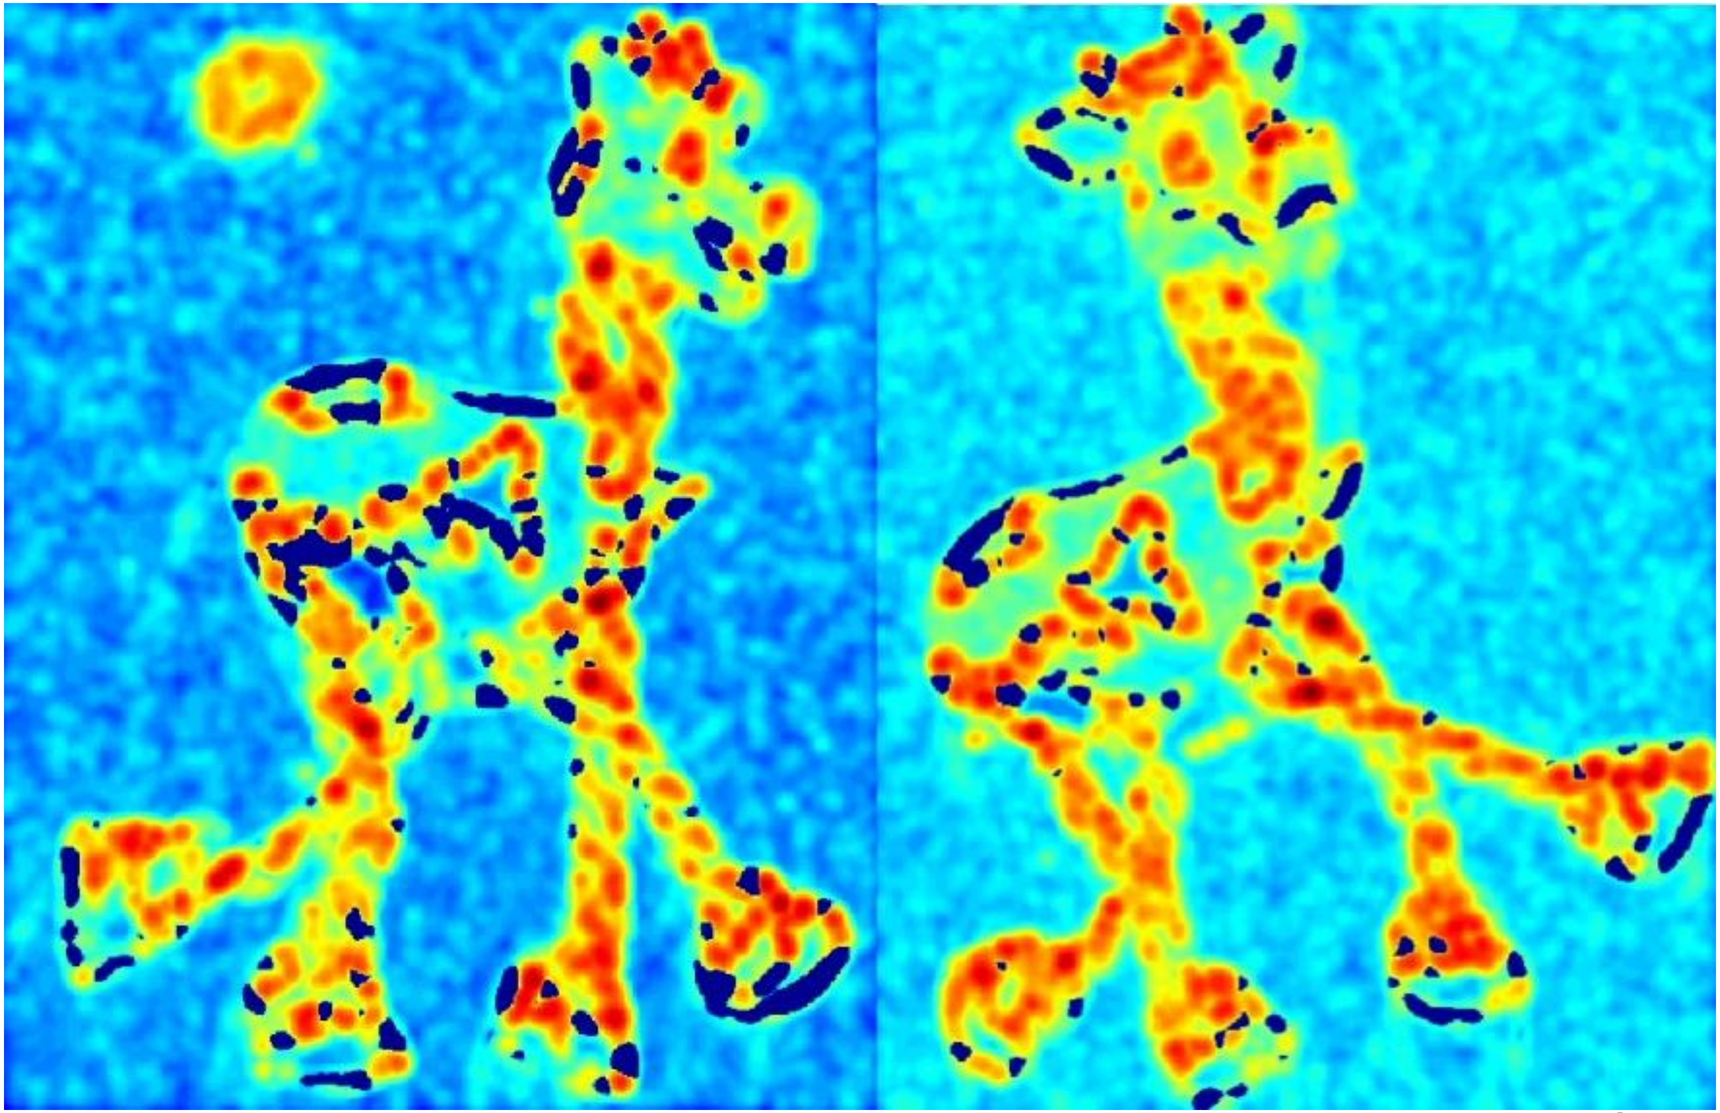

Credit: C. Dyer Compute corner response $R = \lambda_1\lambda_2 - k(\lambda_1 + \lambda_2)^2$

Harris Detector: Example

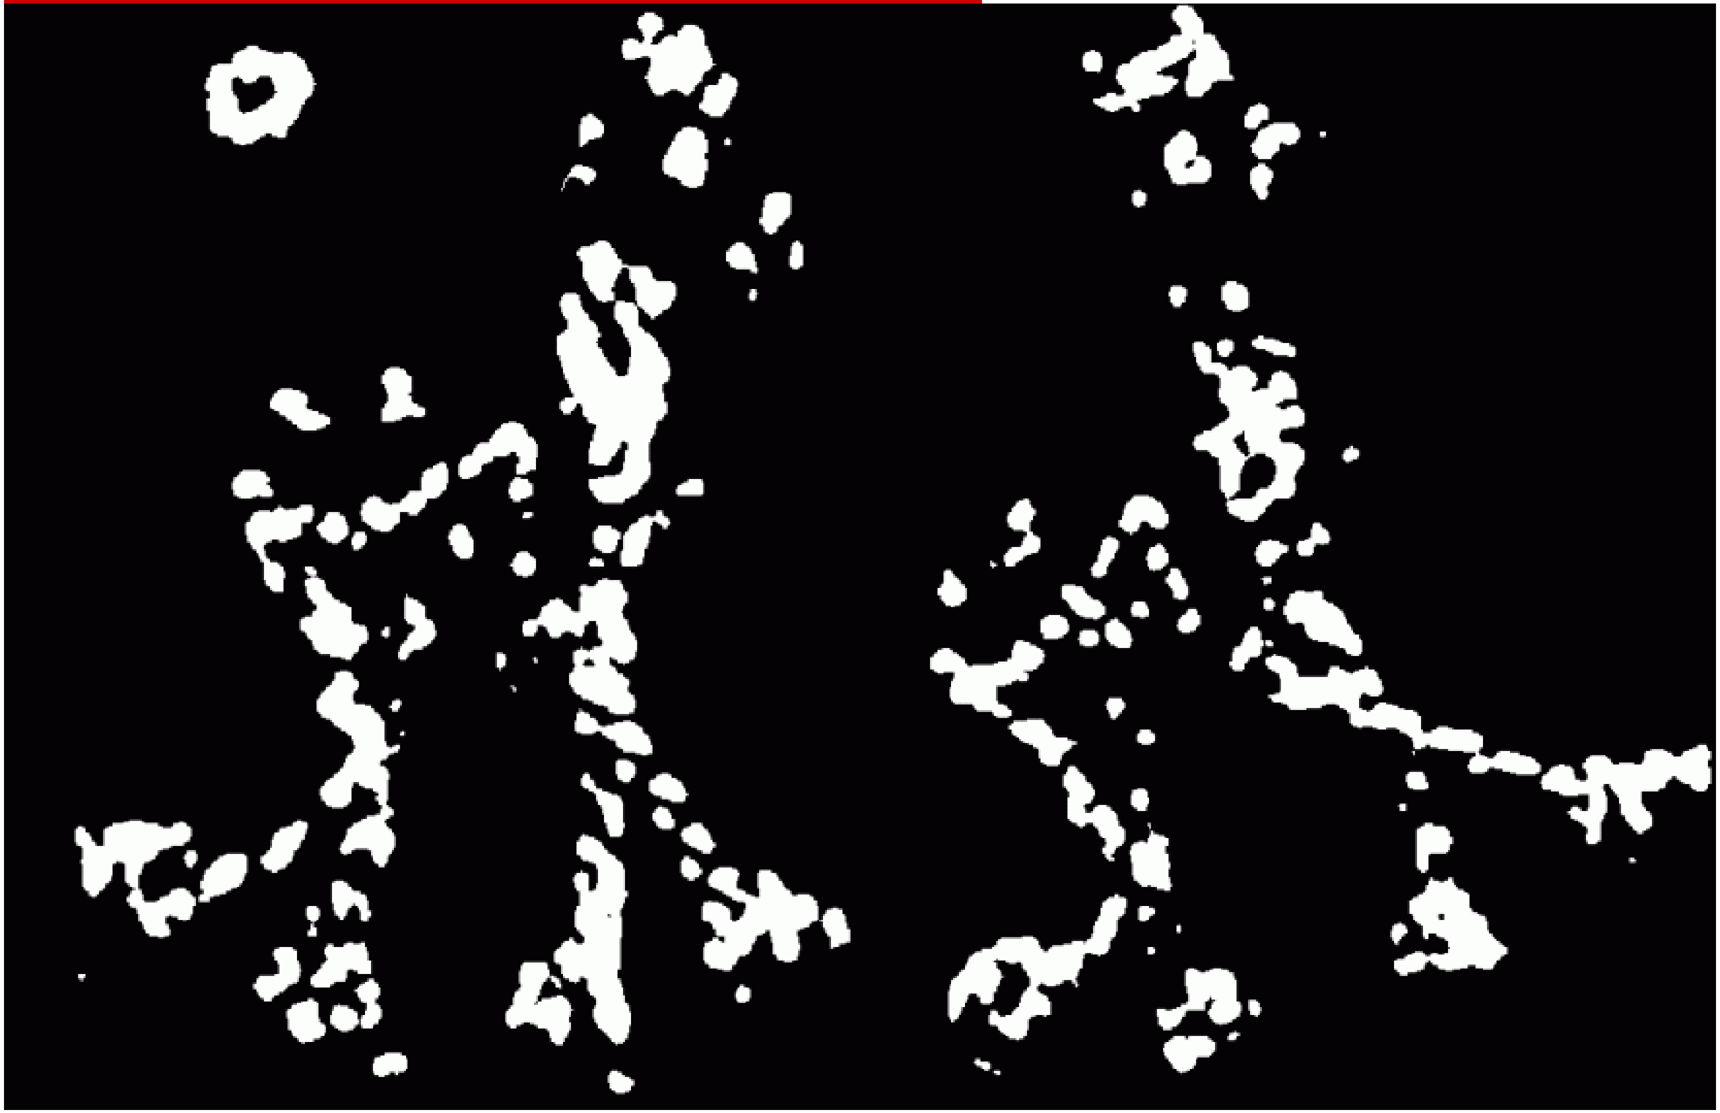

Credit: C. Dyer

Find points with large corner response: $R >$ threshold

Harris Detector: Example

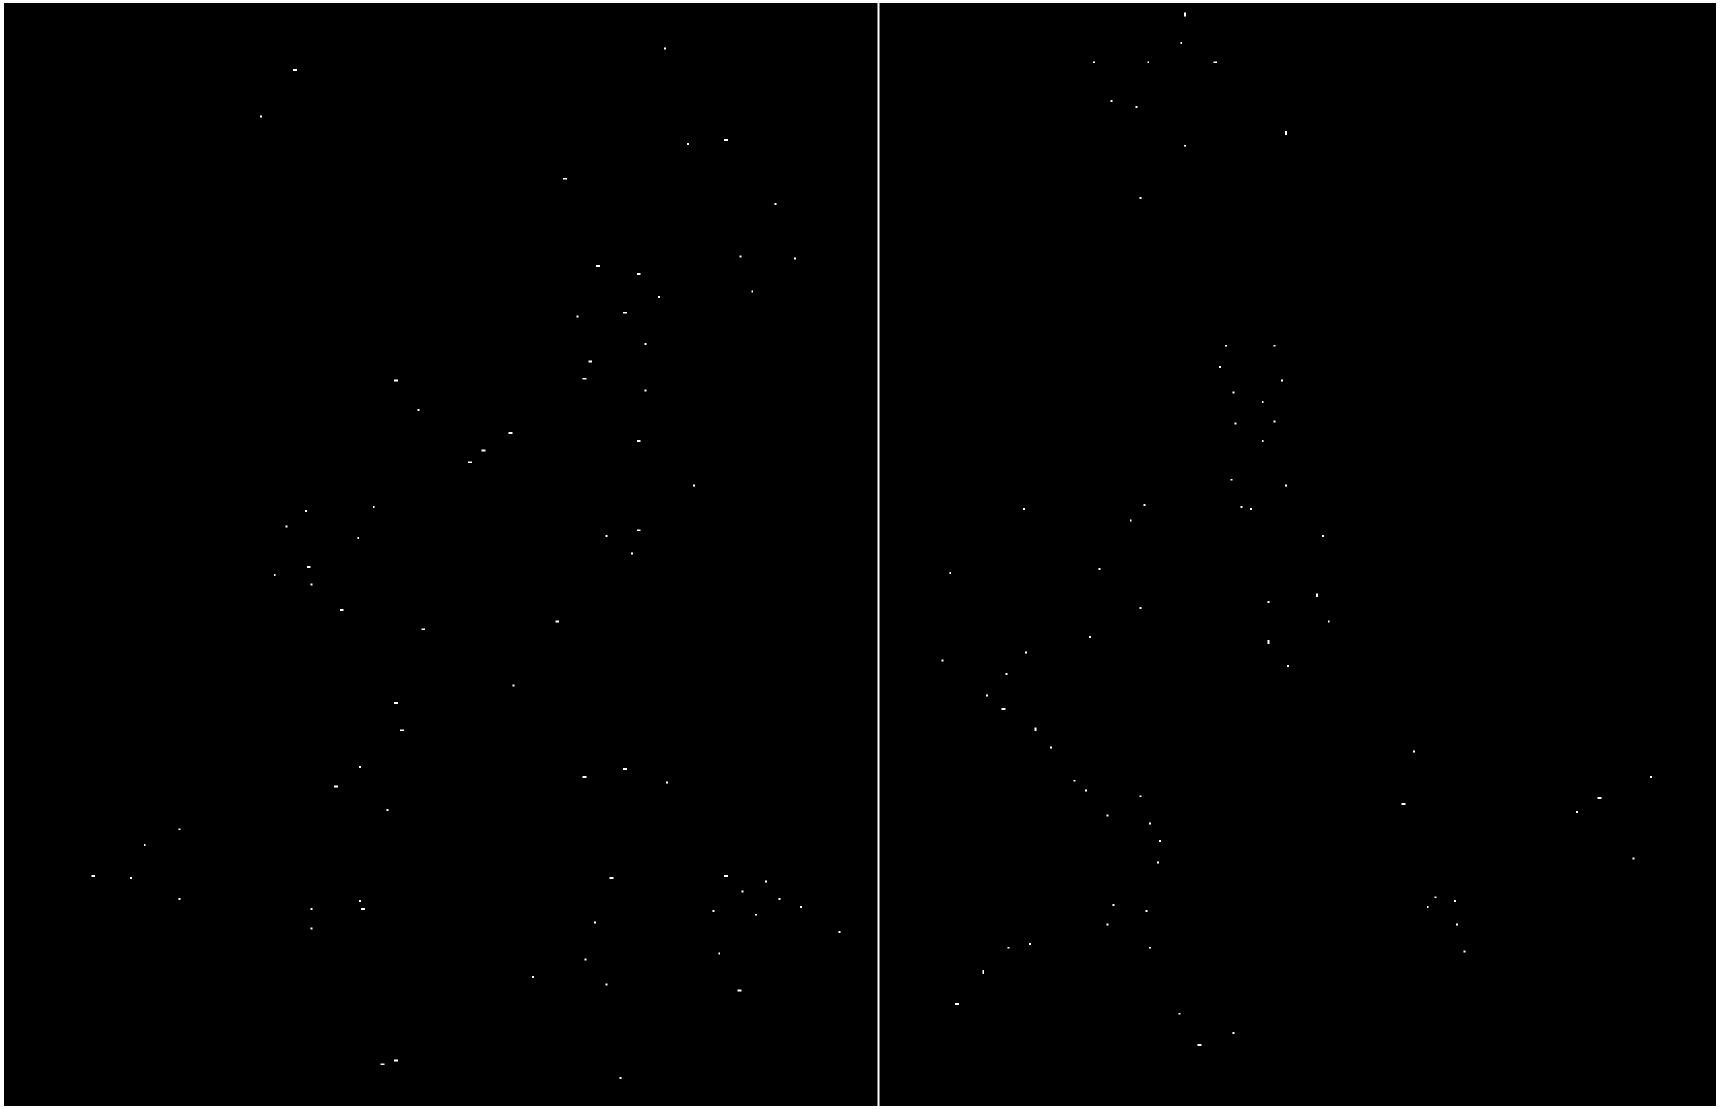

Credit: C. Dyer Take only the points of local maxima of R

Harris Detector: Example

Credit: C. Dyer

Harris Detector Properties

- Rotation invariance

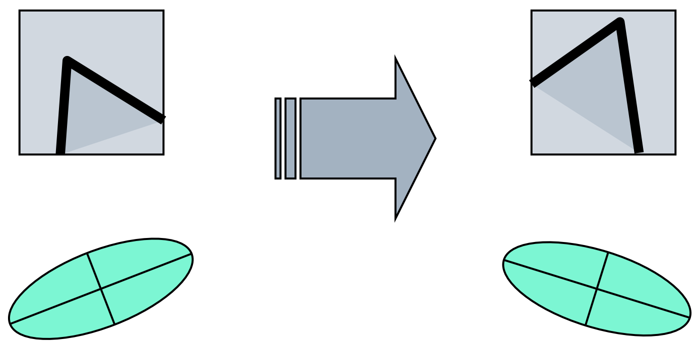

Ellipse rotates but its shape (i.e., eigenvalues) remains the same

Corner response R is invariant to image rotation

Credit: C. Dyer

Harris Detector Properties

- But **not** invariant to *image scale*

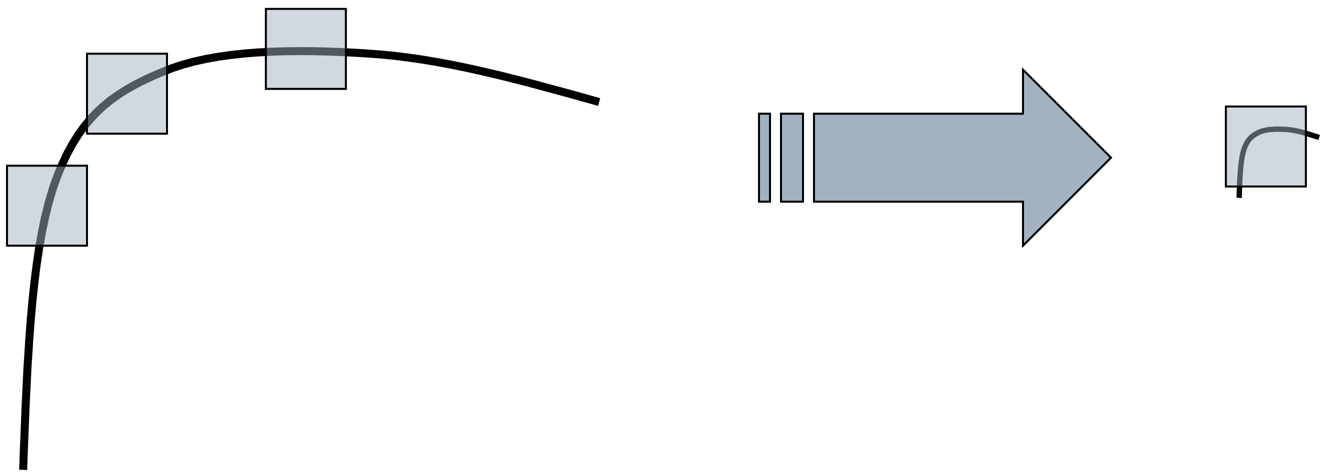

Fine scale: All points will be classified as edges

Coarse scale: Corner

Credit: C. Dyer

Harris Detector Properties

- Quality of Harris detector for different scale changes

Repeatability rate:

$$\frac{\text{\# correct correspondences}}{\text{\# possible correspondences}}$$

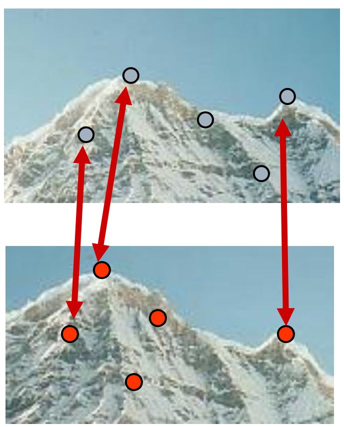

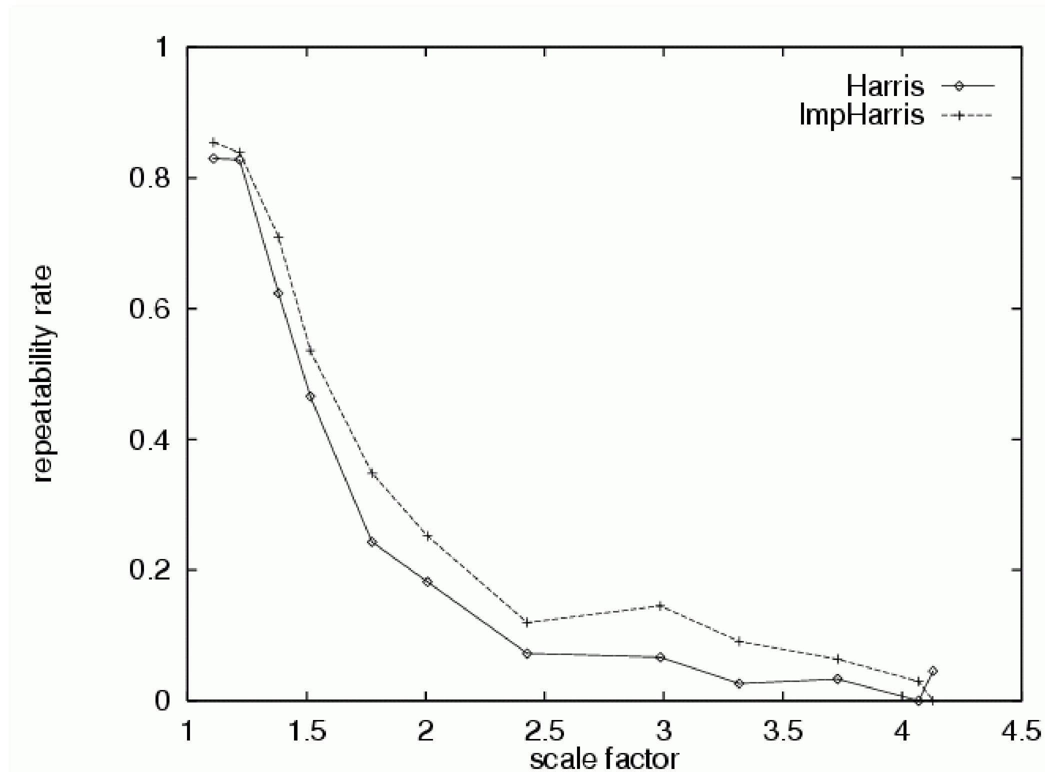

C. Schmid et al., "Evaluation of Interest Point Detectors," IJCV 2000

Credit: C. Dyer

Invariant Local Features

- Goal: Detect *the same* interest points regardless of *image changes* due to translation, rotation, scale, viewpoint

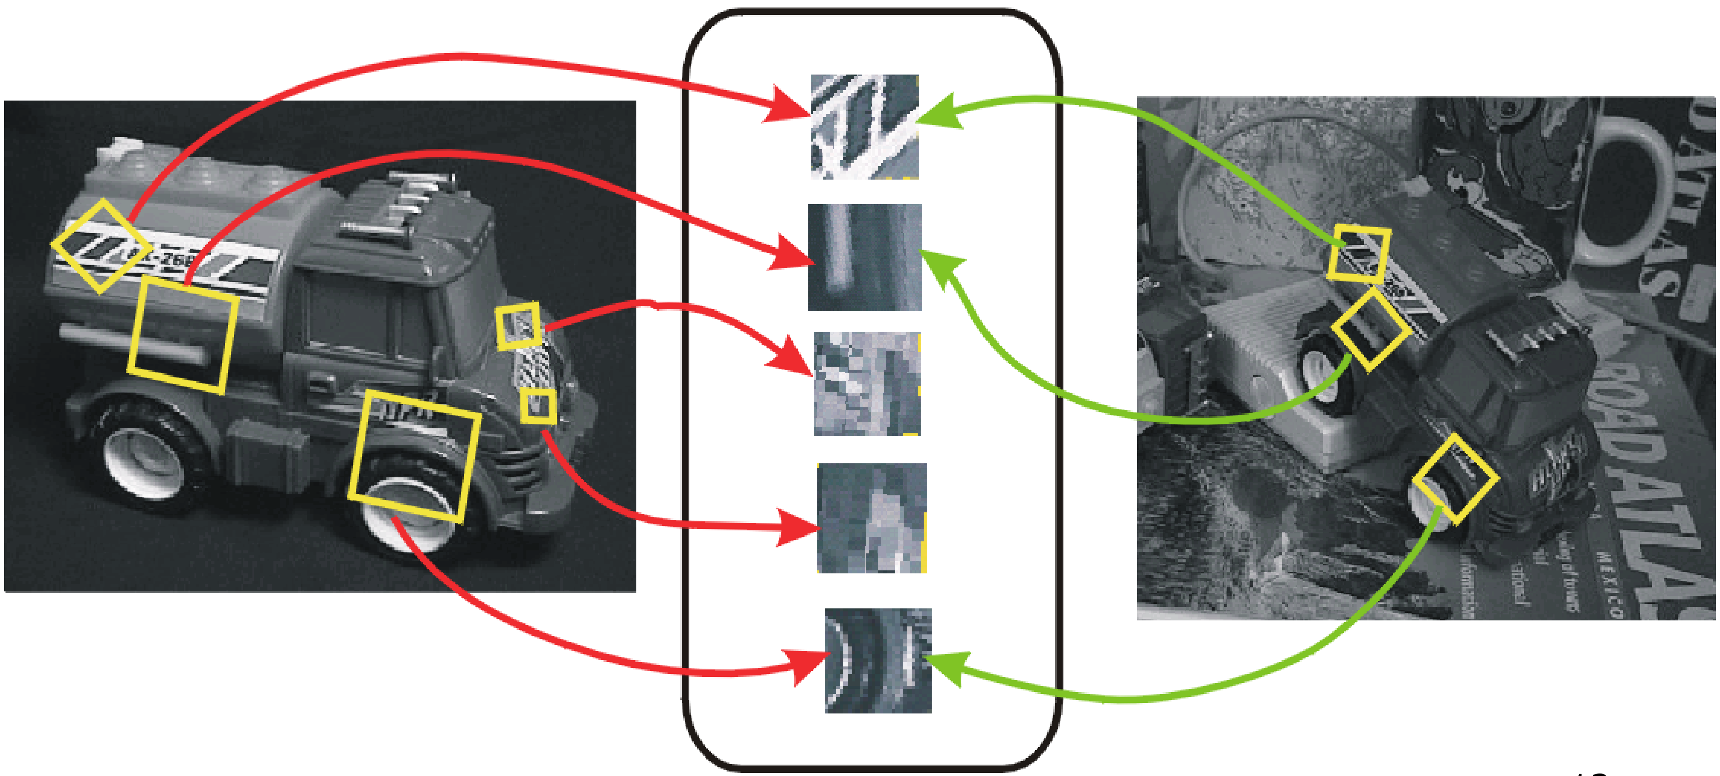

Models of Image Change

- Geometry
  - Rotation

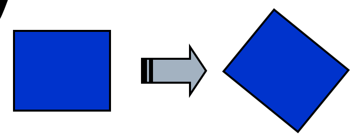

  - Similarity (rotation + uniform scale)

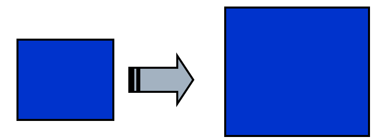

  - Affine (scale dependent on direction)
    valid for: orthographic camera, locally planar object

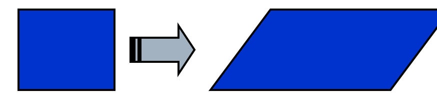

- Photometry
  - Affine intensity change ($I \rightarrow a\,I + b$)

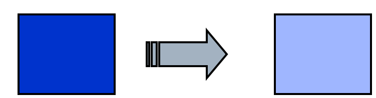

SIFT Detector [Lowe ’04]

- Difference-of-Gaussian (DoG) is an approximation of the Laplacian-of-Gaussian (LoG)

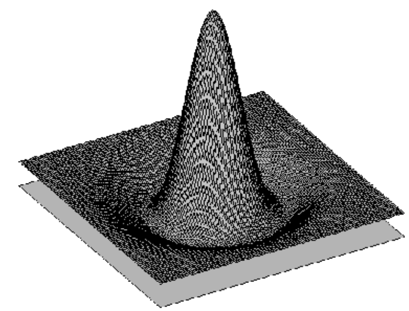

Credit: C. Dyer

Lowe, D. G., “Distinctive Image Features from Scale-Invariant Keypoints”, International Journal of Computer Vision, 60, 2, pp. 91-110, 2004

DoG Pyramid

Credit: C. Dyer

SIFT Detector

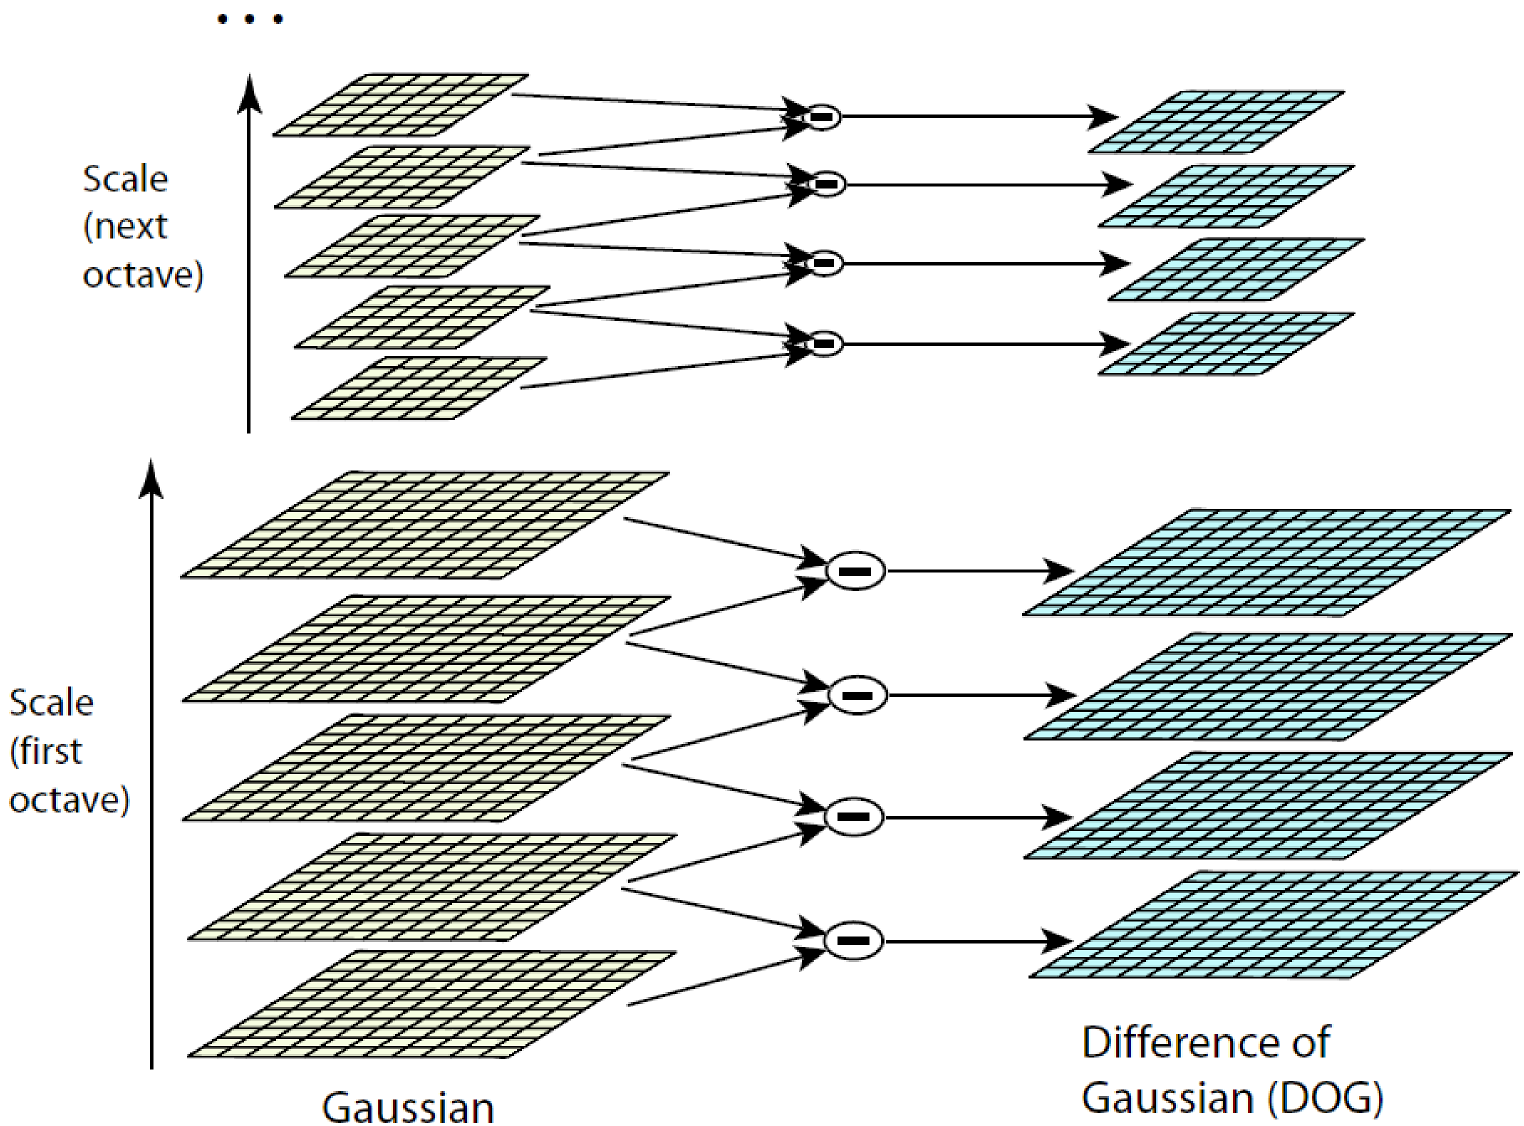

Credit: C. Dyer

SIFT Detector Algorithm Summary

- Detect local maxima in position and scale of squared values of difference-of-Gaussian
- Fit a quadratic to surrounding values for sub-pixel and sub-scale interpolation
- Output = list of (x, y, σ) points

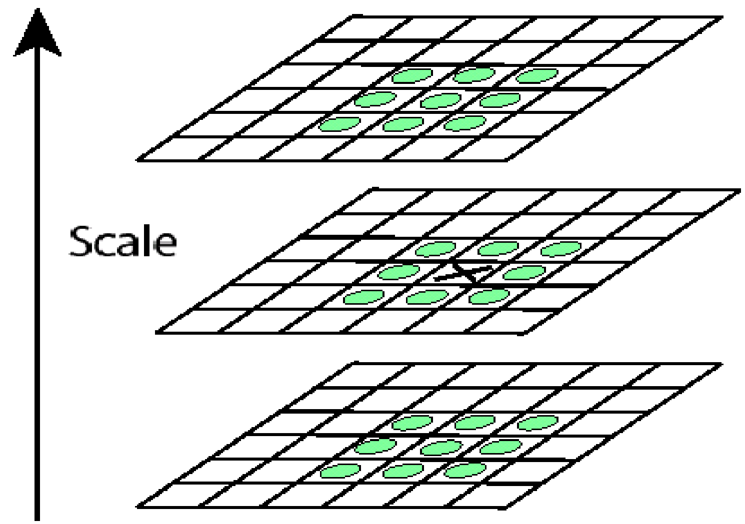

Credit: C. Dyer

References on Feature Descriptors

- A performance evaluation of local descriptors, K. Mikolajczyk and C. Schmid, *IEEE Trans. PAMI* **27**(10), 2005

- Evaluation of features detectors and descriptors based on 3D objects, P. Moreels and P. Perona, *Int. J. Computer Vision* **73**(3), 2007

Credit: C. Dyer

Stitching Recipe

- Align pairs of images
  - Feature Detection
  - Feature Matching
  - Homography Estimation

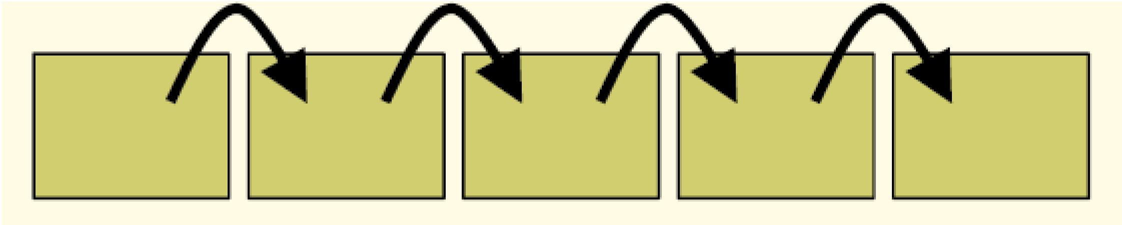

- Align all to a common frame
- Adjust (Global) & Blend

Invariant Local Features

- Goal: Detect *the same* interest points regardless of *image changes* due to translation, rotation, scale, viewpoint

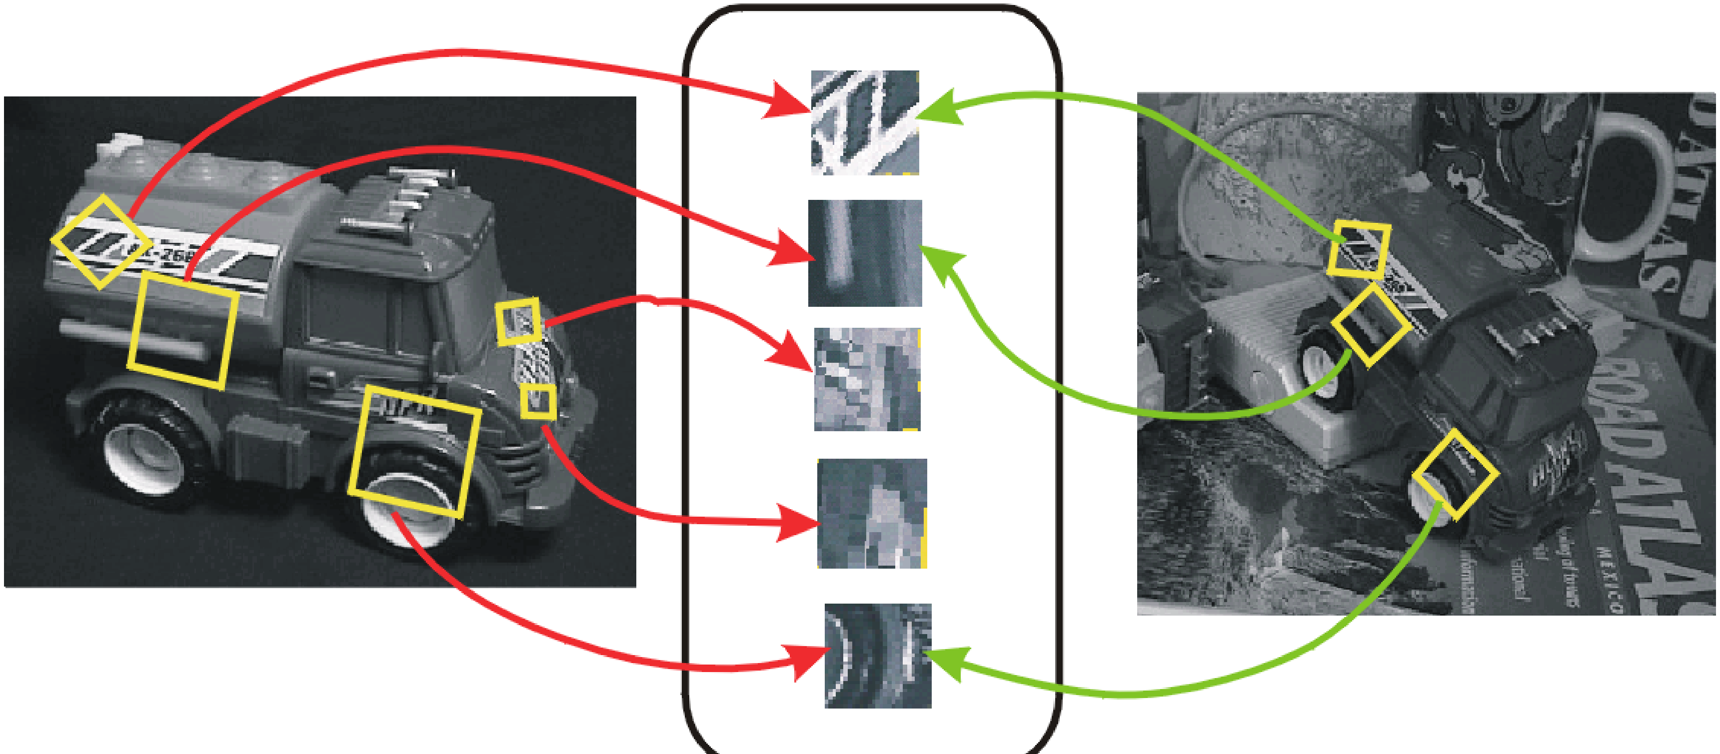

Local Features: Description

1. Detection: Identify the interest points

2. Description: Extract feature vector for each interest point

3. Matching: Determine correspondence between descriptors in two views

$$\mathbf{x}_1 = [x_1^{(1)}, \ldots, x_d^{(1)}]$$

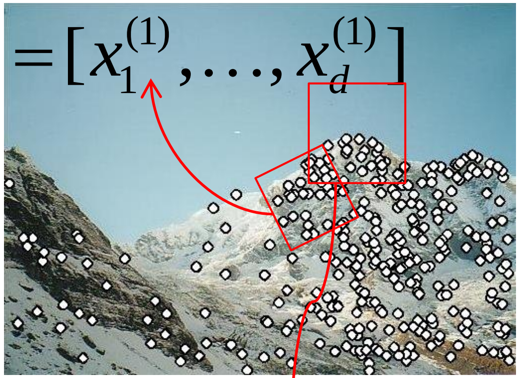

$$\mathbf{x}_2 = [x_1^{(2)}, \ldots, x_d^{(2)}]$$

Feature Point Descriptors

- After detecting points (and patches) in each image,
- Next question: **How to match them?**

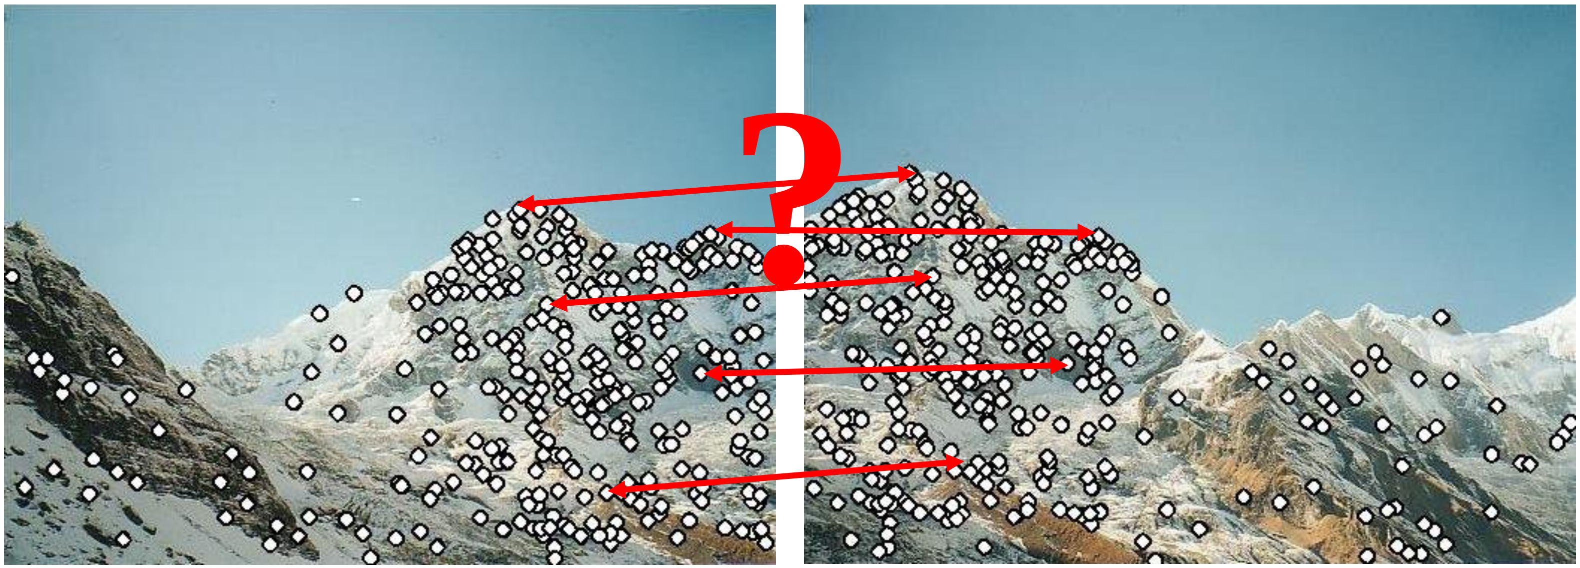

Point descriptor should be:
1. Invariant
2. Distinctive

All the following slides are used from Prof. C. Dyer's relevant course, except those with explicit acknowledgement.

Geometric Transformations

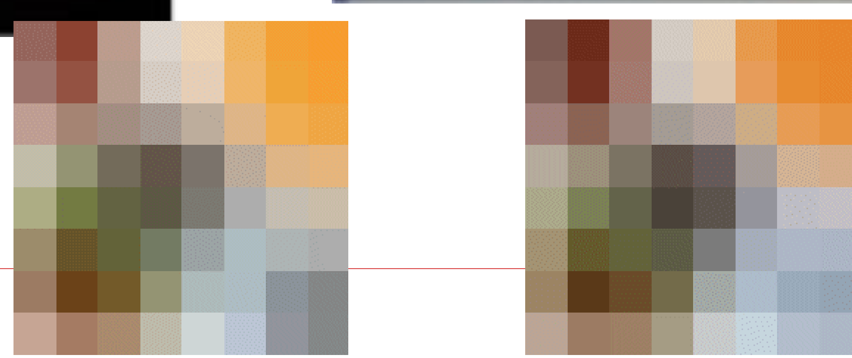

e.g. scale, translation, rotation

Photometric Transformations

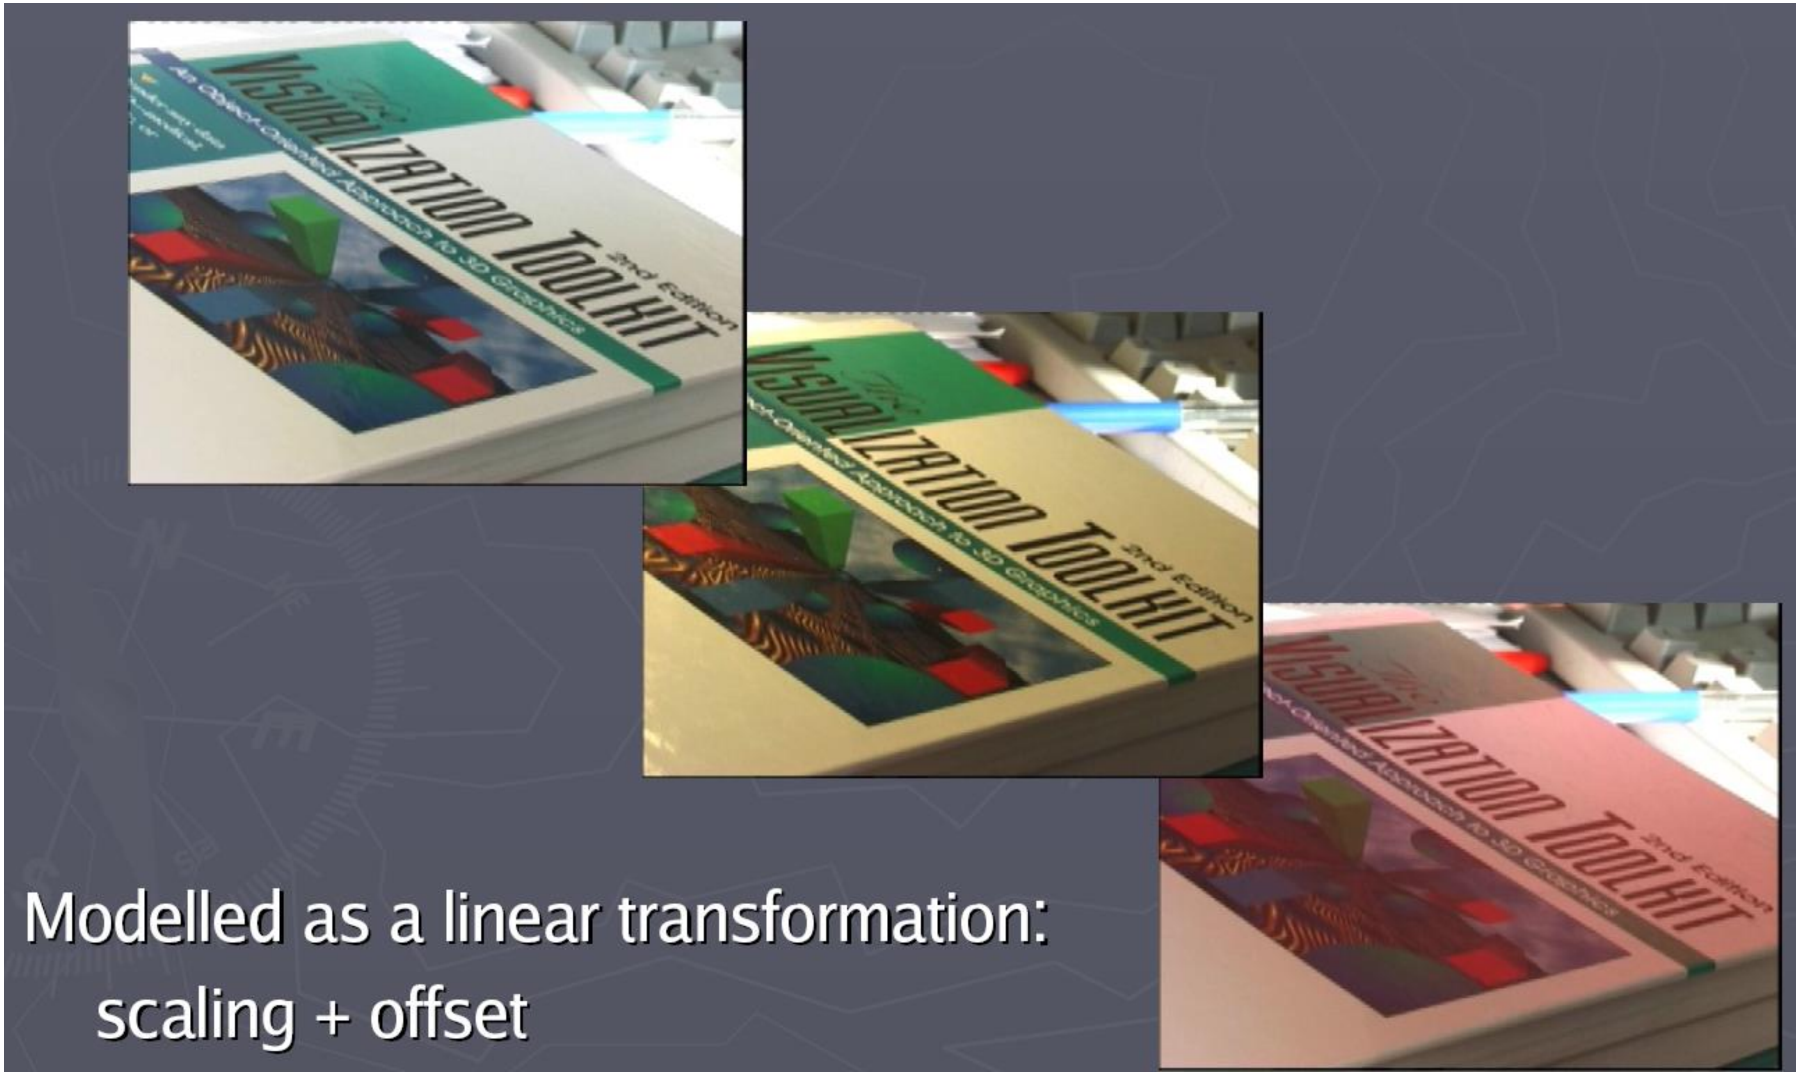

Figure from T. Tuytelaars ECCV 2006 tutorial

Raw Patches as Local Descriptors

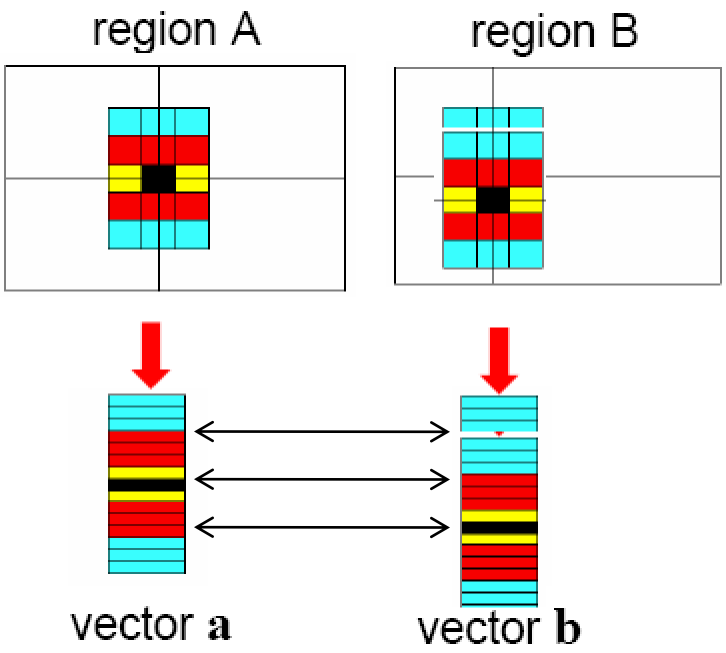

The simplest way to describe the neighborhood around an interest point is to write down the list of intensities to form a feature vector

But this is very sensitive to even small shifts or rotations

Making Descriptors Invariant to Rotation

- Find local orientation

Dominant direction of gradient:

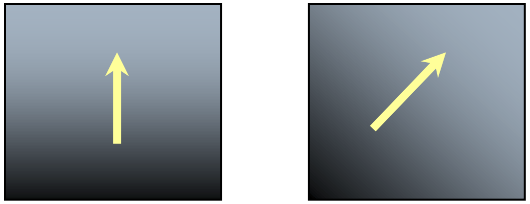

- Compute description **relative** to this orientation

[1] K.Mikolajczyk, C.Schmid. "Indexing Based on Scale Invariant Interest Points". ICCV 2001
[2] D.Lowe. "Distinctive Image Features from Scale-Invariant Keypoints". Accepted to IJCV 2004

Derivatives with convolution

For 2D function $f(x, y)$, the partial derivative is:

$$\frac{\partial f(x, y)}{\partial x} = \lim_{\varepsilon \to 0} \frac{f(x+\varepsilon, y) - f(x, y)}{\varepsilon}$$

For discrete data, we can approximate using finite differences:

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

To implement above as convolution, what would be the associated filter?

Partial derivatives of an image

$$\frac{\partial f(x,y)}{\partial x}$$

-1	1

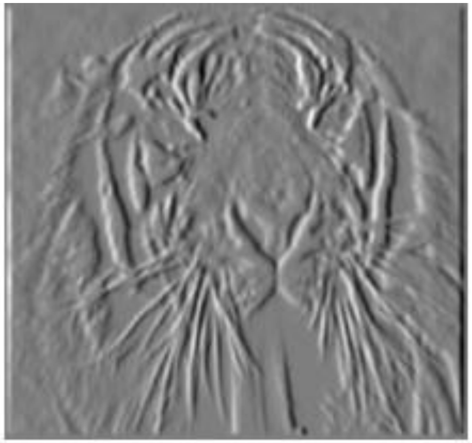

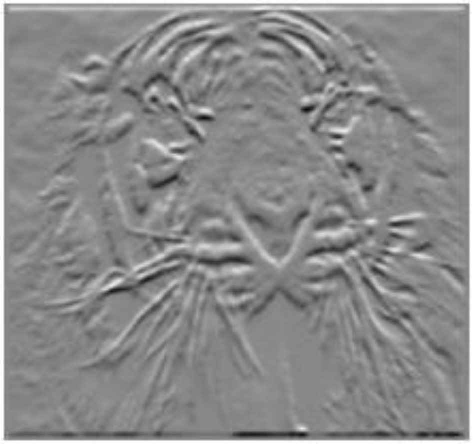

$$\frac{\partial f(x,y)}{\partial y}$$

-1
1

or

1
-1

Which shows changes with respect to x?

Source: S. Lazebnik

Finite difference filters

Other approximations of derivative filters exist:

Prewitt: $M_x = \begin{bmatrix} -1 & 0 & 1 \\ -1 & 0 & 1 \\ -1 & 0 & 1 \end{bmatrix}$; $M_y = \begin{bmatrix} 1 & 1 & 1 \\ 0 & 0 & 0 \\ -1 & -1 & -1 \end{bmatrix}$

Sobel: $M_x = \begin{bmatrix} -1 & 0 & 1 \\ -2 & 0 & 2 \\ -1 & 0 & 1 \end{bmatrix}$; $M_y = \begin{bmatrix} 1 & 2 & 1 \\ 0 & 0 & 0 \\ -1 & -2 & -1 \end{bmatrix}$

Roberts: $M_x = \begin{bmatrix} 0 & 1 \\ -1 & 0 \end{bmatrix}$; $M_y = \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix}$

Source: K. Grauman

Image gradient

The gradient of an image: $$\nabla f = \left[\frac{\partial f}{\partial x}, \frac{\partial f}{\partial y}\right]$$

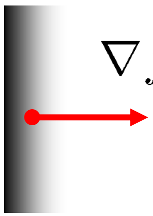

$$\nabla f = \left[\frac{\partial f}{\partial x}, 0\right]$$

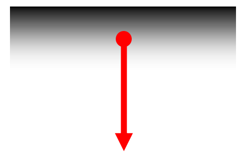

$$\nabla f = \left[0, \frac{\partial f}{\partial y}\right]$$

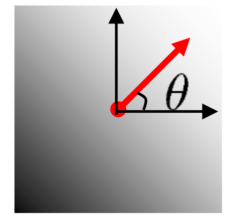

$$\nabla f = \left[\frac{\partial f}{\partial x}, \frac{\partial f}{\partial y}\right]$$

The gradient points in the direction of most rapid increase in intensity

- How does this direction relate to the direction of the edge?

The gradient direction is given by $\theta = \tan^{-1}\left(\frac{\partial f}{\partial y} / \frac{\partial f}{\partial x}\right)$

The *edge strength* is given by the gradient magnitude

$$\|\nabla f\| = \sqrt{\left(\frac{\partial f}{\partial x}\right)^2 + \left(\frac{\partial f}{\partial y}\right)^2}$$

Source: Steve Seitz

SIFT Descriptor: Select Major Orientation

- Compute **histogram of local gradient directions** computed at selected scale in neighborhood of a feature point relative to dominant local orientation
- Compute gradients within sub-patches, and compute histogram of orientations using discrete “bins”
- Descriptor is rotation and scale invariant, and also has some illumination invariance (why?)

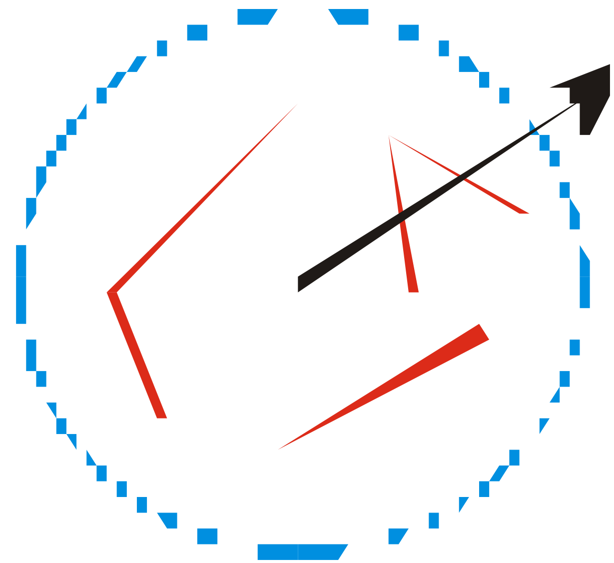

SIFT Descriptor

- Compute gradient orientation histograms on 4 x 4 neighborhoods over 16 x 16 array of locations in scale space around each keypoint position, relative to the keypoint orientation using thresholded image gradients from Gaussian pyramid level at keypoint’s scale
- Quantize orientations to 8 values
- 4 x 4 array of histograms
- SIFT feature vector of length 4 x 4 x 8 = 128 values for each keypoint
- Normalize the descriptor to make it invariant to intensity change

D.Lowe. “Distinctive Image Features from Scale-Invariant Keypoints,” IJCV 2004

SIFT Descriptor

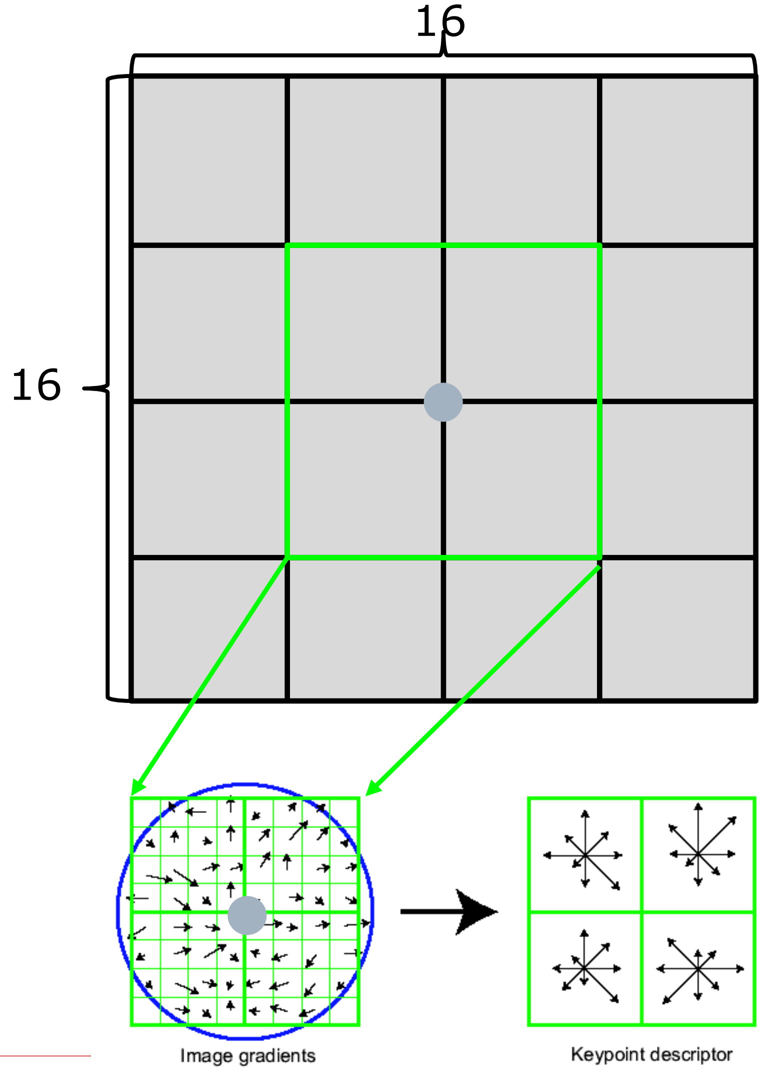

D.Lowe. “Distinctive Image Features from Scale-Invariant Keypoints,” IJCV 2004

Feature Detection and Description Summary

- Stable (repeatable) feature points can currently be detected that are invariant to
  - Rotation, scale, and affine transformations, but ***not*** to more general perspective and projective transformations

- Feature point descriptors can be computed, ***but***
  - are *noisy* due to use of differential operators
  - are *not* invariant to projective transformations

Feature Matching

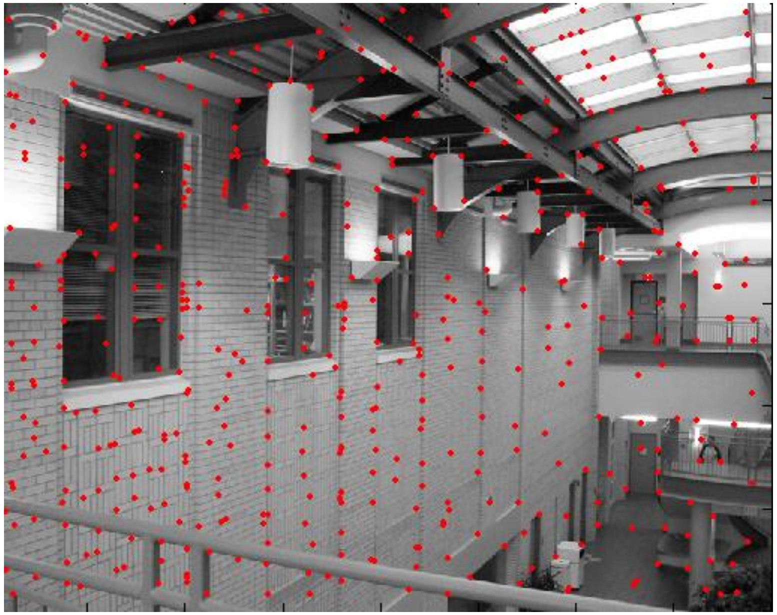

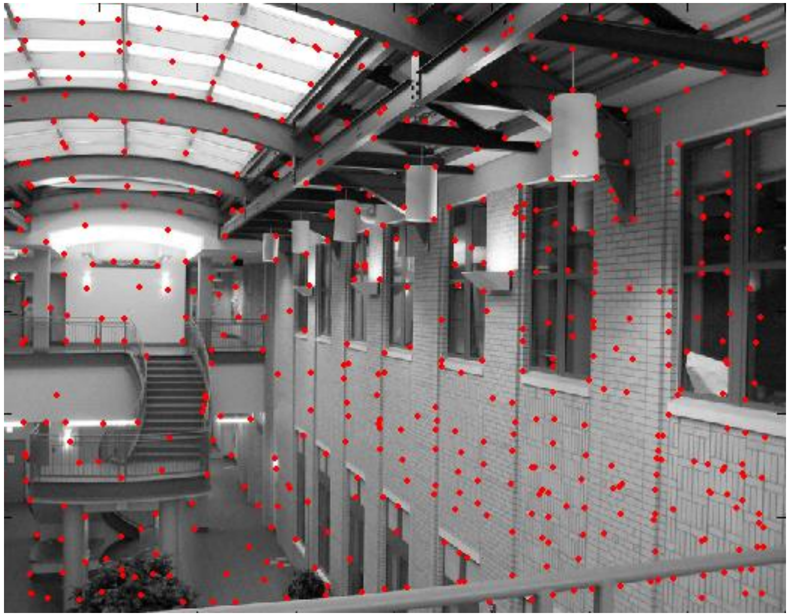

Wide-Baseline Feature Matching

- Standard approach for pair-wise matching:
  - For each feature point in image A
  - Find the feature point with the closest descriptor in image B

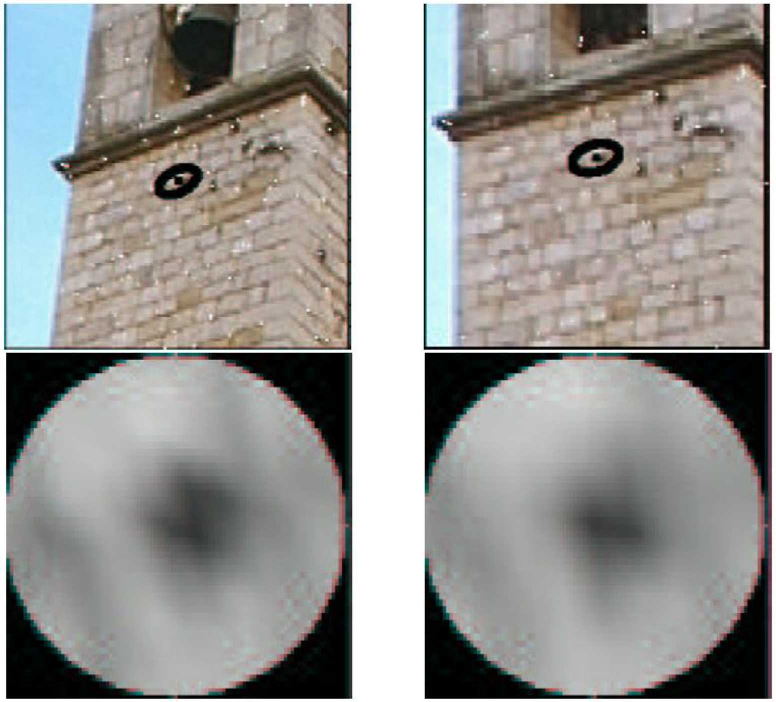

From Schaffalitzky and Zisserman ’02

Wide-Baseline Feature Matching

- Compare the distance, d1, to the closest feature, to the distance, d2, to the second closest feature
- Accept if d1/d2 < 0.6
  - If the ratio of distances is less than a threshold, keep the feature
- Why the ratio test?
  - Eliminates hard-to-match repeated features
  - Distances in SIFT descriptor space seem to be non-uniform

Feature Matching

- Exhaustive search
  - for each feature in one image, look at all the other features in the other image(s)
- Hashing
  - compute a short descriptor from each feature vector, or hash longer descriptors (randomly)
- Nearest neighbor techniques
  - k-trees and their variants

Wide-Baseline Feature Matching

- Because of the high dimensionality of features, *approximate nearest neighbors* are necessary for efficient performance

- See ANN package, Mount and Arya
  http://www.cs.umd.edu/~mount/ANN/

Student paper presentation

Color Conceptualization

X. Hou and L. Zhang

ACM Multimedia 2007

Presenter: Li, Chi

Student paper presentation

Photographic Tone Reproduction for Digital Images

E. Reinhard, M. Stark, P. Shirley, and J. Ferwerda
ACM SIGGRAPH 2012

Presenter: Liu, Jiawei

Next Time

- Panorama
  - Homography estimation
  - Cylindrical panorama
  - Blending
- Student paper presentations
  - 04/28: Seward, Garrett
    - Color image colorization
      R. Zhang, P. Isola, and A. Efros, ECCV 2016
  - 04/28: Peters, Emerson
    - Burst photography for high dynamic range and low-light imaging on mobile cameras.
      Samuel W. Hasinoff, et al. SIGGRAPH Asia 2016