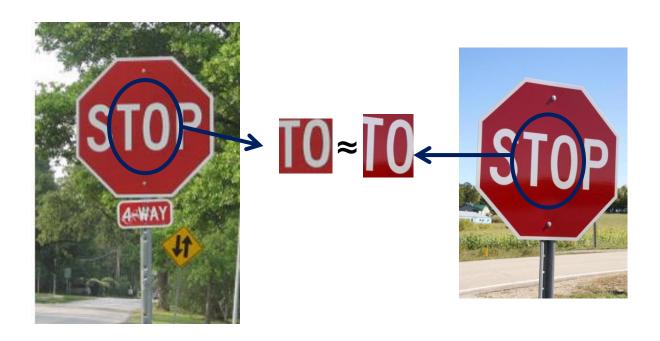
# Locating and Describing Interest Points

Computer Vision
CS 543 / ECE 549
University of Illinois

Derek Hoiem

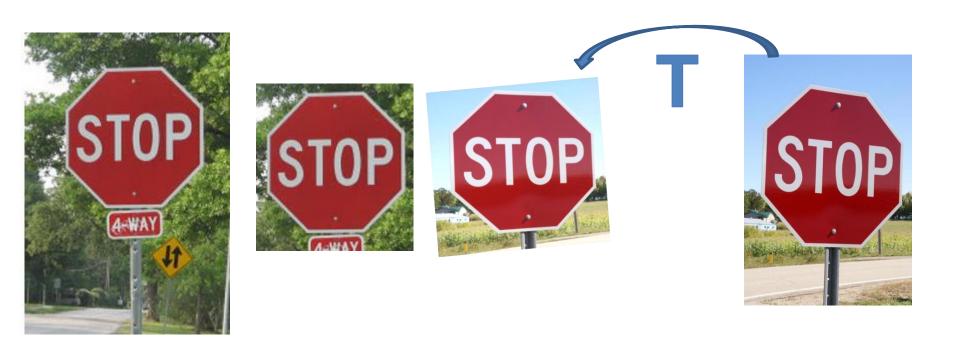
# This section: correspondence and alignment

 Correspondence: matching points, patches, edges, or regions across images

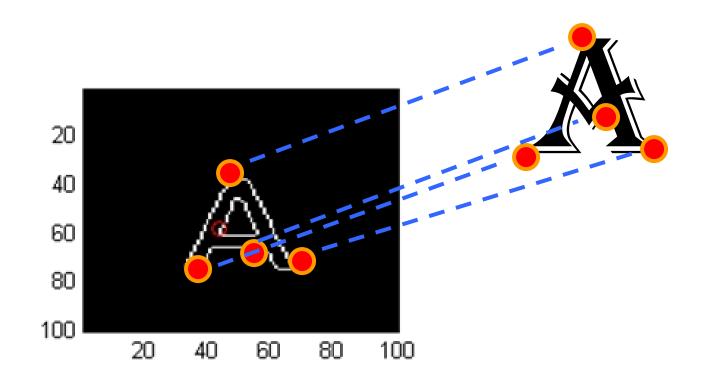


# This section: correspondence and alignment

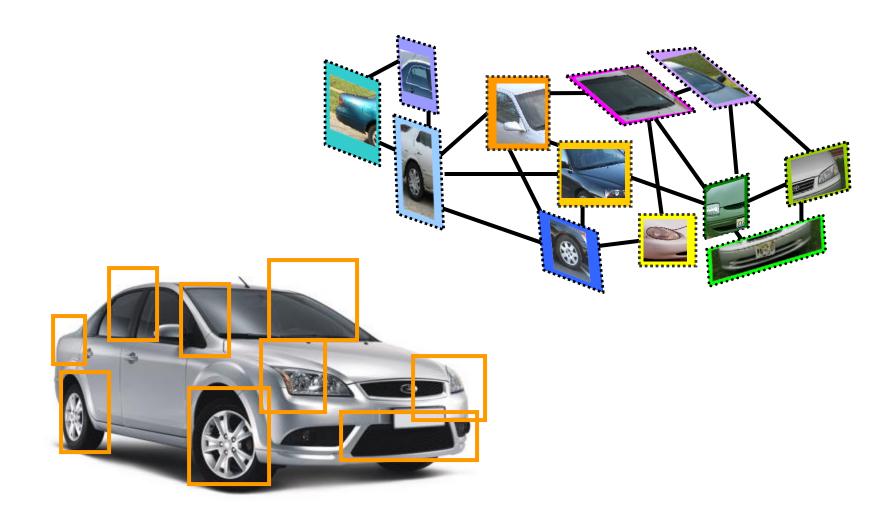
 Alignment: solving the transformation that makes two things match better



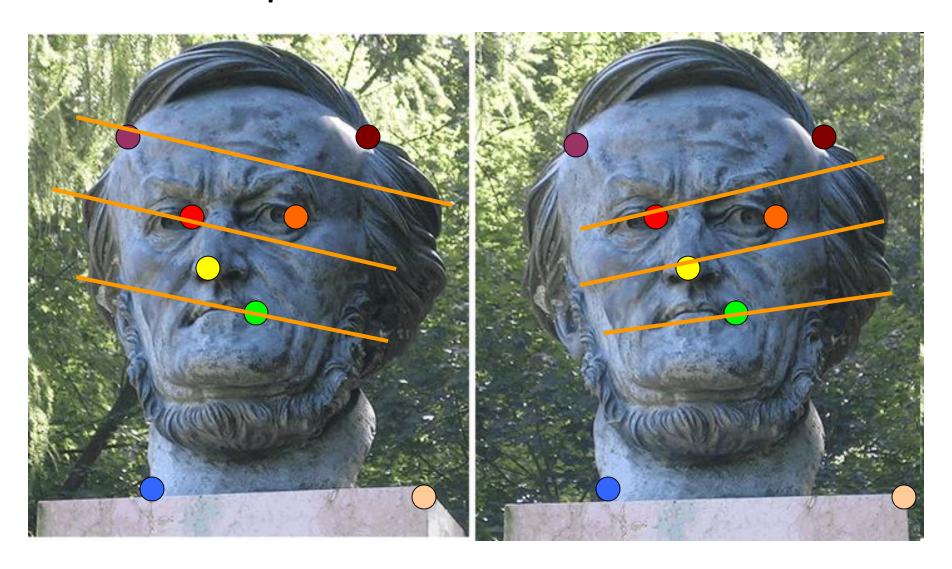
#### Example: fitting an 2D shape template



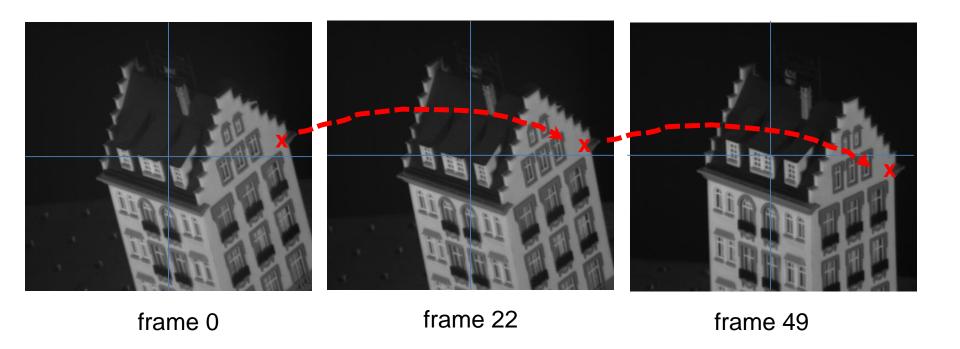
## Example: fitting a 3D object model



# Example: estimating "fundamental matrix" that corresponds two views



### Example: tracking points



Your problem 1 for HW 2!

#### HW 2

- Interest point detection and tracking
  - Detect trackable points
  - Track them across 50 frames
  - In HW 3, you will use these tracked points for structure from motion



frame 0



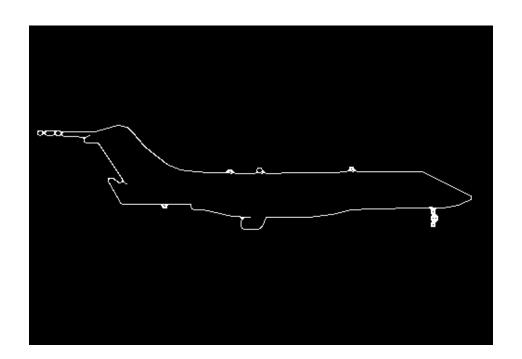
frame 22



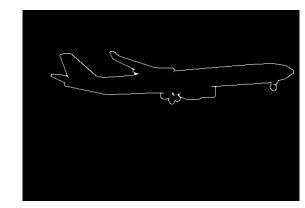
frame 49

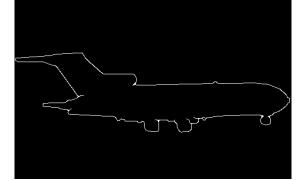
#### HW<sub>2</sub>

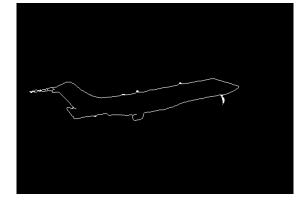
- Alignment of object edge images
  - Compute a transformation that aligns two edge maps





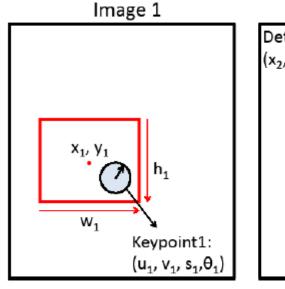


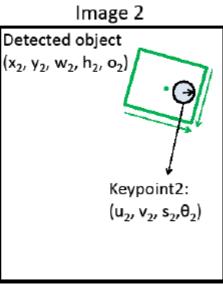




#### HW 2

- Initial steps of object alignment
  - Derive basic equations for interest-point based alignment



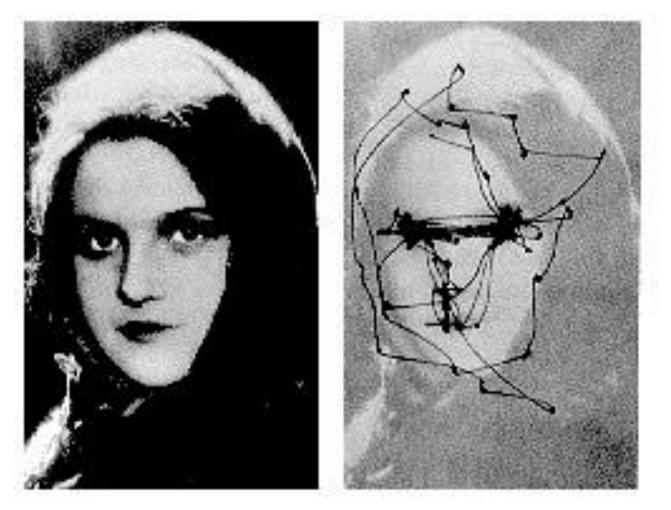


## This class: interest points

 Note: "interest points" = "keypoints", also sometimes called "features"

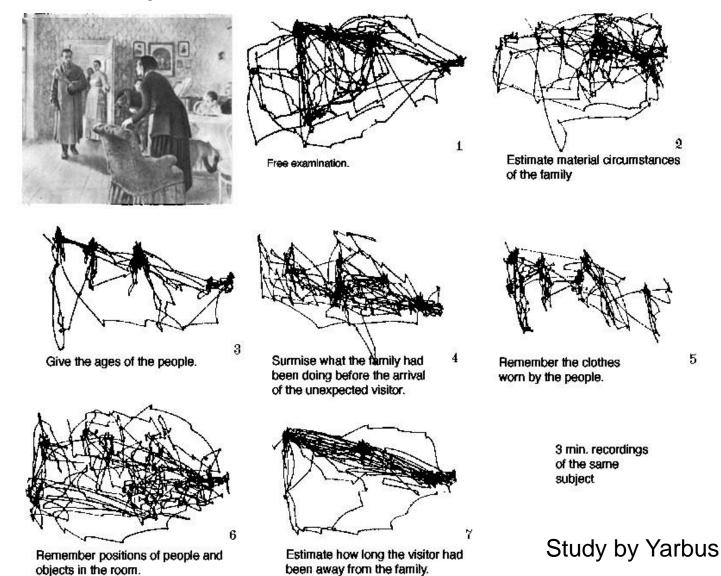
- Many applications
  - tracking: which points are good to track?
  - recognition: find patches likely to tell us something about object category
  - 3D reconstruction: find correspondences across different views

# Human eye movements



Yarbus eye tracking

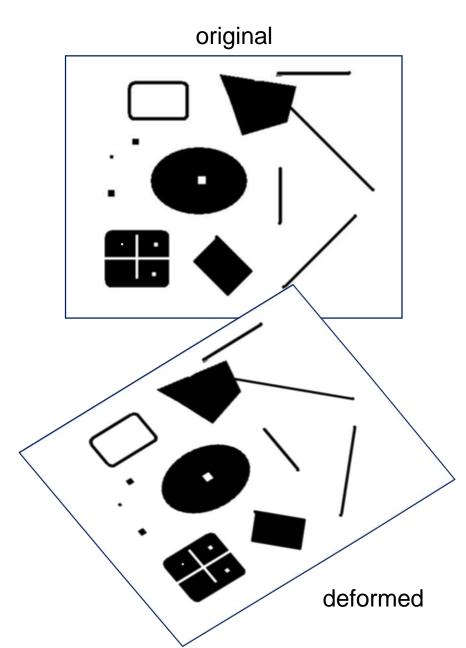
### Human eye movements



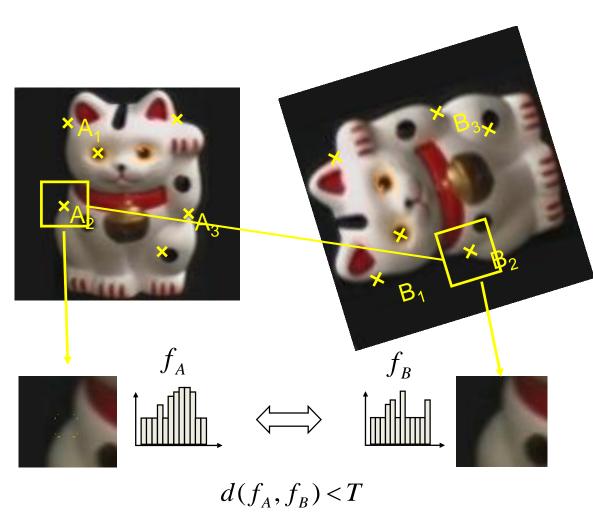
Change blindness: <a href="http://www.simonslab.com/videos.html">http://www.simonslab.com/videos.html</a>

### This class: interest points

- Suppose you have to click on some point, go away and come back after I deform the image, and click on the same points again.
  - Which points would you choose?



## Overview of Keypoint Matching



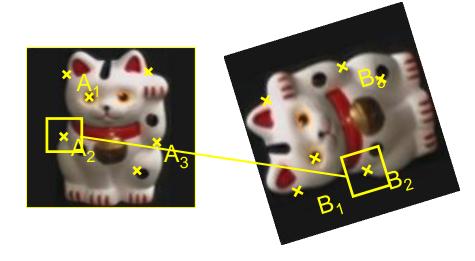
- 1. Find a set of distinctive key-points
- 2. Define a region around each keypoint
- 3. Extract and normalize the region content
- 4. Compute a local descriptor from the normalized region
- 5. Match local descriptors

## Goals for Keypoints



Detect points that are repeatable and distinctive

## Key trade-offs



#### Detection

More Repeatable

Robust detection
Precise localization

**More Points** 

Robust to occlusion
Works with less texture

#### Description

More Distinctive
Minimize wrong matches

More Flexible
Robust to expected variations

Maximize correct matches

## Choosing interest points

Where would you tell your friend to meet you?



## Choosing interest points

Where would you tell your friend to meet you?



## Many Existing Detectors Available

**Hessian & Harris** 

Laplacian, DoG

Harris-/Hessian-Laplace

Harris-/Hessian-Affine

EBR and IBR

**MSER** 

Salient Regions

Others...

[Beaudet '78], [Harris '88]

[Lindeberg '98], [Lowe 1999]

[Mikolajczyk & Schmid '01]

[Mikolajczyk & Schmid '04]

[Tuytelaars & Van Gool '04]

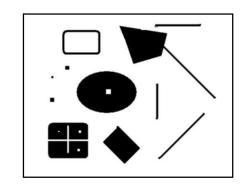
[Matas '02]

[Kadir & Brady '01]

### Harris Detector [Harris88]

Second moment matrix

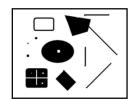
$$\mu(\sigma_I, \sigma_D) = g(\sigma_I) * \begin{bmatrix} I_x^2(\sigma_D) & I_x I_y(\sigma_D) \\ I_x I_y(\sigma_D) & I_y^2(\sigma_D) \end{bmatrix}$$



Intuition: Search for local neighborhoods where the image content has two main directions (eigenvectors).

#### Harris Detector [Harris88]

Second moment matrix



$$\mu(\sigma_{I},\sigma_{D}) = g(\sigma_{I}) * \begin{bmatrix} I_{x}^{2}(\sigma_{D}) & I_{x}I_{y}(\sigma_{D}) \\ I_{x}I_{y}(\sigma_{D}) & I_{y}^{2}(\sigma_{D}) \end{bmatrix}$$
 1. Image derivatives (optionally, blur first)





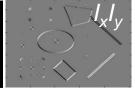
$$\det M = \lambda_1 \lambda_2$$

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

2. Square of derivatives



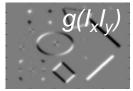




3. Gaussian filter  $g(\sigma_i)$ 







4. Cornerness function – both eigenvalues are strong

$$har = \det[\mu(\sigma_{I}, \sigma_{D})] - \alpha[\operatorname{trace}(\mu(\sigma_{I}, \sigma_{D}))^{2}] =$$

$$g(I_{x}^{2})g(I_{y}^{2}) - [g(I_{x}I_{y})]^{2} - \alpha[g(I_{x}^{2}) + g(I_{y}^{2})]^{2}$$

5. Non-maxima suppression



#### Harris Detector: Mathematics

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

1. Want large eigenvalues, and small ratio  $\frac{\lambda_1}{\lambda_2} < t$ 

2. We know

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

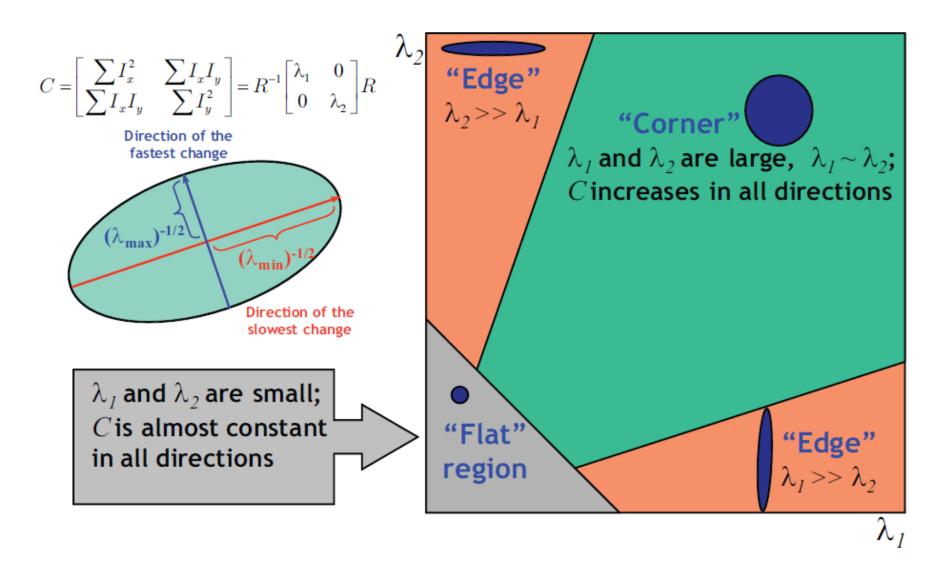
$$\operatorname{trace} M = \lambda_1 + \lambda_2$$

3. Leads to

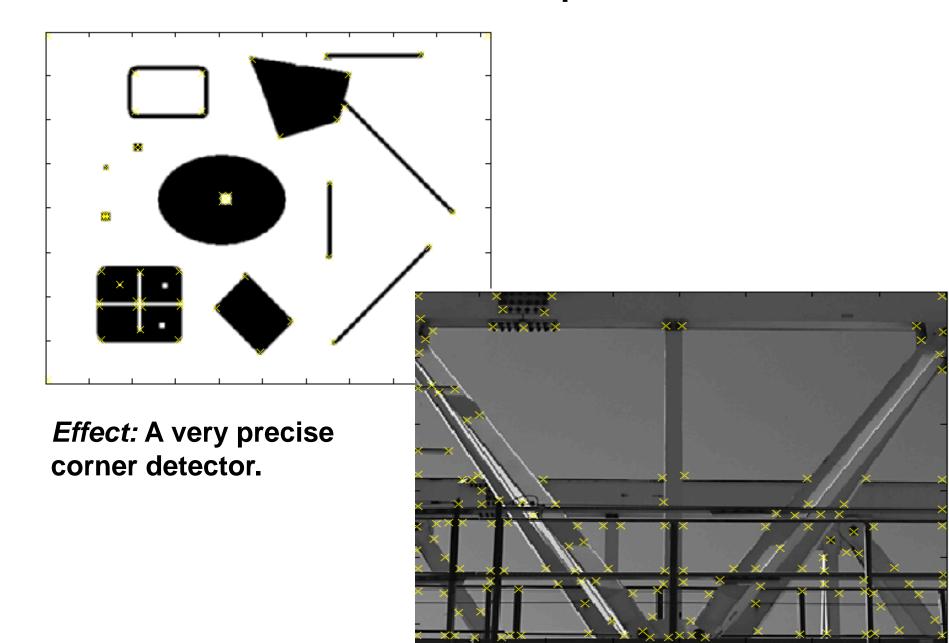
$$\det M - k \cdot \operatorname{trace}^{2}(M) > t$$
(k :empirical constant,  $k = 0.04-0.06$ )

Nice brief derivation on wikipedia

## **Explanation of Harris Criterion**



# Harris Detector – Responses [Harris88]



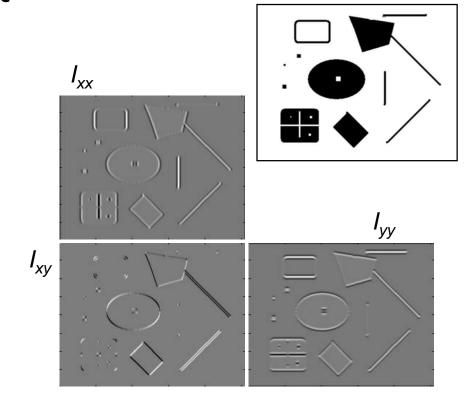
# Harris Detector - Responses [Harris88]



#### Hessian Detector [Beaudet78]

Hessian determinant

$$Hessian(I) = \begin{bmatrix} I_{xx} & I_{xy} \\ I_{xy} & I_{yy} \end{bmatrix}$$



Intuition: Search for strong curvature in two orthogonal directions

#### Hessian Detector [Beaudet78]

Hessian determinant

$$Hessian(x,\sigma) = \begin{bmatrix} I_{xx}(x,\sigma) & I_{xy}(x,\sigma) \\ I_{xy}(x,\sigma) & I_{yy}(x,\sigma) \end{bmatrix}$$

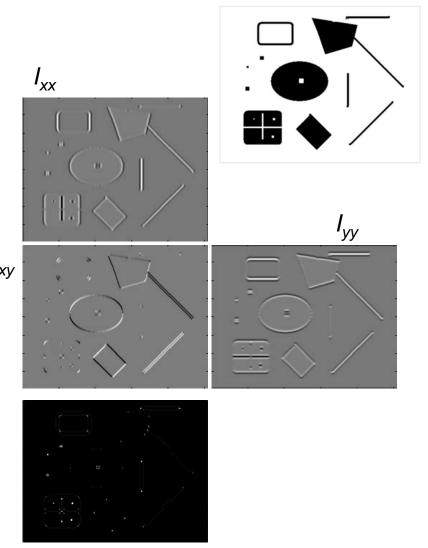
$$\det M = \lambda_1 \lambda_2$$
$$\operatorname{trace} M = \lambda_1 + \lambda_2$$

Find maxima of determinant

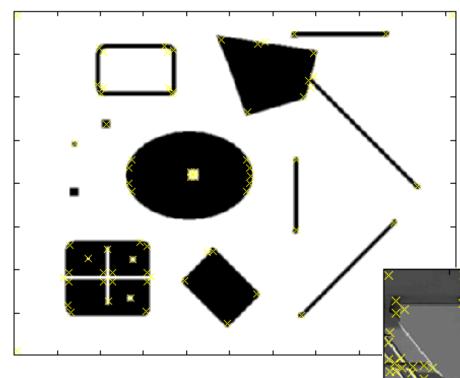
$$\det(Hessian(x)) = I_{xx}(x)I_{yy}(x) - I_{xy}^{2}(x)$$

In Matlab:

$$I_{xx} * I_{yy} - (I_{xy})^2$$



## Hessian Detector – Responses [Beaudet78]



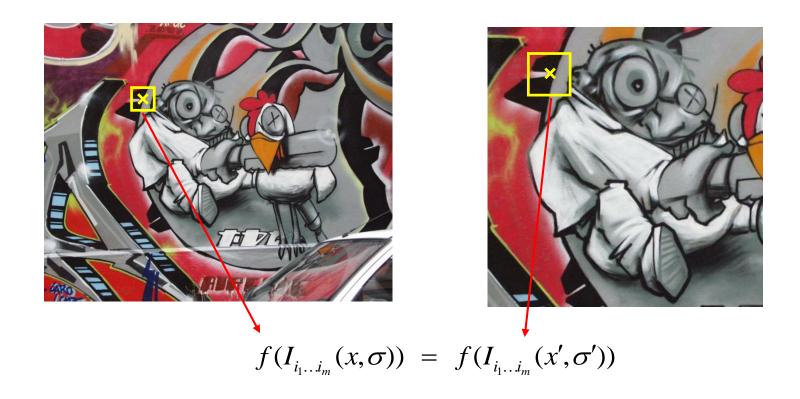
Effect: Responses mainly on corners and strongly textured areas.

# Hessian Detector – Responses [Beaudet78]



# So far: can localize in x-y, but not scale

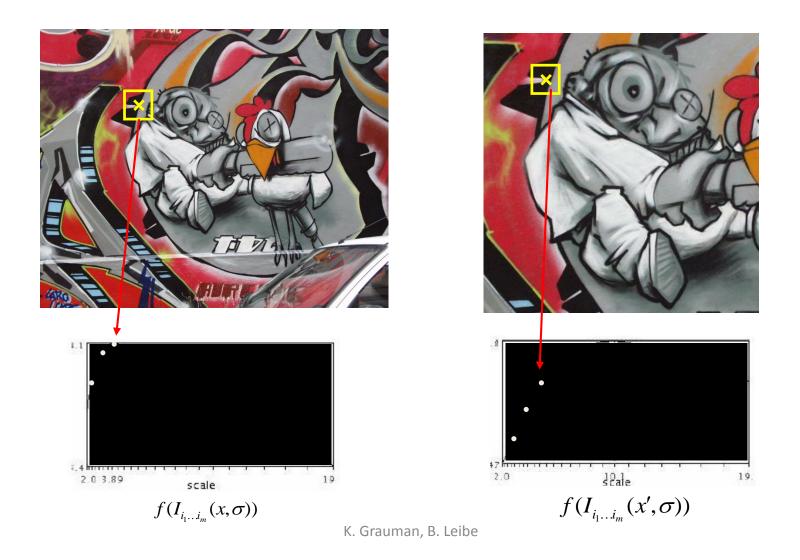


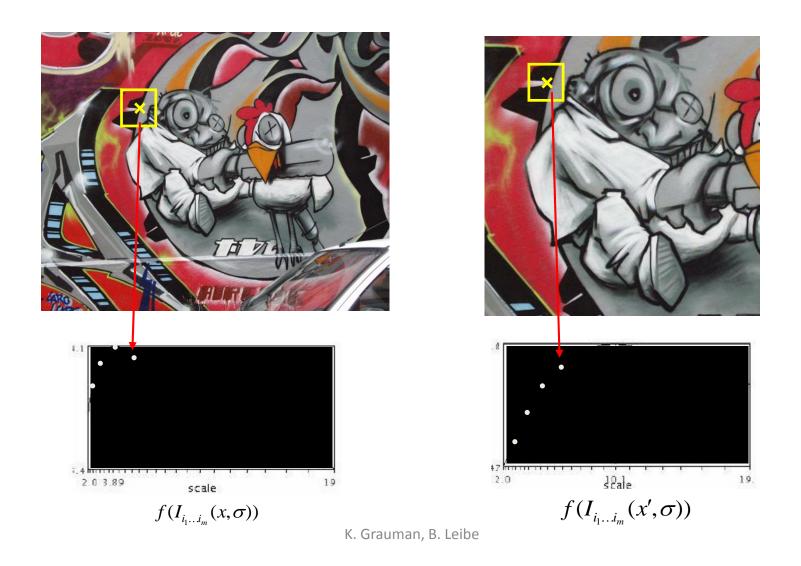


How to find corresponding patch sizes?



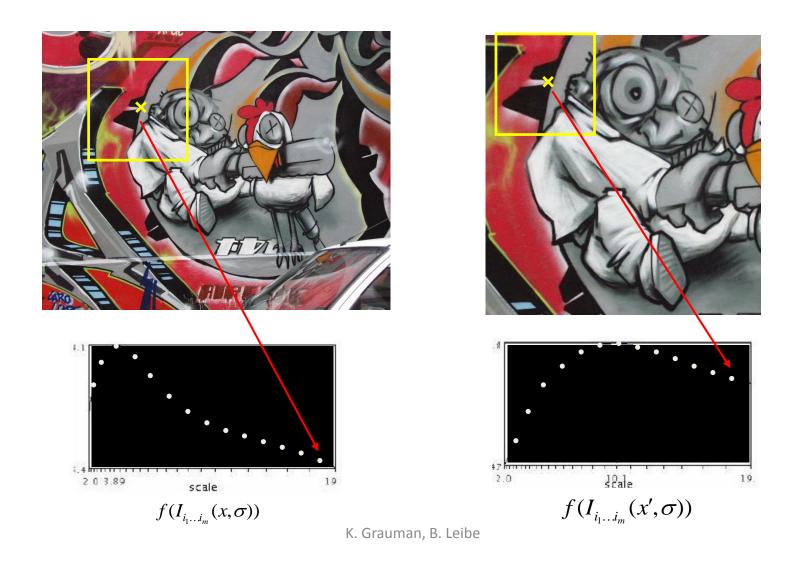






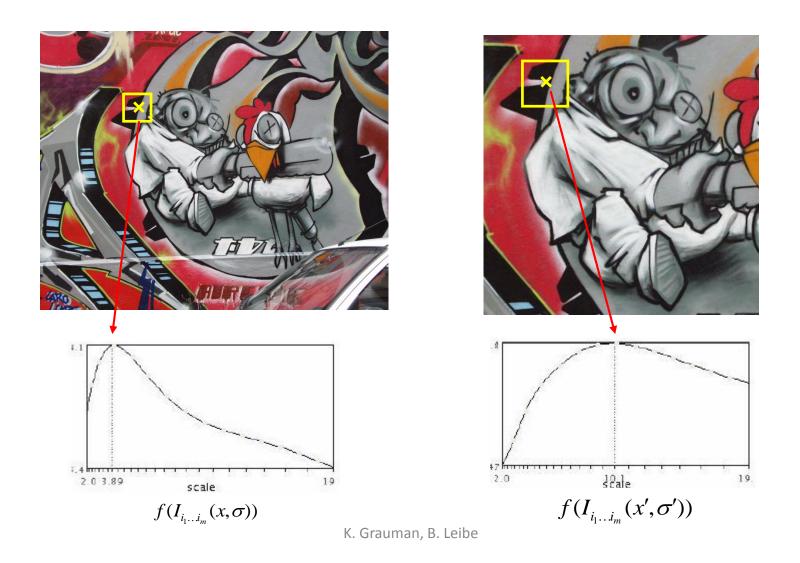
#### **Automatic Scale Selection**

Function responses for increasing scale (scale signature)



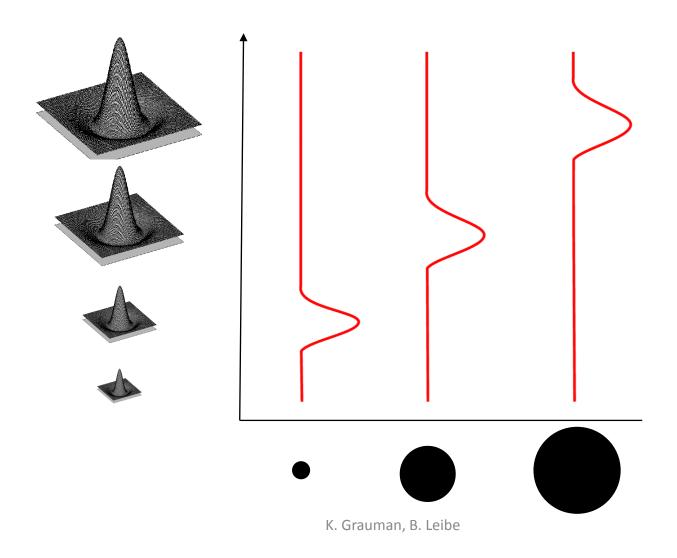
#### **Automatic Scale Selection**

Function responses for increasing scale (scale signature)

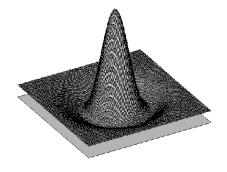


#### What Is A Useful Signature Function?

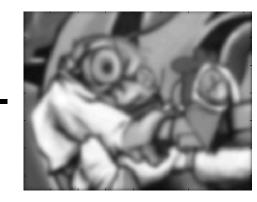
• Difference-of-Gaussian = "blob" detector



# Difference-of-Gaussian (DoG)





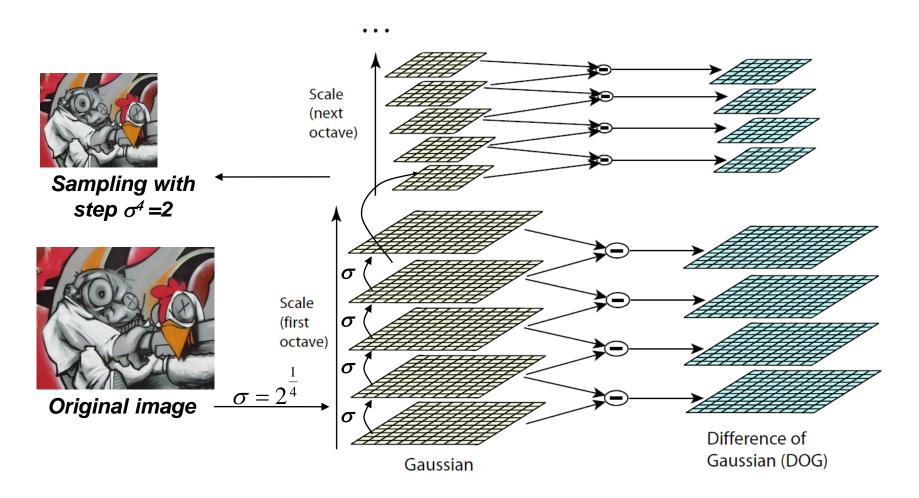




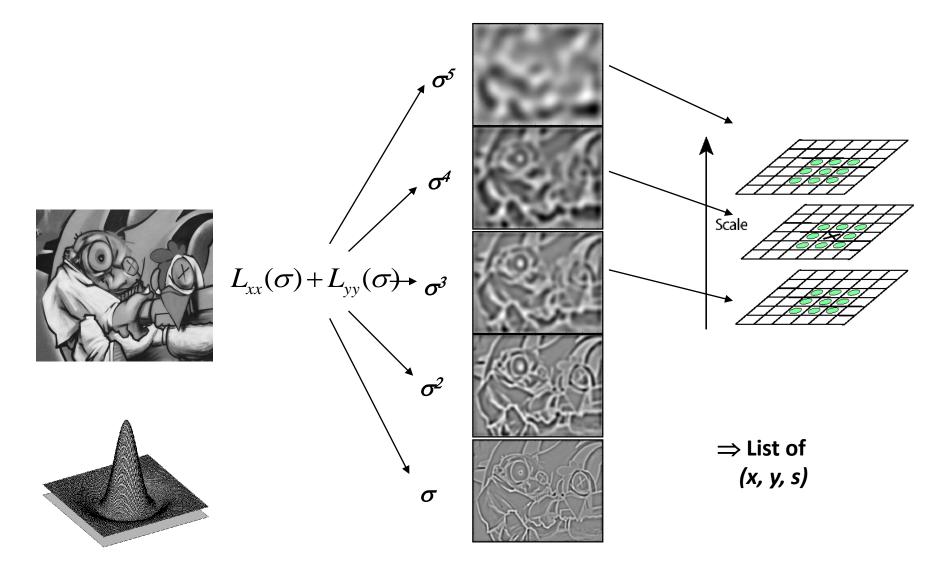


# DoG – Efficient Computation

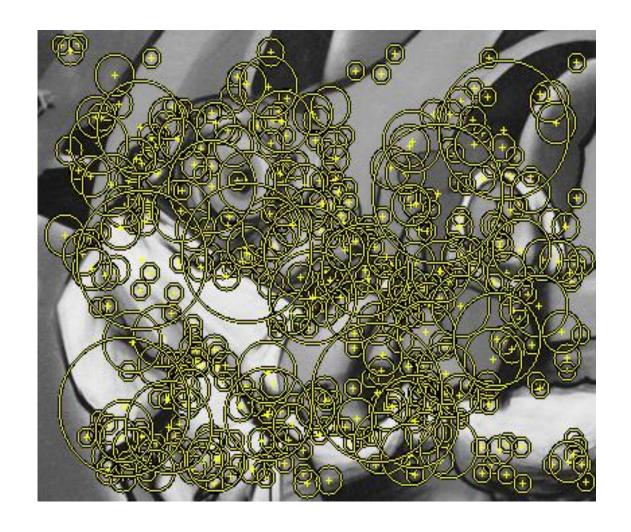
Computation in Gaussian scale pyramid



# Find local maxima in position-scale space of Difference-of-Gaussian



#### Results: Difference-of-Gaussian

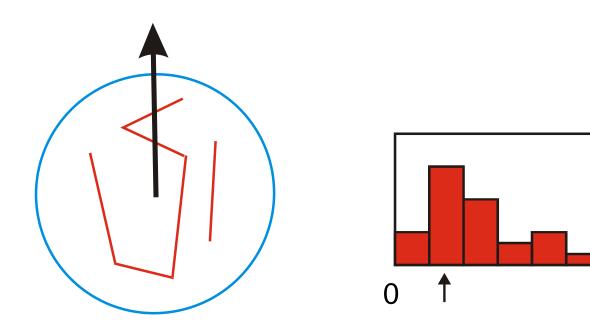


#### **Orientation Normalization**

Compute orientation histogram

[Lowe, SIFT, 1999]

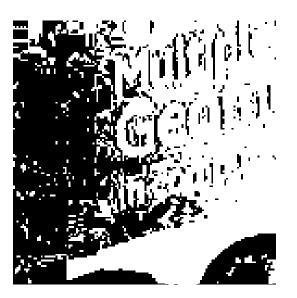
- Select dominant orientation
- Normalize: rotate to fixed orientation



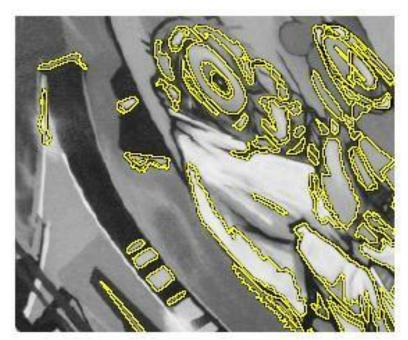
#### Maximally Stable Extremal Regions [Matas '02]

- Based on Watershed segmentation algorithm
- Select regions that stay stable over a large parameter range





# Example Results: MSER





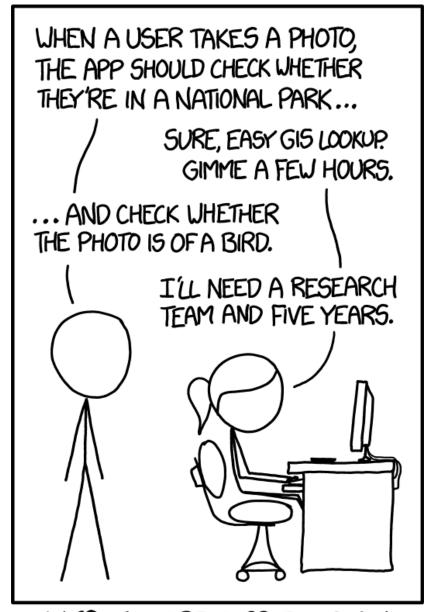




#### Available at a web site near you...

- 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

#### Interlude



IN CS, IT CAN BE HARD TO EXPLAIN THE DIFFERENCE BETWEEN THE EASY AND THE VIRTUALLY IMPOSSIBLE.

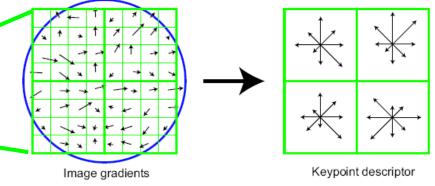
#### **Local Descriptors**

- The ideal descriptor should be
  - Robust
  - Distinctive
  - Compact
  - Efficient

- Most available descriptors focus on edge/gradient information
  - Capture texture information
  - Color rarely used

#### Local Descriptors: SIFT Descriptor





# Histogram of oriented gradients

- Captures important texture information
- Robust to small translations / affine deformations

[Lowe, ICCV 1999]

# Details of Lowe's SIFT algorithm

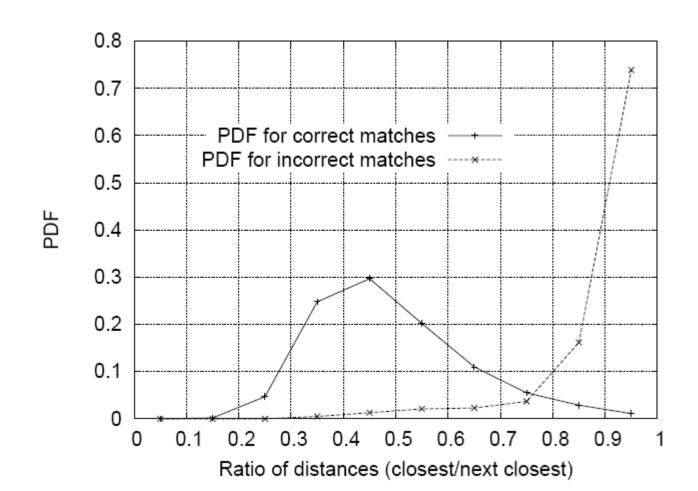
- Run DoG detector
  - Find maxima in location/scale space
  - Remove edge points
- Find all major orientations

- $\mathbf{H} = \begin{bmatrix} D_{xx} & D_{xy} \\ D_{xy} & D_{yy} \end{bmatrix}$  $\mathbf{Tr}(\mathbf{H})^2 \qquad (r+1)^2$
- $\frac{\mathrm{Tr}(\mathbf{H})^2}{\mathrm{Det}(\mathbf{H})} < \frac{(r+1)^2}{r}$

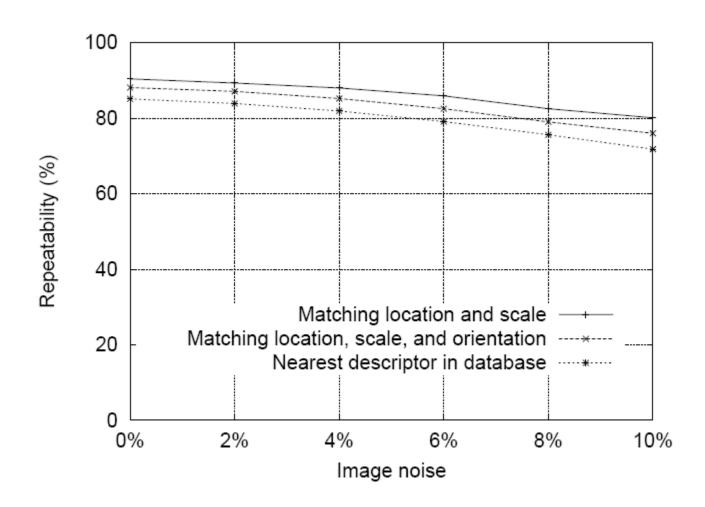
- Bin orientations into 36 bin histogram
  - Weight by gradient magnitude
  - Weight by distance to center (Gaussian-weighted mean)
- Return orientations within 0.8 of peak
  - Use parabola for better orientation fit
- For each (x,y,scale,orientation), create descriptor:
  - Sample 16x16 gradient mag. and rel. orientation
  - Bin 4x4 samples into 4x4 histograms
  - Threshold values to max of 0.2, divide by L2 norm
  - Final descriptor: 4x4x8 normalized histograms

#### Matching SIFT Descriptors

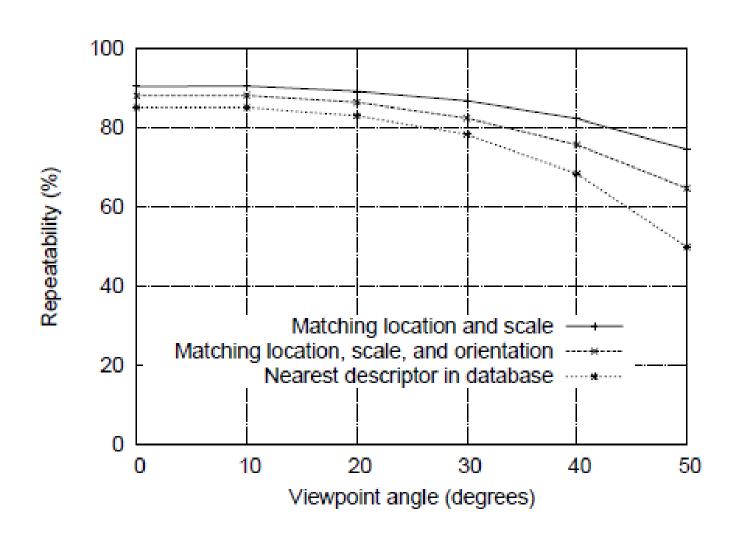
- Nearest neighbor (Euclidean distance)
- Threshold ratio of nearest to 2<sup>nd</sup> nearest descriptor



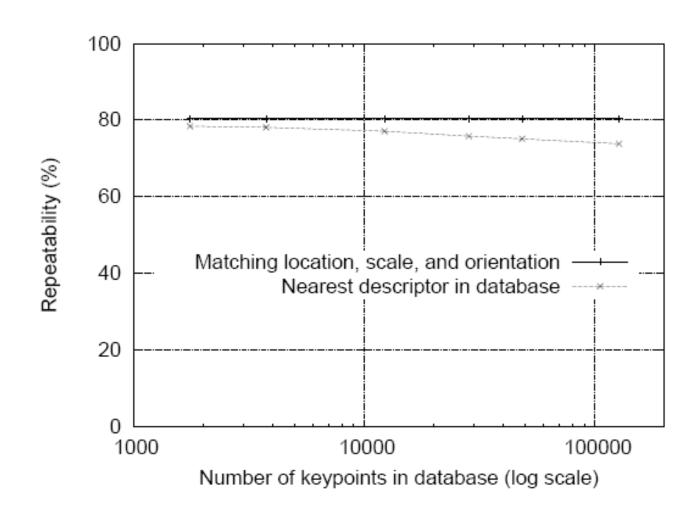
# SIFT Repeatability



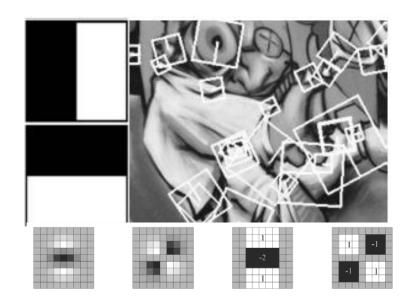
#### SIFT Repeatability



# SIFT Repeatability



#### Local Descriptors: SURF



#### Fast approximation of SIFT idea

Efficient computation by 2D box filters & integral images

 $\Rightarrow$  6 times faster than SIFT

**Equivalent quality for object identification** 

Many other efficient descriptors are also available

#### **GPU** implementation available

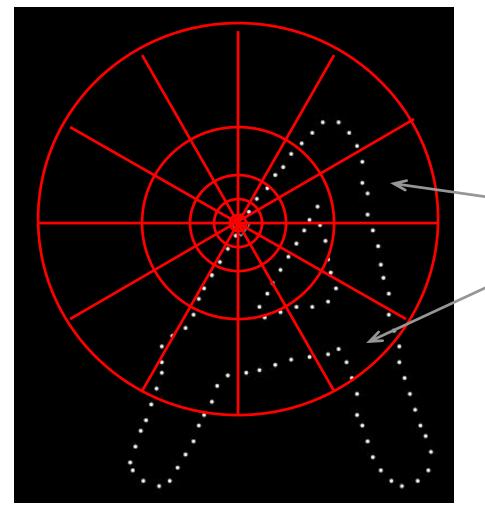
Feature extraction @ 200Hz (detector + descriptor, 640×480 img) http://www.vision.ee.ethz.ch/~surf

#### Local Descriptors: ORB

- Many similarities to SIFT/SURF
- Designed for efficiency and robustness to orientation
- Not designed for scale robustness
- Used for tracking and long-range matching in ORB-SLAM

http://www.willowgarage.com/sites/default/files/orb\_final.pdf (ICCV 2011) http://webdiis.unizar.es/~raulmur/orbslam/

#### **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.

#### Choosing a detector

- What do you want it for?
  - Precise localization in x-y: Harris
  - Good localization in scale: Difference of Gaussian
  - Flexible region shape: MSER
- Best choice often application dependent
  - Harris-/Hessian-Laplace/DoG work well for many natural categories
  - MSER works well for buildings and printed things
- Why choose?
  - Get more points with more detectors
- There have been extensive evaluations/comparisons
  - [Mikolajczyk et al., IJCV'05, PAMI'05]
  - All detectors/descriptors shown here work well

# Comparison of Keypoint Detectors

Table 7.1 Overview of feature detectors.

				Rotation	Scale	Affine		Localization		
Feature Detector	Corner	$_{\mathrm{Blob}}$	Region	invariant	invariant	invariant	Repeatability	accuracy	Robustness	Efficiency
Harris	√			√			+++	+++	+++	++
Hessian		$\checkmark$		$\checkmark$			++	++	++	+
SUSAN				√			++	++	++	+++
Harris-Laplace	$\checkmark$	(√)		√	√		+++	+++	++	+
Hessian-Laplace	(√)	$\checkmark$		$\checkmark$	$\checkmark$		+++	+++	+++	+
DoG	(√)	$\checkmark$		$\checkmark$	$\checkmark$		++	++	++	++
SURF	(√)	$\checkmark$		√	$\checkmark$		++	++	++	+++
Harris-Affine	√	(√)		<b>√</b>	√	<b>√</b>	+++	+++	++	++
Hessian-Affine	(√)	$\checkmark$		$\checkmark$	$\checkmark$	$\checkmark$	+++	+++	+++	++
Salient Regions	(√)	$\checkmark$		$\checkmark$	$\checkmark$	(√)	+	+	++	+
Edge-based	$\checkmark$			√	$\checkmark$	$\checkmark$	+++	+++	+	+
MSER				√	√	<b>√</b>	+++	+++	++	+++
Intensity-based			$\checkmark$	$\checkmark$	$\checkmark$	$\checkmark$	++	++	++	++
Superpixels			$\checkmark$	$\checkmark$	(√)	()	+	+	+	+

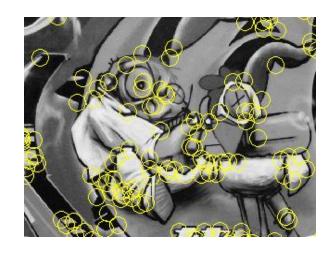
# Choosing a descriptor

 Consider efficiency of computation and lookup and robustness to orientation/scale

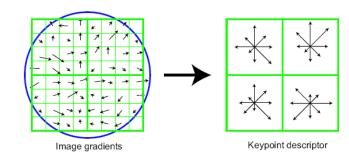
 For object instance recognition, SfM, or stitching, SIFT or variant is a good choice

#### Things to remember

- Keypoint detection: repeatable and distinctive
  - Corners, blobs, stable regions
  - Harris, DoG



- Descriptors: robust and selective
  - spatial histograms of orientation
  - SIFT



#### Next time

Feature tracking