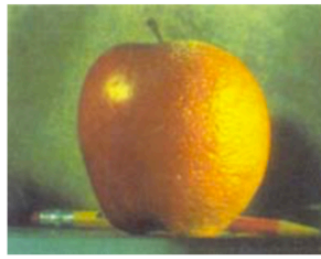


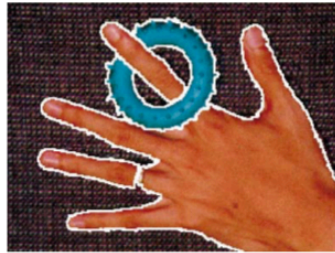
2. Image Formation



3. Image Processing



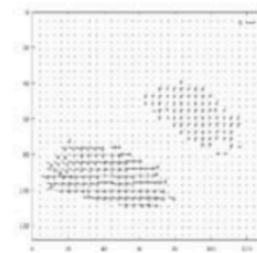
4. Features



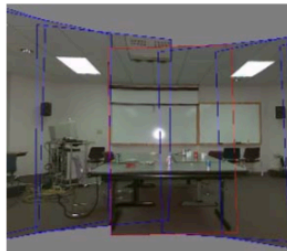
5. Segmentation



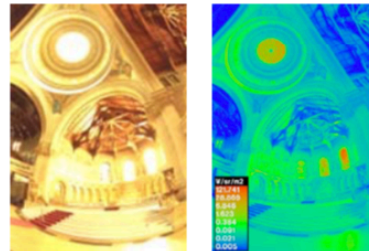
6-7. Structure from Motion



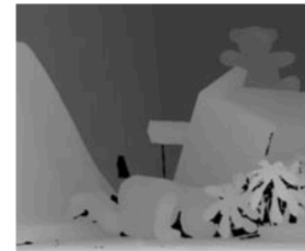
8. Motion



9. Stitching



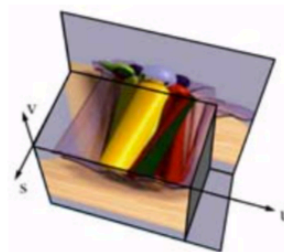
10. Computational Photography



11. Stereo



12. 3D Shape



13. Image-based Rendering



14. Recognition

|       |  |     |
|-------|--|-----|
| 4.1   | Points and patches . . . . .                               | 207 |
| 4.1.1 | Feature detectors . . . . .                                | 209 |
| 4.1.2 | Feature descriptors . . . . .                              | 222 |
| 4.1.3 | Feature matching . . . . .                                 | 225 |
| 4.1.4 | Feature tracking . . . . .                                 | 235 |
| 4.1.5 | <i>Application: Performance-driven animation</i> . . . . . | 237 |
| 4.2   | Edges . . . . .  | 238 |
| 4.2.1 | Edge detection . . . . .                                   | 238 |
| 4.2.2 | Edge linking . . . . .                                     | 244 |
| 4.2.3 | <i>Application: Edge editing and enhancement</i> . . . . . | 249 |
| 4.3   | Lines . . . . .  | 250 |
| 4.3.1 | Successive approximation . . . . .                         | 250 |
| 4.3.2 | Hough transforms . . . . .                                 | 251 |
| 4.3.3 | Vanishing points . . . . .                                 | 254 |
| 4.3.4 | <i>Application: Rectangle detection</i> . . . . .          | 257 |

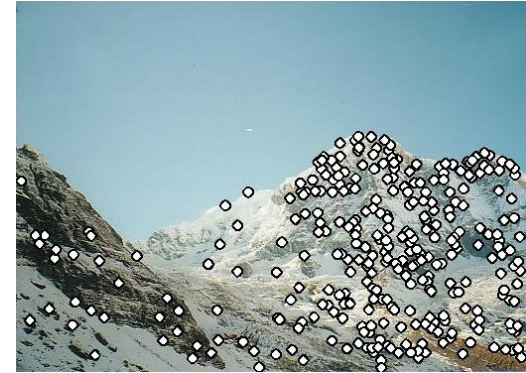
|       |  |     |
|-------|--|-----|
| 4.1   | Points and patches . . . . .                               | 207 |
| 4.1.1 | Feature detectors . . . . .                                | 209 |
| 4.1.2 | Feature descriptors . . . . .                              | 222 |
| 4.1.3 | Feature matching . . . . .                                 | 225 |
| 4.1.4 | Feature tracking . . . . .                                 | 235 |
| 4.1.5 | <i>Application: Performance-driven animation</i> . . . . . | 237 |
| 4.2   | Edges . . . . .  | 238 |
| 4.2.1 | Edge detection . . . . .                                   | 238 |
| 4.2.2 | Edge linking . . . . .                                     | 244 |
| 4.2.3 | <i>Application: Edge editing and enhancement</i> . . . . . | 249 |
| 4.3   | Lines . . . . .  | 250 |
| 4.3.1 | Successive approximation . . . . .                         | 250 |
| 4.3.2 | Hough transforms . . . . .                                 | 251 |
| 4.3.3 | Vanishing points . . . . .                                 | 254 |
| 4.3.4 | <i>Application: Rectangle detection</i> . . . . .          | 257 |

# Detectors



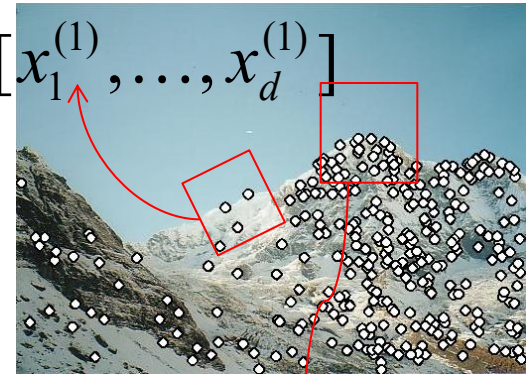
# Local features: main components

1) **Detection:** Identify the interest points



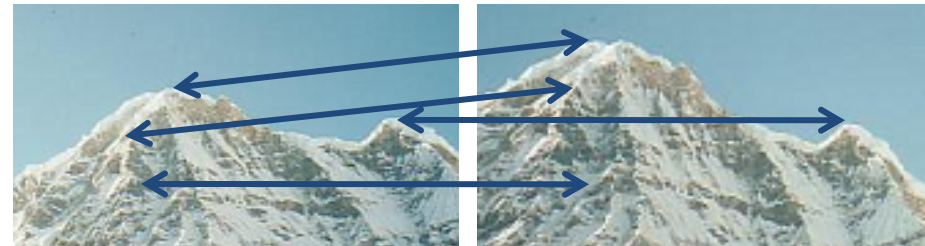
2) **Description:** Extract vector feature descriptor surrounding each interest point.

feature descriptor surrounding  $\mathbf{x}_1 = [x_1^{(1)}, \dots, x_d^{(1)}]$



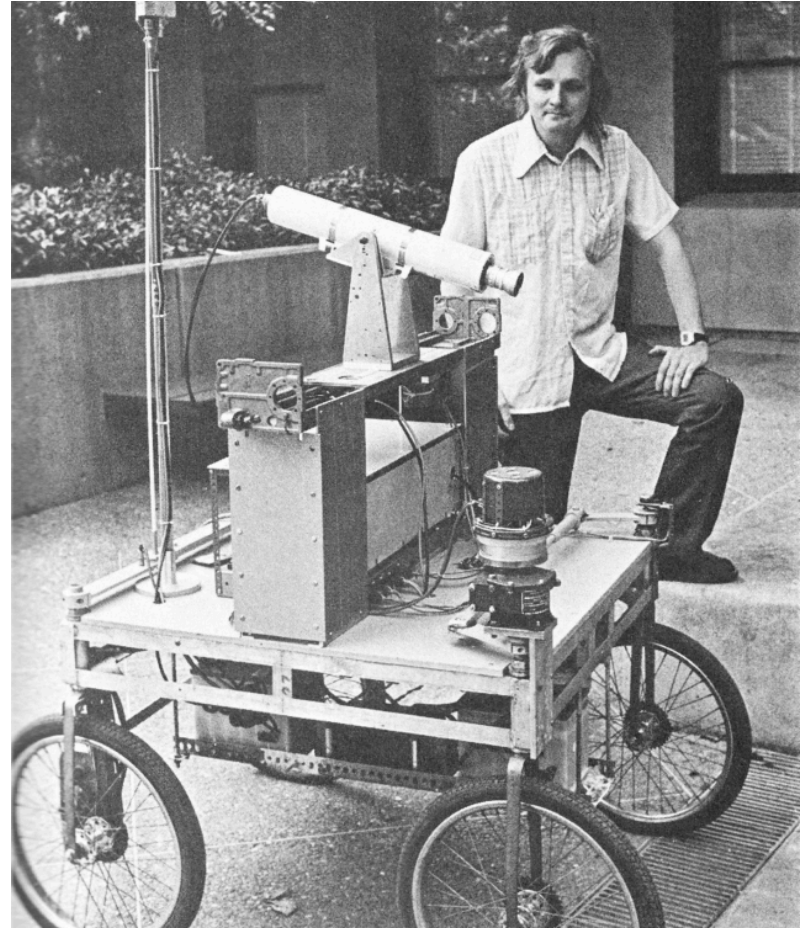
3) **Matching:** Determine correspondence between descriptors in two views

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



# History

- Moravec 1980
- Harris Corners 1988
- [Wolf & Platt 1993: FCN!]
- SIFT (Lowe) 2004
- FAST 2006 (learning!)
- SURF 2006
- ORB 2011



# Harris corner detector

- 1) Compute  $M$  matrix for each image window to get their *cornerness* scores.
- 2) Find points whose surrounding window gave large corner response ( $f >$  threshold)
- 3) Take the points of local maxima, i.e., perform non-maximum suppression

C.Harris and M.Stephens. [“A Combined Corner and Edge Detector.”](#)  
*Proceedings of the 4th Alvey Vision Conference*: pages 147—151, 1988.

# Harris Detector [Harris88]

- Second moment matrix

$$\mu(\sigma_I, \sigma_D) = g(\sigma_I) * \begin{bmatrix} I_x^2(\sigma_D) & I_x I_y(\sigma_D) \\ I_x I_y(\sigma_D) & I_y^2(\sigma_D) \end{bmatrix}$$

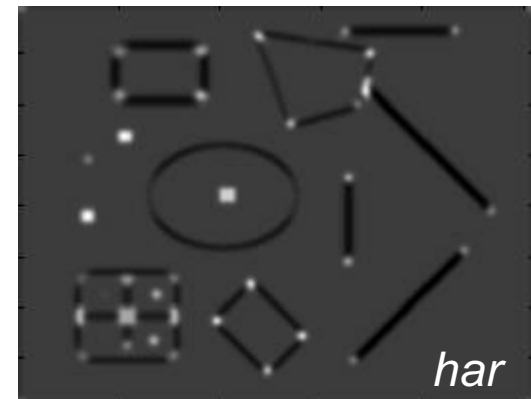
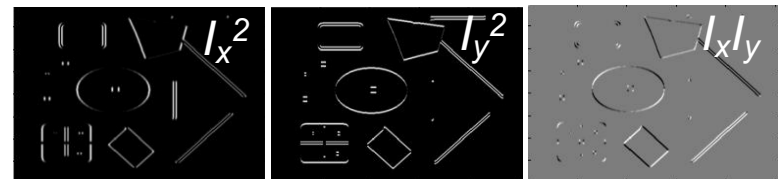
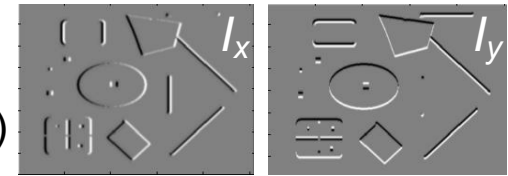
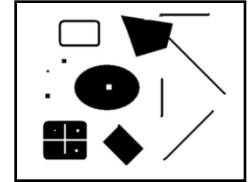
$$\det M = \lambda_1 \lambda_2$$

$$\text{trace } M = \lambda_1 + \lambda_2$$

2. Square of derivatives

3. Gaussian filter  $g(\sigma)$

1. Image derivatives  
(optionally, blur first)



4. Cornerness function – both eigenvalues are strong

$$\text{har} = \det[\mu(\sigma_I, \sigma_D)] - \alpha [\text{trace}(\mu(\sigma_I, \sigma_D))]^2 = g(I_x^2)g(I_y^2) - [g(I_x I_y)]^2 - \alpha [g(I_x^2) + g(I_y^2)]^2$$

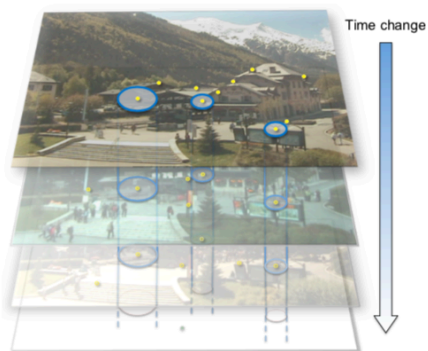
5. Non-maxima suppression

# Deep Detectors

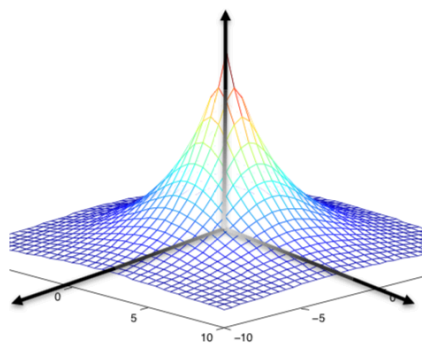
# TILDE: A Temporally Invariant Learned DETector

CVPR 2015

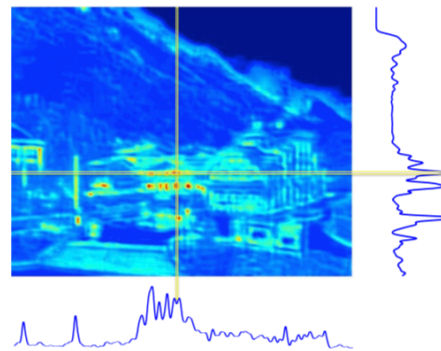
Yannick Verdie<sup>1,\*</sup> Kwang Moo Yi<sup>1,\*</sup> Pascal Fua<sup>1</sup> Vincent Lepetit<sup>2</sup>  
<sup>1</sup>Computer Vision Laboratory, École Polytechnique Fédérale de Lausanne (EPFL)  
<sup>2</sup>Institute for Computer Graphics and Vision, Graz University of Technology



(a) Stack of training images



(b) Desired response on positive samples



(c) Regressor response for a new image



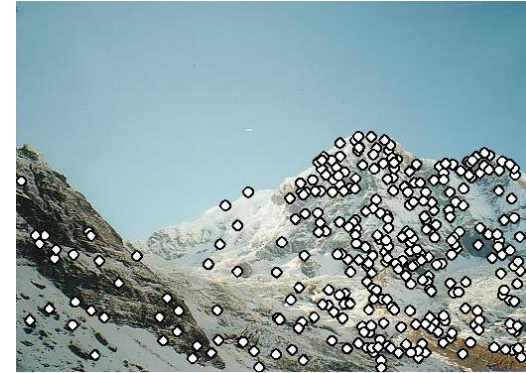
(d) Keypoints detected in the new image

- Train on images from webcams: fixed view, different times
- Learn CNN-like regressor
- Loss = repeatability

# Descriptors

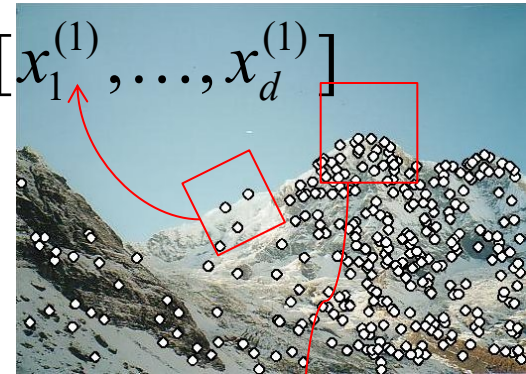
# Local features: main components

1) Detection: Identify the interest points



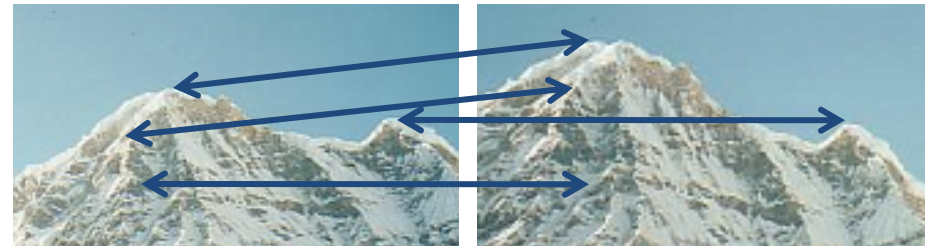
2) **Description:** Extract vector feature descriptor surrounding each interest point.

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



3) Matching: Determine correspondence between descriptors in two views

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





# Image representations

- Templates

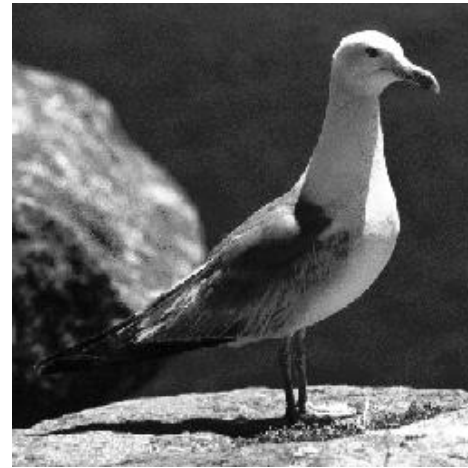
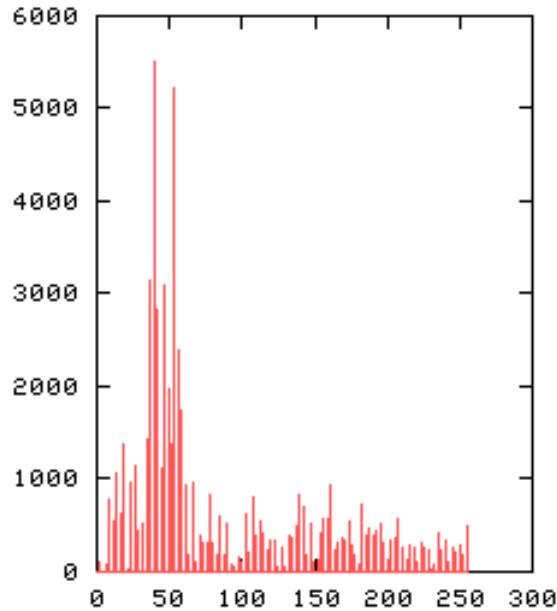
- Intensity, gradients, etc.



- Histograms

- Color, texture, SIFT descriptors, etc.

# Image Representations: Histograms

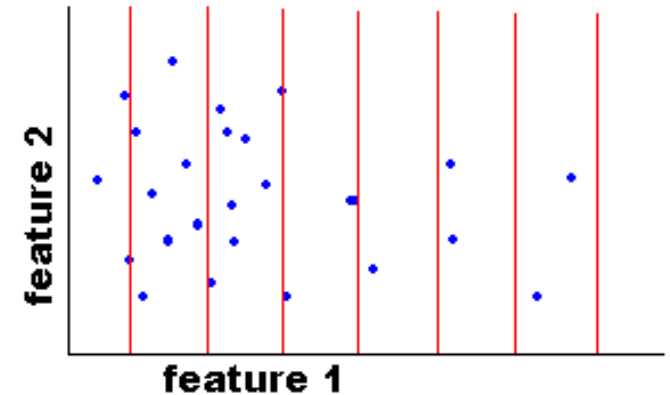
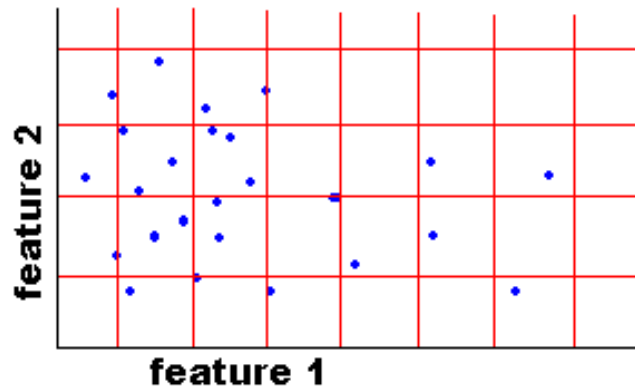
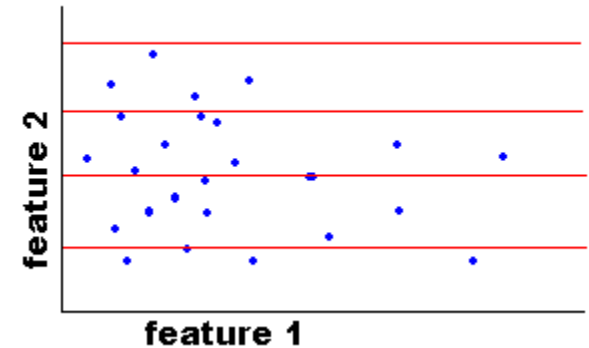
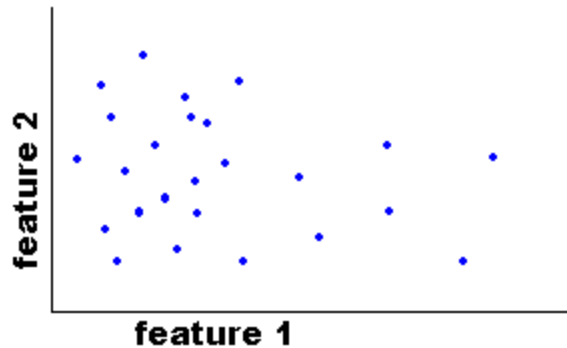


## Global histogram

- Represent distribution of features
  - Color, texture, depth, ...

# Image Representations: Histograms

Histogram: Probability or count of data in each bin



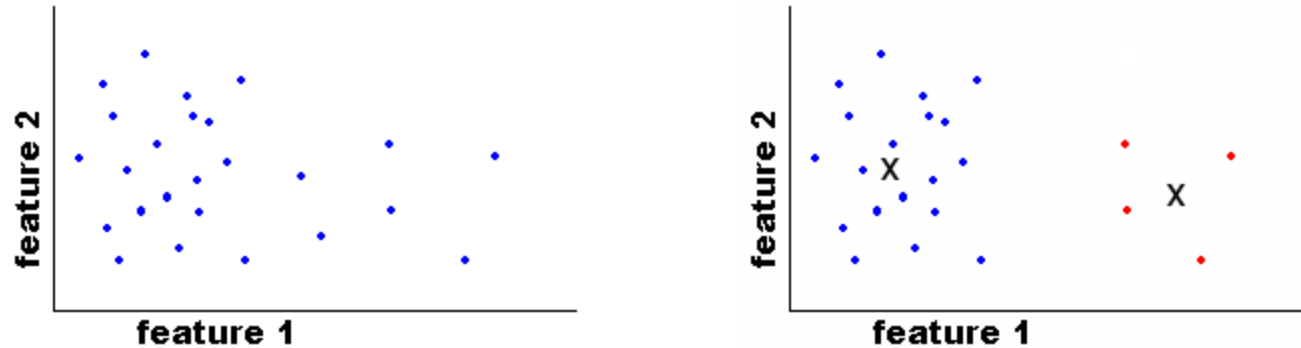
- Joint histogram
  - Requires lots of data
  - Loss of resolution to avoid empty bins

## Marginal histogram

- Requires independent features
- More data/bin than joint histogram

# Image Representations: Histograms

## Clustering



Use the same cluster centers for all images

# What kind of things do we compute histograms of?

- Histograms of oriented gradients

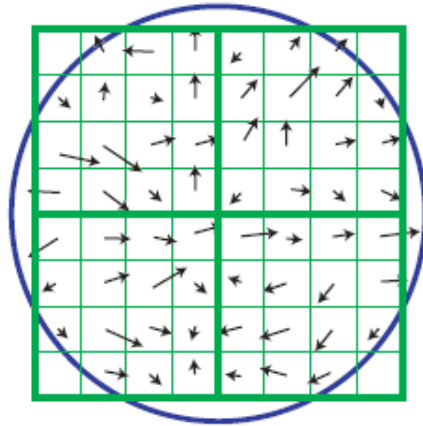
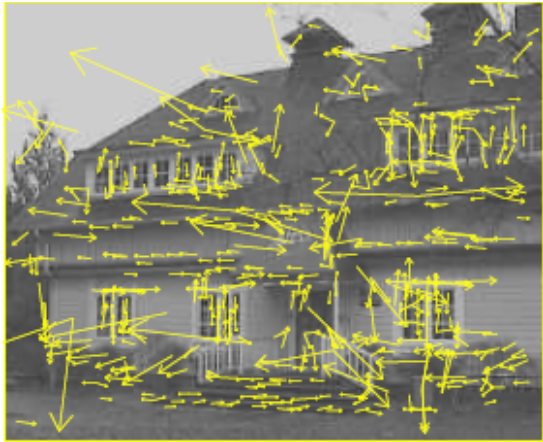
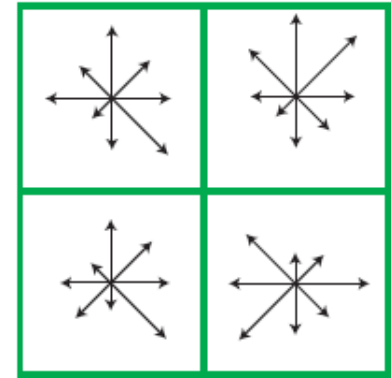


Image gradients

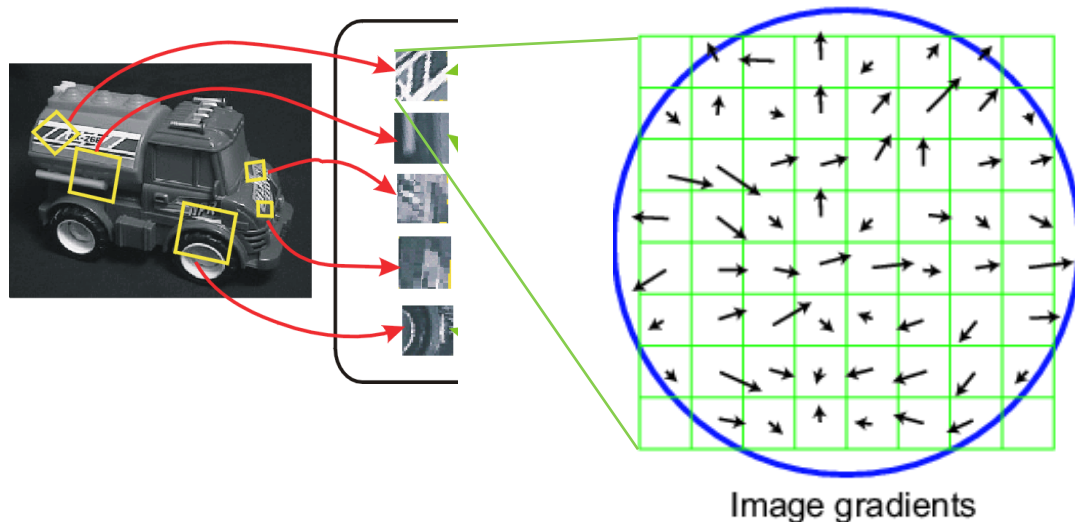


Keypoint descriptor

SIFT – Lowe IJCV 2004

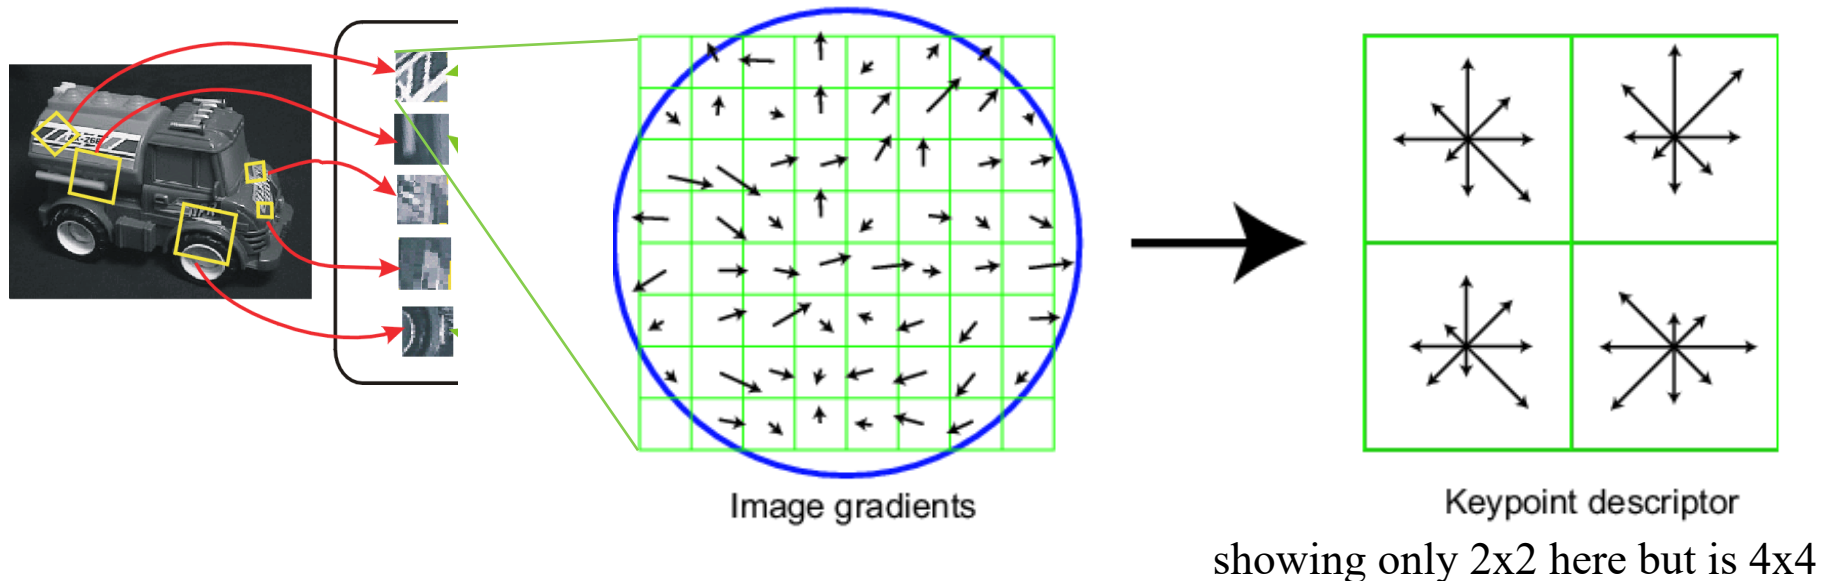
# SIFT vector formation

- Computed on rotated and scaled version of window according to computed orientation & scale
  - resample the window
- Based on gradients weighted by a Gaussian of variance half the window (for smooth falloff)



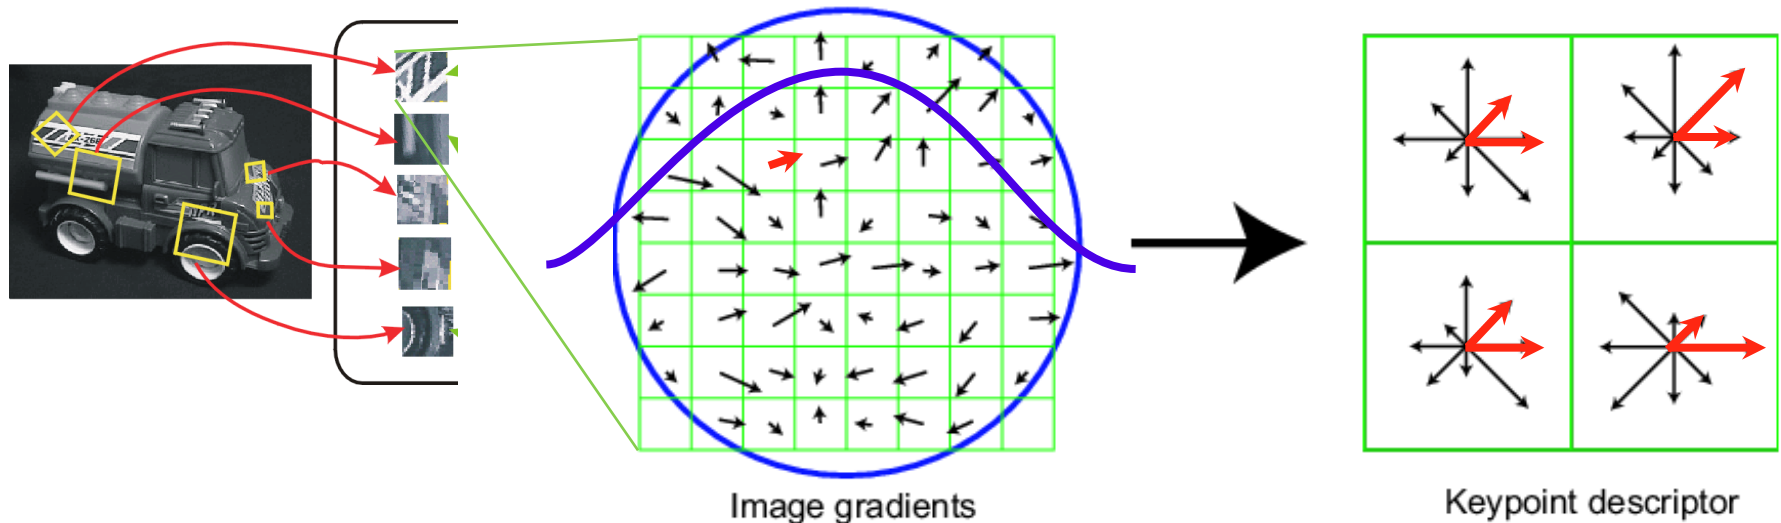
# SIFT vector formation

- 4x4 array of gradient orientation histogram weighted by magnitude
- 8 orientations x 4x4 array = 128 dimensions
- Motivation: some sensitivity to spatial layout, but not too much.



# Ensure smoothness

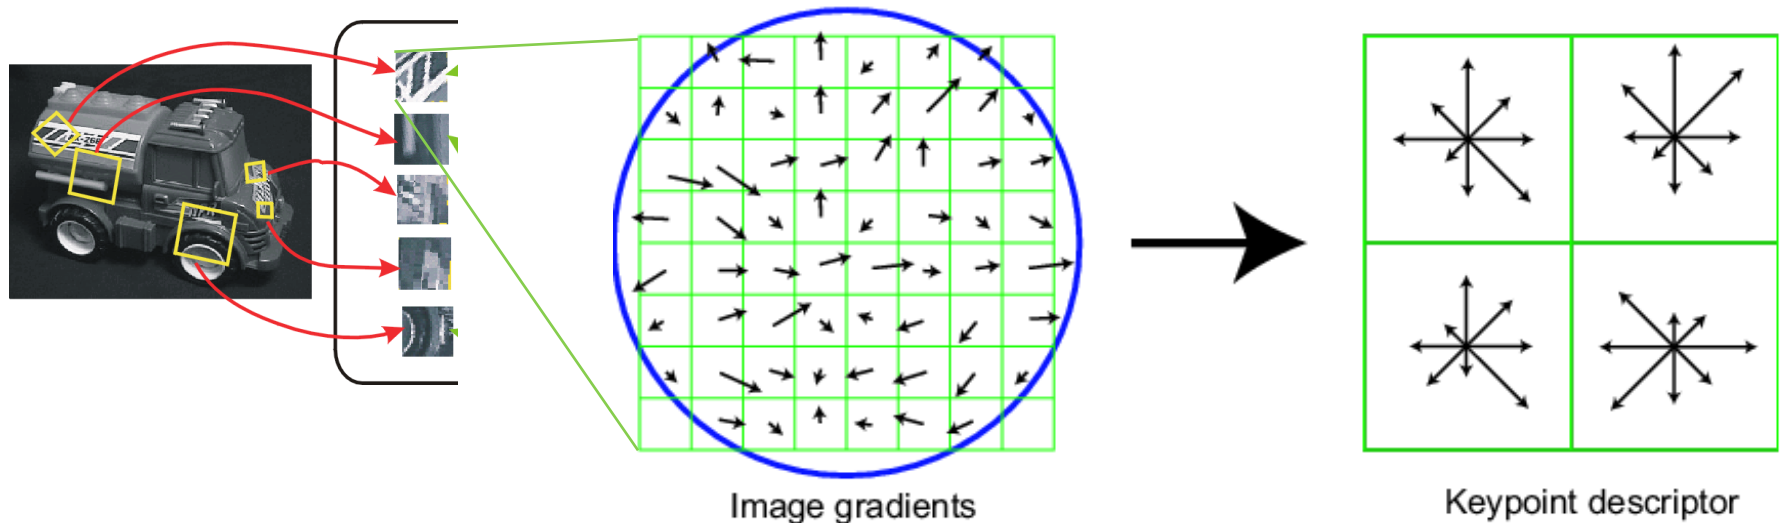
- Gaussian weight
- Interpolation
  - a given gradient contributes to 8 bins:  
4 in space times 2 in orientation



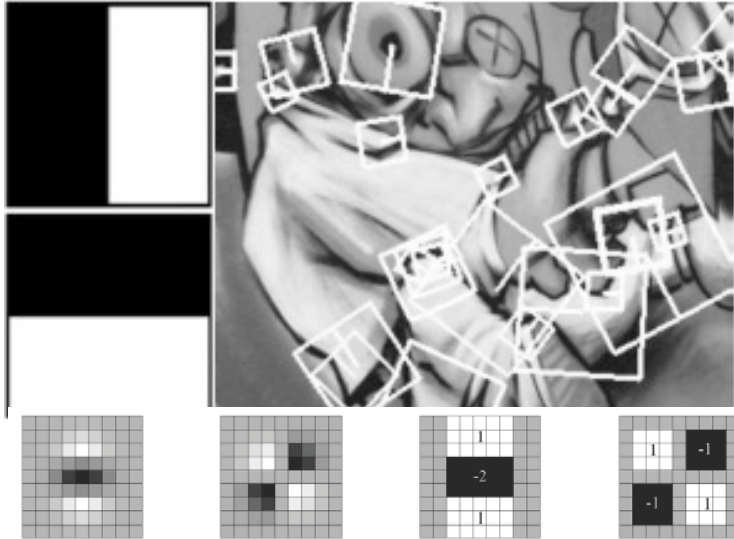


# Reduce effect of illumination

- 128-dim vector normalized to 1
- Threshold gradient magnitudes to avoid excessive influence of high gradients
  - after normalization, clamp gradients  $>0.2$
  - renormalize



# Local Descriptors: SURF



## Fast approximation of SIFT idea

Efficient computation by 2D box filters & integral images

⇒ 6 times faster than SIFT

Equivalent quality for object identification

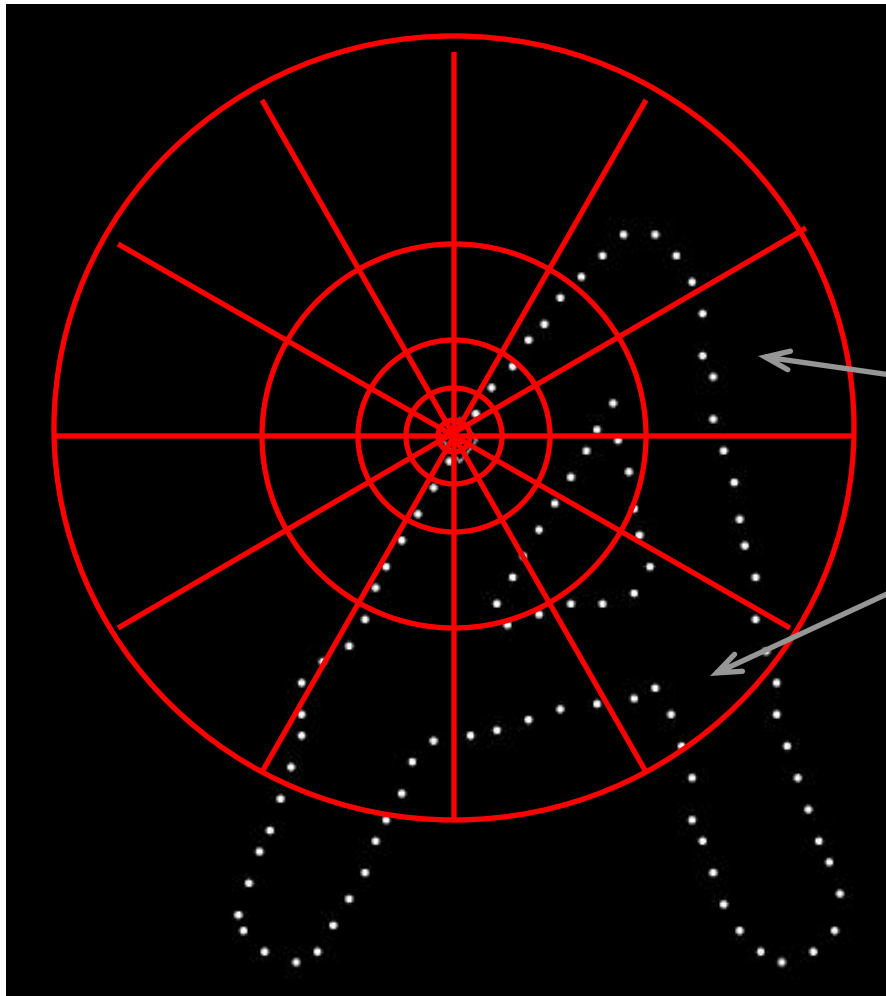
## GPU implementation available

Feature extraction @ 200Hz

(detector + descriptor, 640×480 img)

<http://www.vision.ee.ethz.ch/~surf>

# Local Descriptors: Shape Context



Count the number of points inside each bin, e.g.:

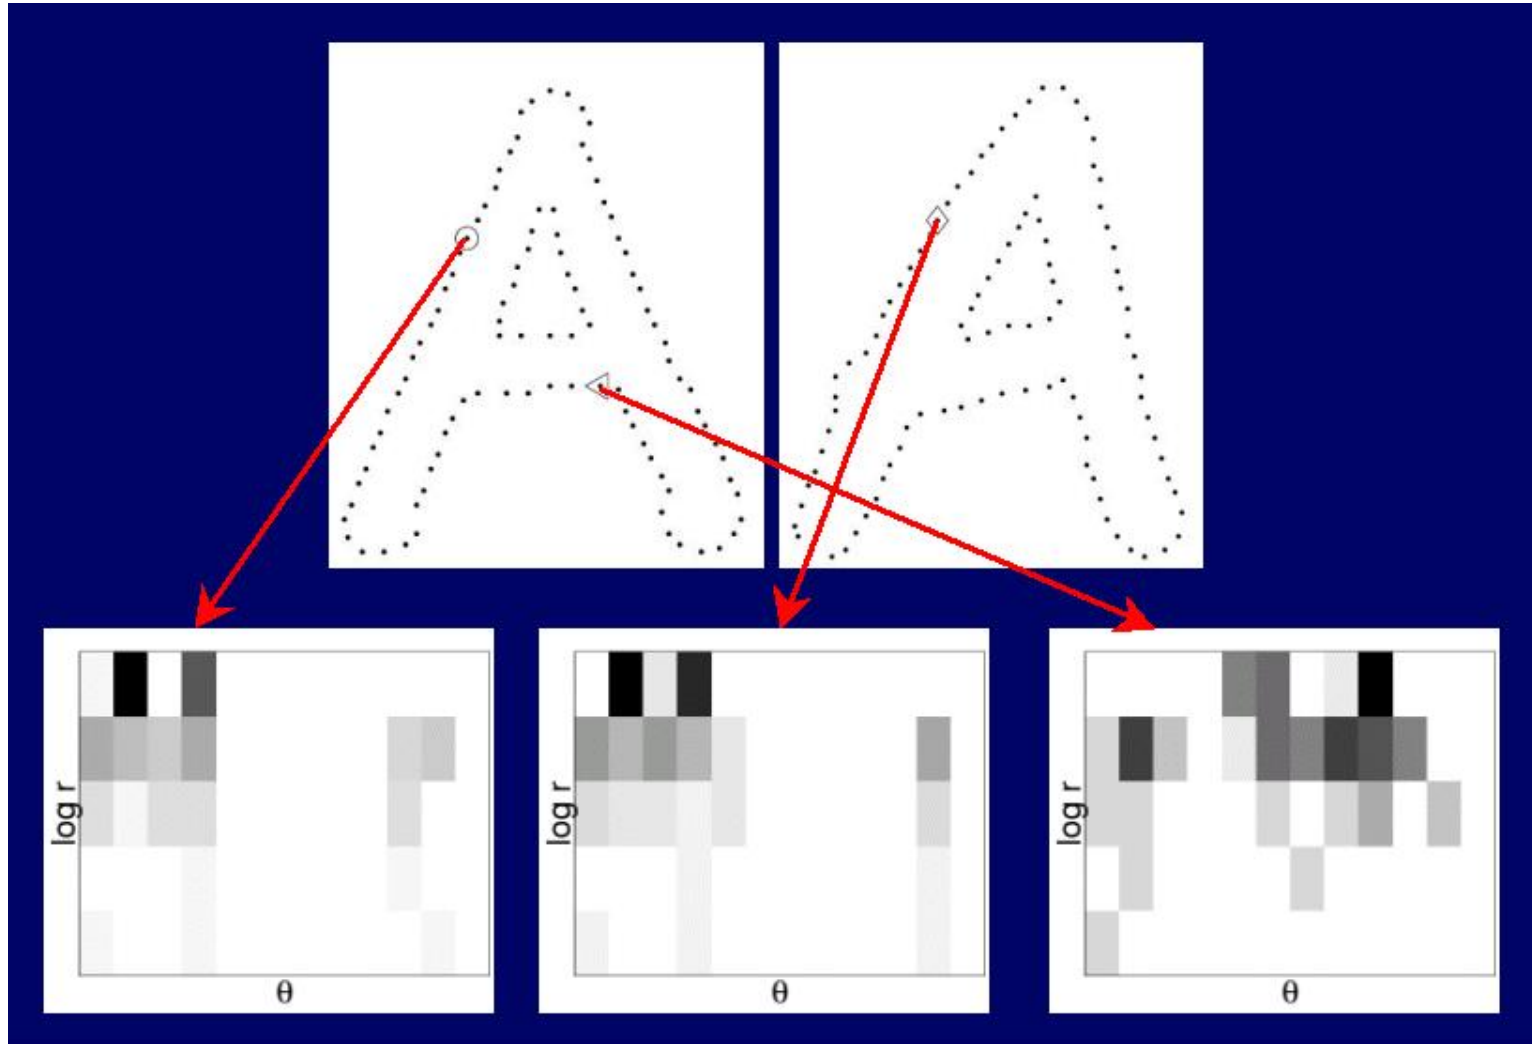
Count = 4

⋮

Count = 10

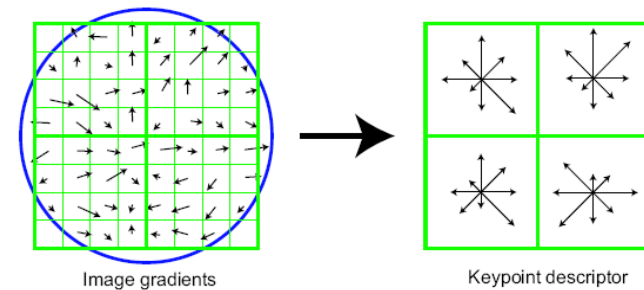
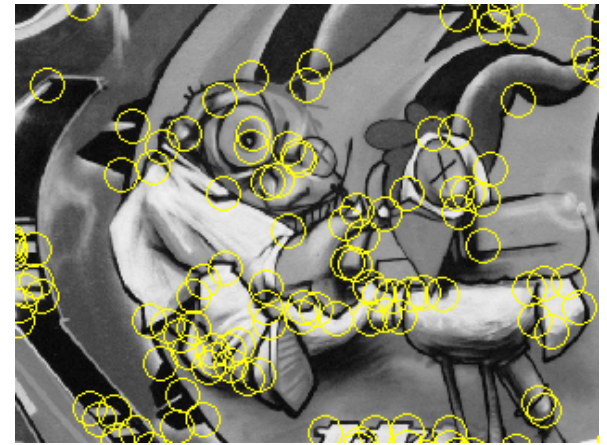
Log-polar binning: more precision for nearby points, more flexibility for farther points.

# Shape Context Descriptor



# Things to remember

- Keypoint detection: repeatable and distinctive
  - Corners, blobs, stable regions
  - Harris, DoG
- Descriptors: robust and selective
  - spatial histograms of orientation
  - SIFT



# Deep Descriptors

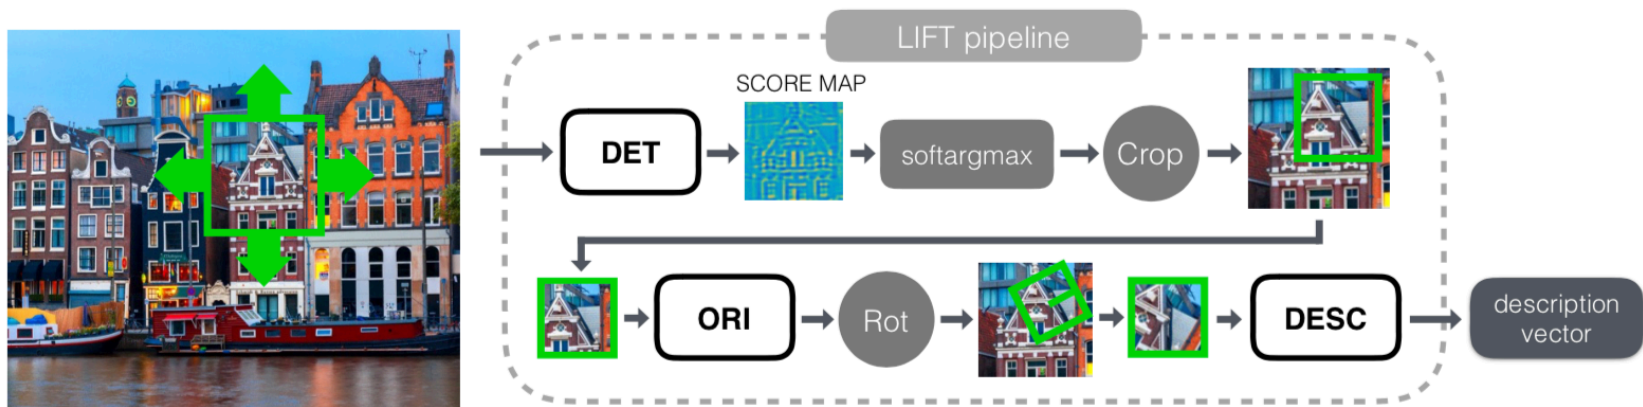
# LIFT: Learned Invariant Feature Transform

## ECCV 2016

Kwang Moo Yi<sup>\*,1</sup>, Eduard Trulls<sup>\*,1</sup>, Vincent Lepetit<sup>2</sup>, Pascal Fua<sup>1</sup>

<sup>1</sup>Computer Vision Laboratory, Ecole Polytechnique Fédérale de Lausanne (EPFL)

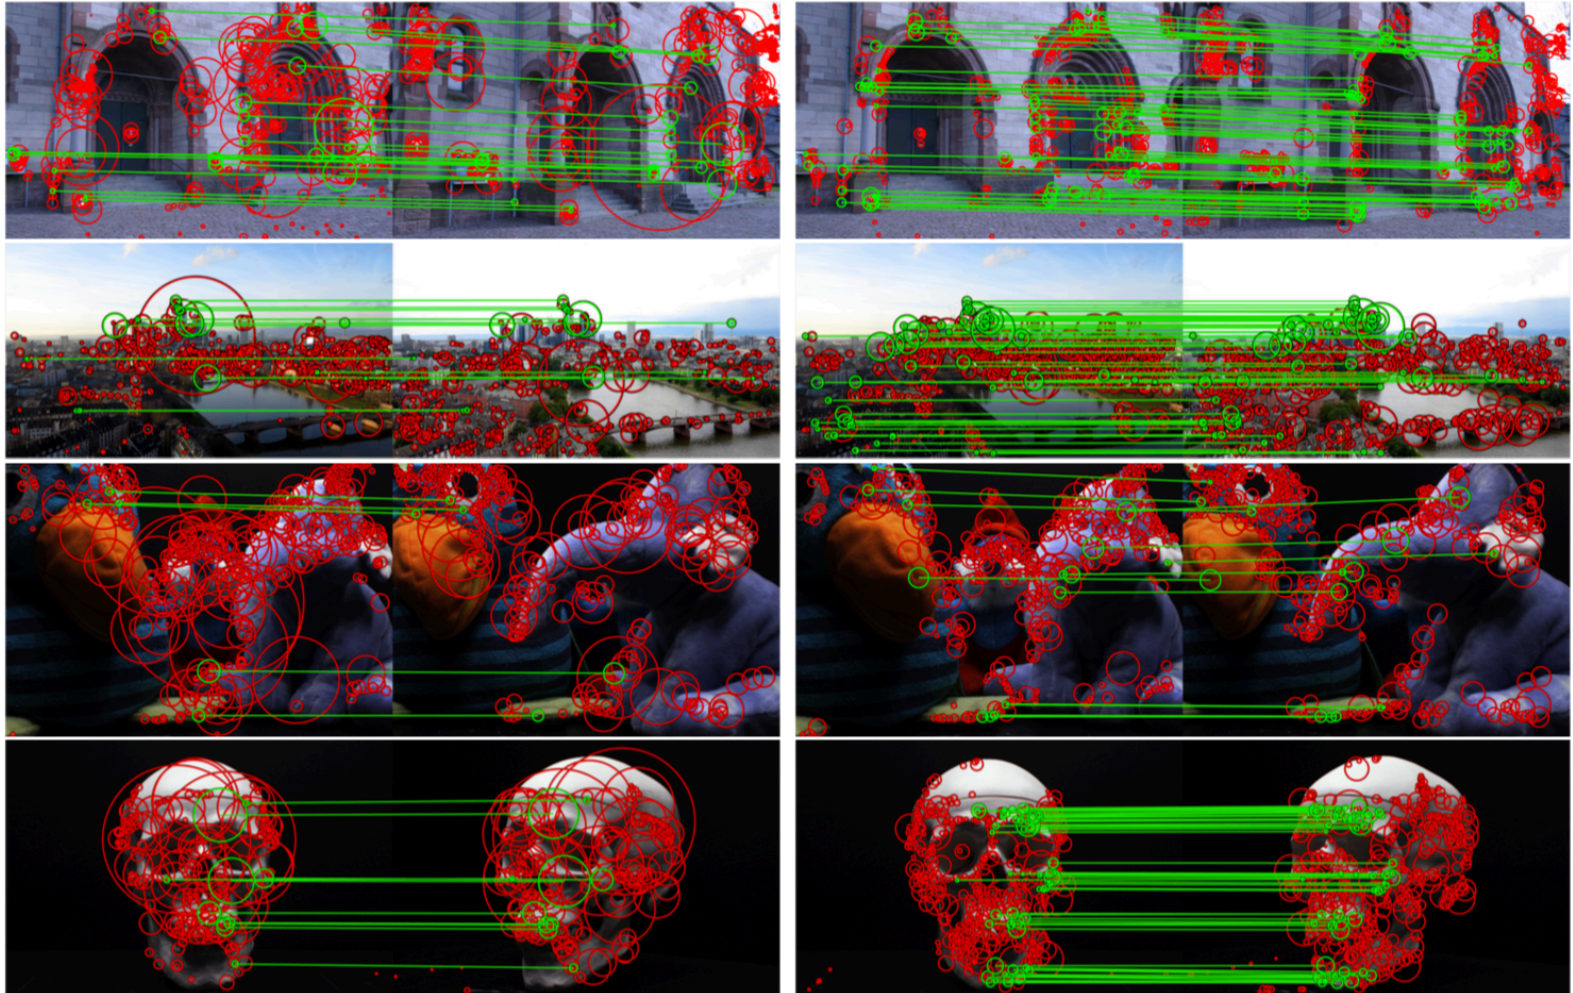
<sup>2</sup>Institute for Computer Graphics and Vision, Graz University of Technology



- Three networks: detection, orientation, description
- detection+orientation -> STN -> descriptor
- Trained separately :-)



# SIFT vs. LIFT





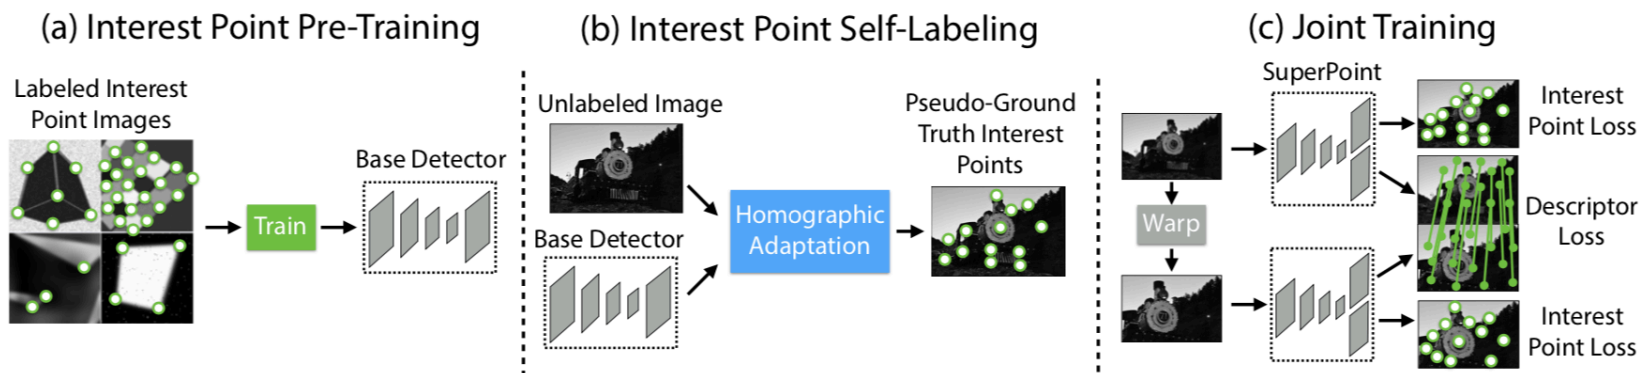
# SuperPoint: Self-Supervised Interest Point Detection and Description

2018 CVPR Workshop

Daniel DeTone  
Magic Leap  
Sunnyvale, CA

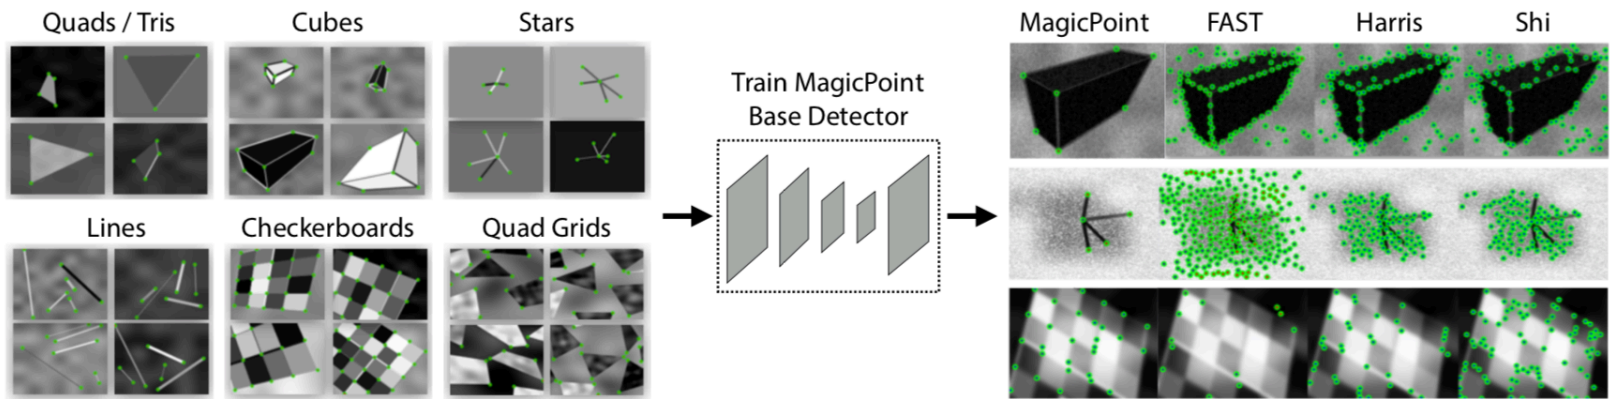
Tomasz Malisiewicz  
Magic Leap  
Sunnyvale, CA

Andrew Rabinovich  
Magic Leap  
Sunnyvale, CA

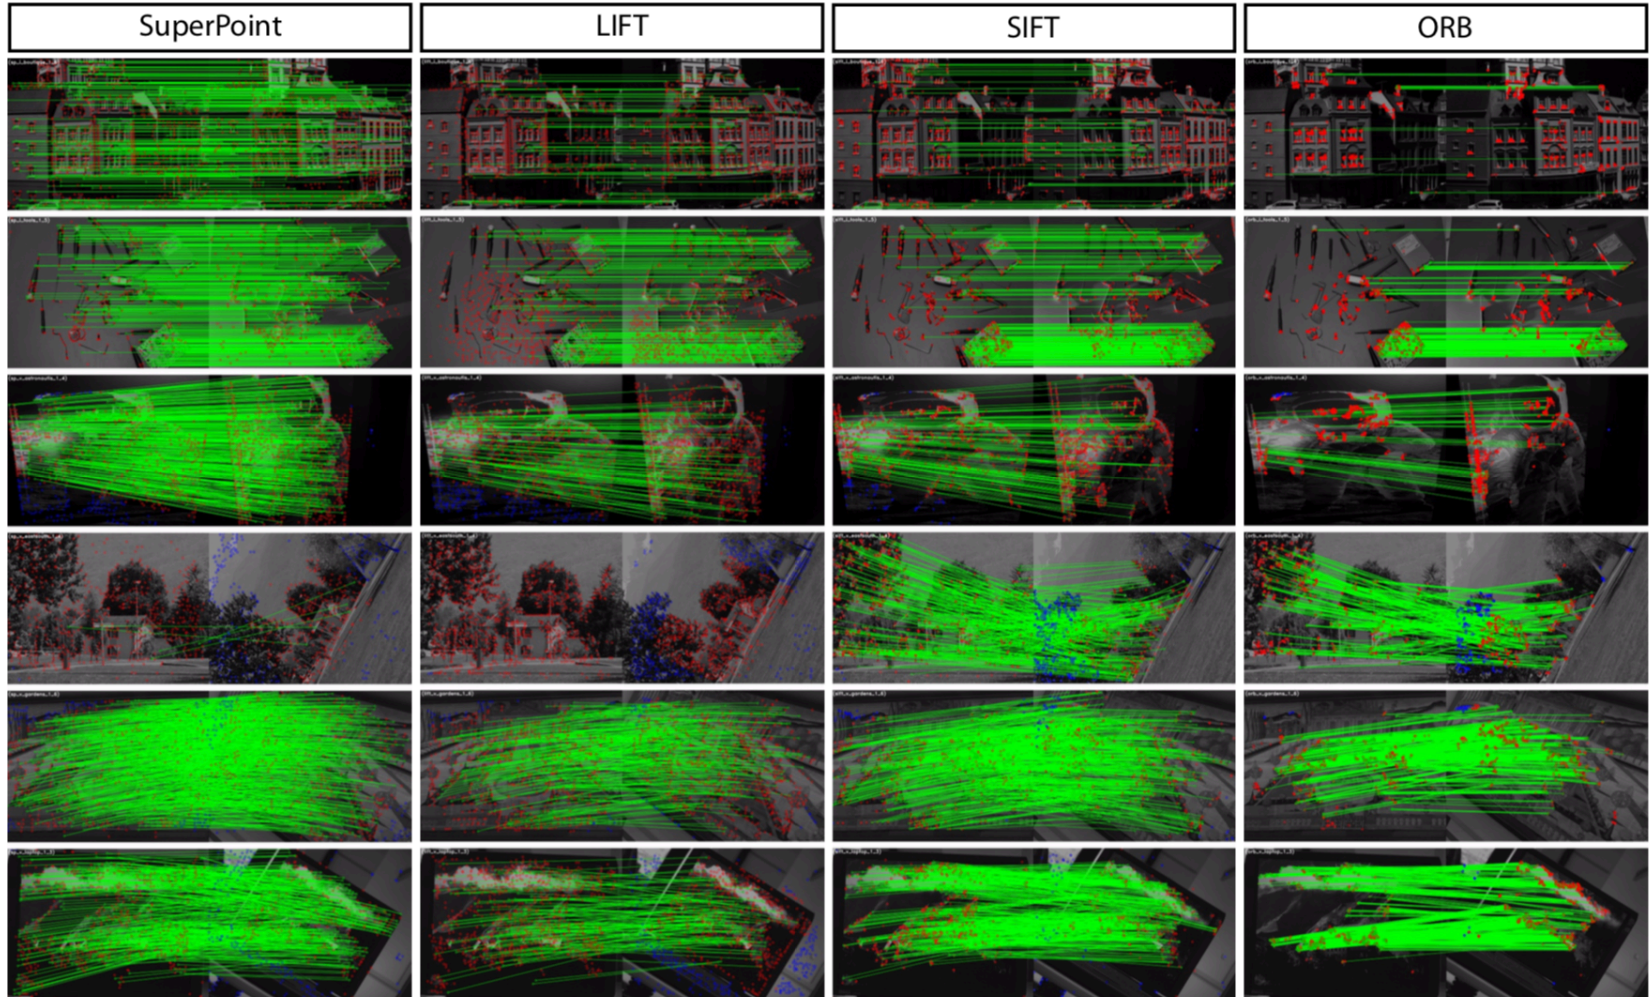


- Interest point = ill-defined -> self-supervised
- MagicPoint -> SuperPoint

# MagicPoint



# SuperPoint Results





# D2-Net: A Trainable CNN for *Joint Description and Detection of Local Features*

CVPR 2019

Mihai Dusmanu<sup>1,2,3</sup>

Ignacio Rocco<sup>1,2</sup>

Tomas Pajdla<sup>4</sup>

Marc Pollefeys<sup>3,5</sup>

Josef Sivic<sup>1,2,4</sup>

Akihiko Torii<sup>6</sup>

Torsten Sattler<sup>7</sup>

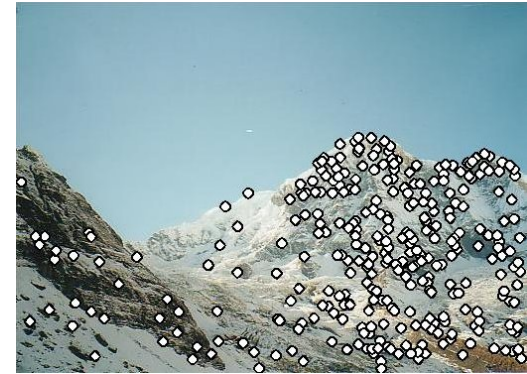
- Tensor viewed as descriptors and detector maps
- VGG16-based, loss encourages distinctiveness and repeatability
- Results beat the star of the art in day-night and indoor localization, but not in more traditional settings (Superpoint shines for Hatches, **GeoDesc** for SFM)



# Matching

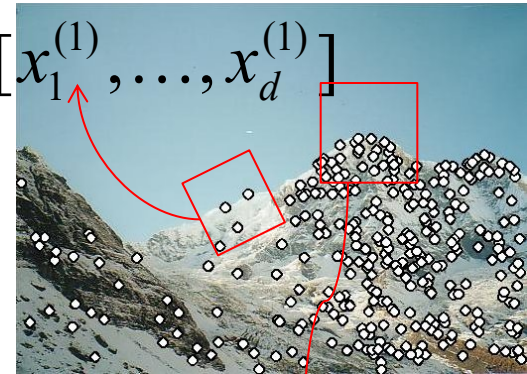
# Local features: main components

1) **Detection:** Identify the interest points



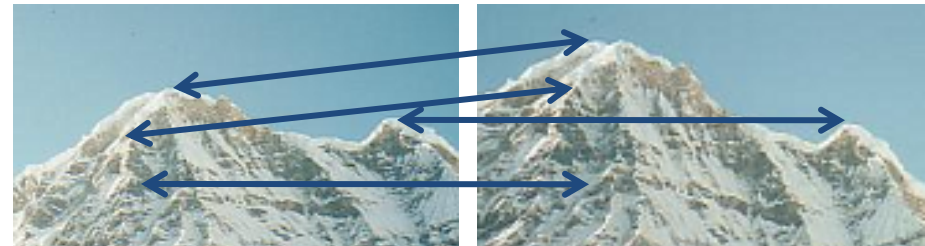
2) **Description:** Extract vector feature descriptor surrounding each interest point.

feature descriptor surrounding  $\mathbf{x}_1 = [x_1^{(1)}, \dots, x_d^{(1)}]$



3) **Matching:** Determine correspondence between descriptors in two views

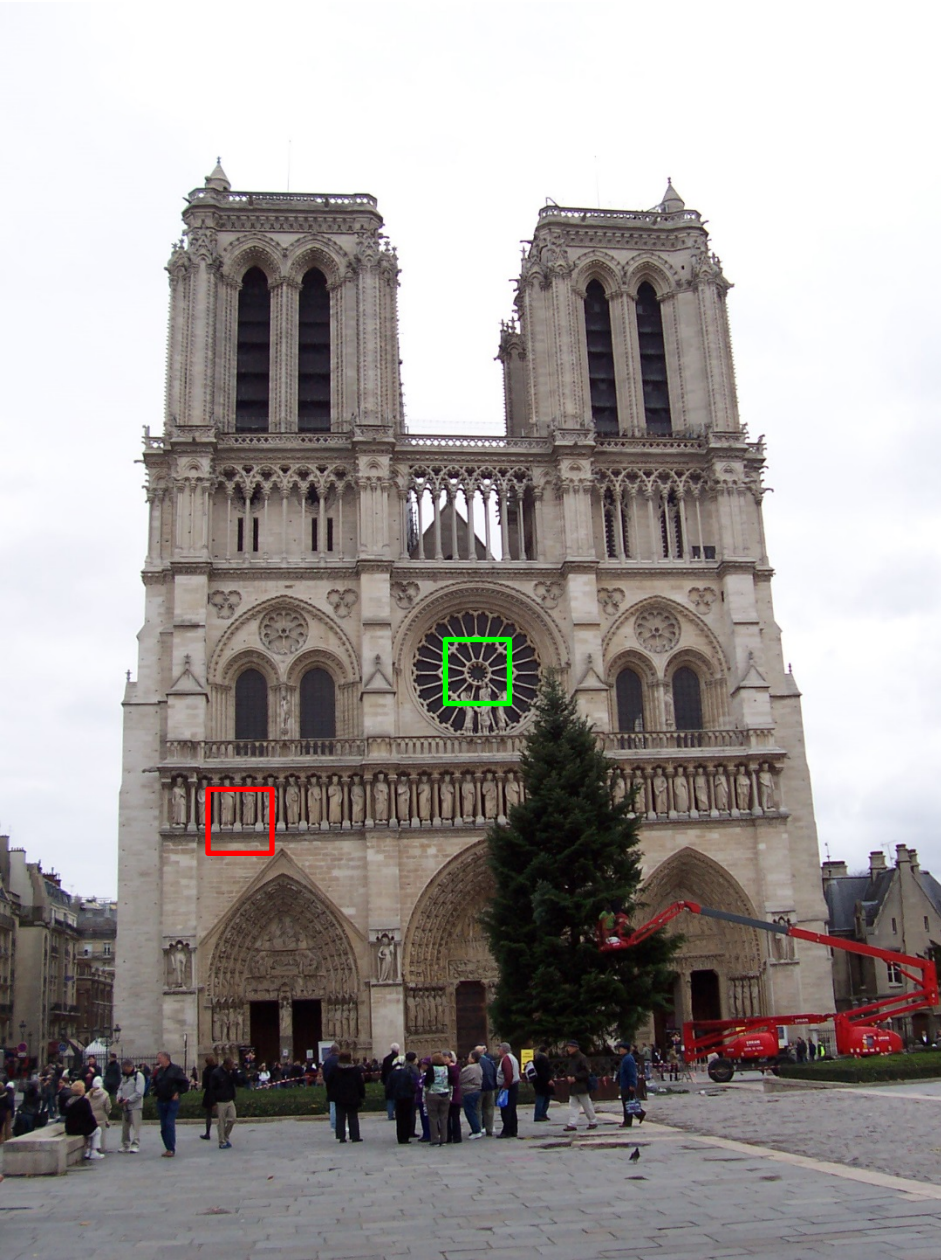
$\mathbf{x}_2 = [x_1^{(2)}, \dots, x_d^{(2)}]$



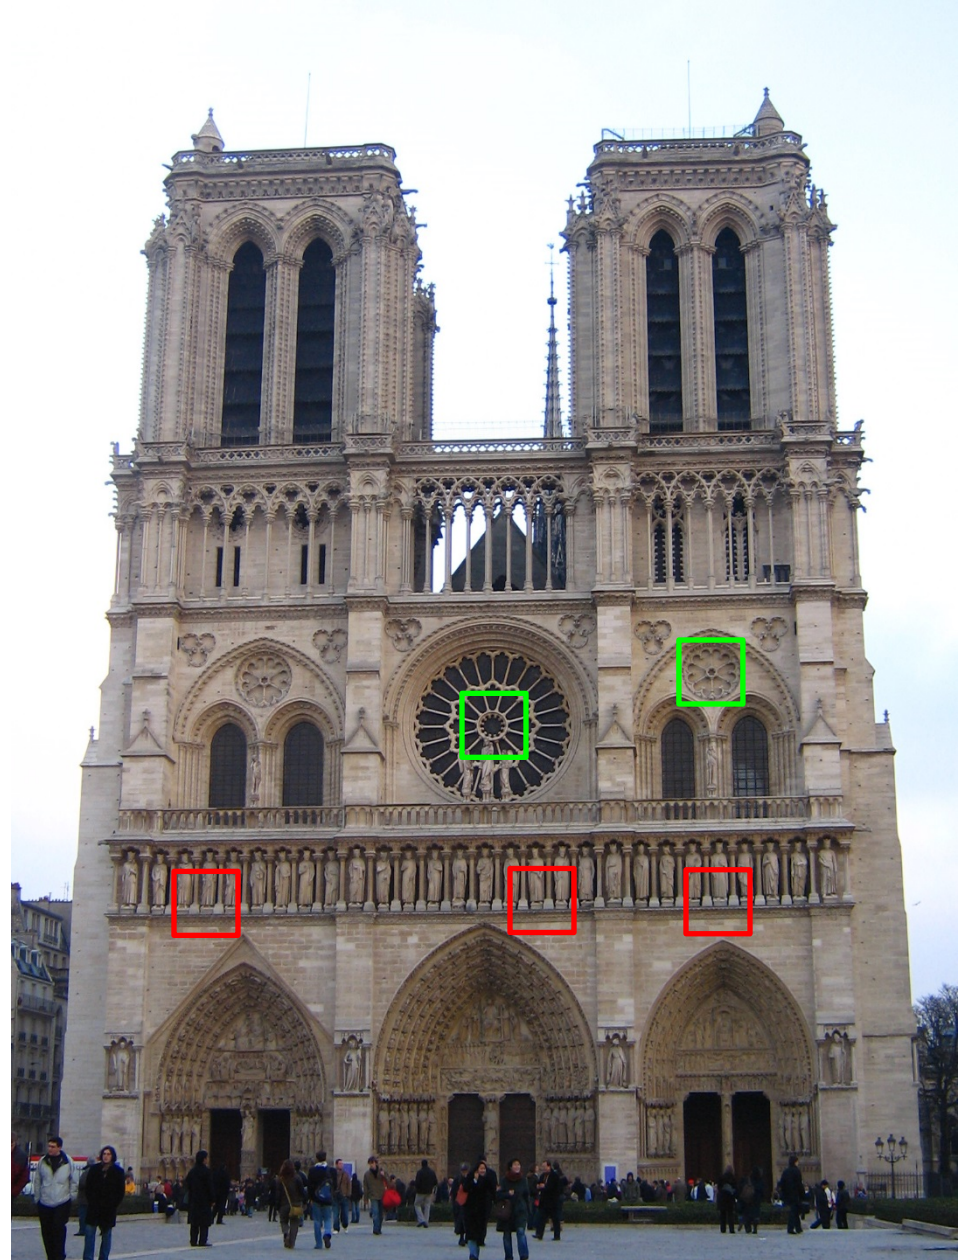
# Matching

- Simplest approach: Pick the nearest neighbor.  
Threshold on absolute distance
- Problem: Lots of self similarity in many photos





Distance: 0.34, 0.30, 0.40



Distance: 0.61  
Distance: 1.22

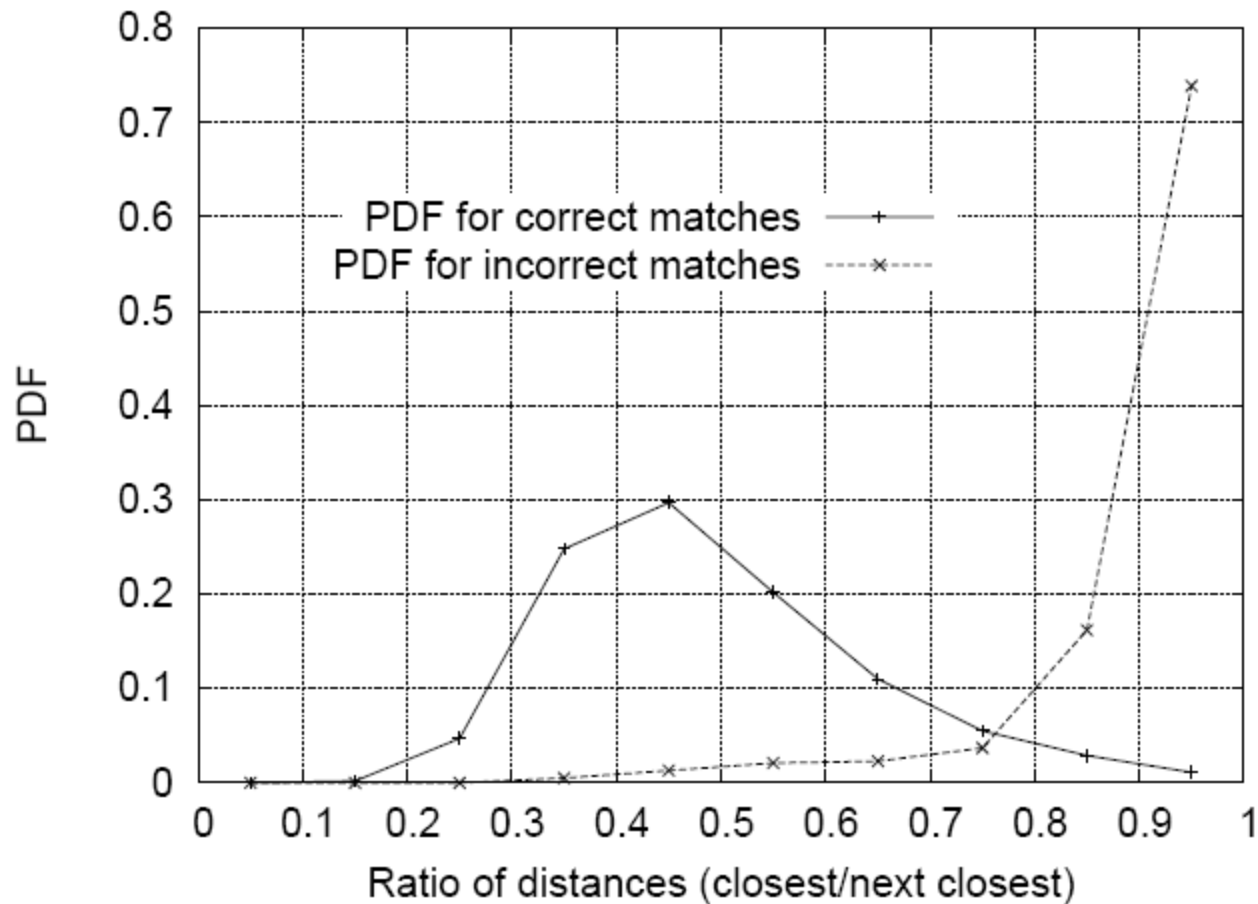


# Nearest Neighbor Distance Ratio

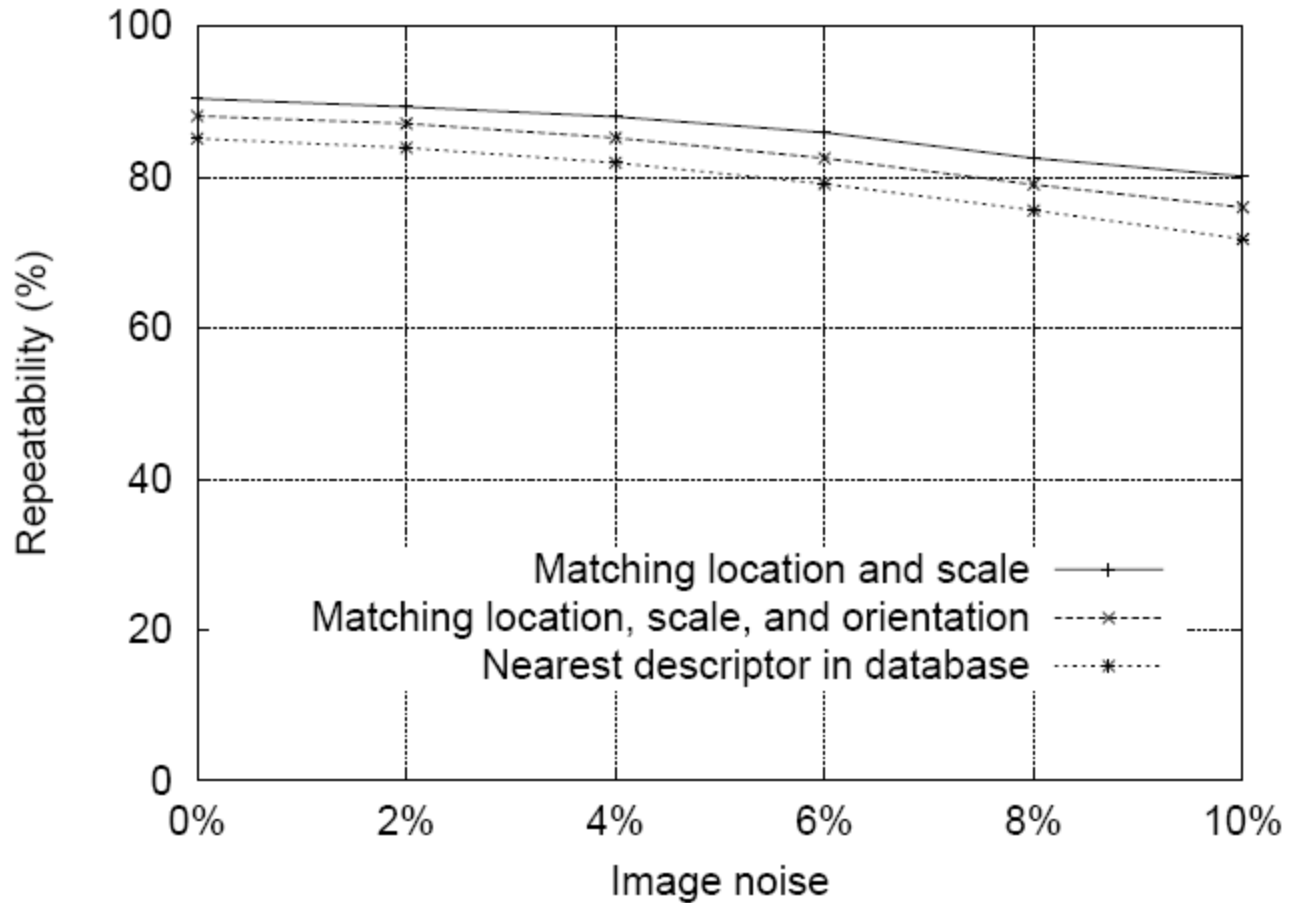
- $\frac{NN1}{NN2}$  where NN1 is the distance to the first nearest neighbor and NN2 is the distance to the second nearest neighbor.
- Sorting by this ratio puts matches in order of confidence.

# Matching Local Features

- Nearest neighbor (Euclidean distance)
- Threshold ratio of nearest to 2<sup>nd</sup> nearest descriptor



# SIFT Repeatability



# SIFT Repeatability

