CS688: Web-Scale Image Search

Classical Keypoint Localization

Sung-Eui Yoon
(윤성의)

Course URL:
http://sgvr.kaist.ac.kr/~sungeui/IR

KAIST

Class Objective (Ch. 2 of My Draft)

- Get to know related conferences
- Understand locally invariant features
  - Key point localization
  - Harris detector

- At Last class:
  - Student activities including paper presentations and mid- and final-term project presentations
  - Grading policy

Homework for Every Class

- **Go over the next lecture slides**
- **Come up with one question on what we have discussed today**
  - **1 for typical questions (that were answered in the class)**
  - **2 for questions with thoughts or that surprised me**
- **Write questions 3 times before the mid-term**
  - **Multiple questions in one time will be counted as one time**
- **Common questions are addressed at my draft**
  - **Some of questions will be discussed in the class**
- **If you want to know the answer of your question, ask me or TA on person**

Homework for Every Class

- Go over recent papers on image search
  - High quality papers: Papers published at the top-tier conf.; e.g., CVPR, ICCV, ECCV, ACM ICMR, NeuroIPS, ICML, MM, SIGGRAPH
  - Recent publication: papers published since 2016
  - Find and browse two papers, and submit two summaries before every beginning of the Tue. Class
  - Online submission is possible
- Think about possible team members
- Too late if you think them later..

Content-Based Image Retrieval (CBIR)

- **Identify similar images given a user-specified image or other types of inputs**

Extract image descriptors (e.g., CNN or SIFT)

Input

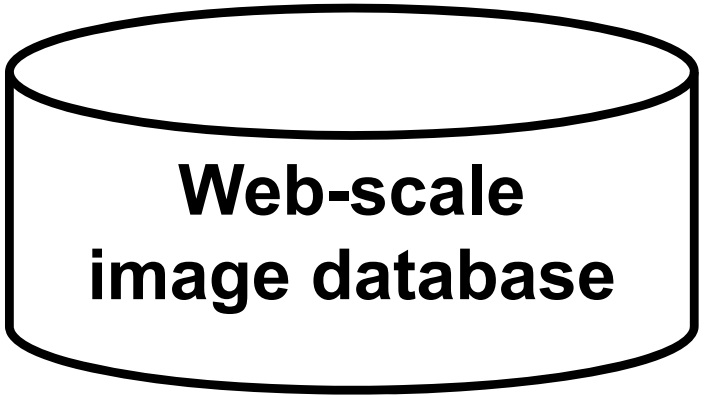

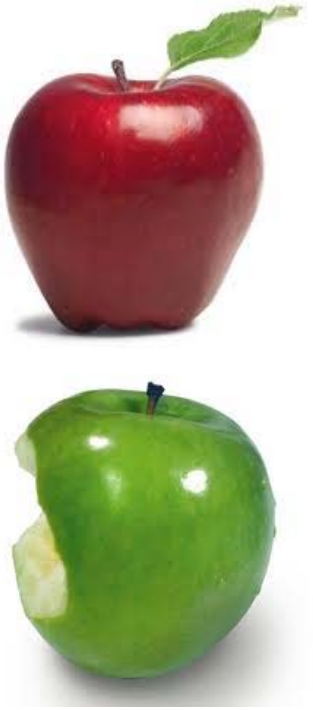

Output

KAIST

Key Components of Image Search

- Image representations
- Indexing algorithms
- Matching methods
- Classification, Localization, etc.
  - Can improve image search or improve these techniques utilizing image search

Image Representations

- **SIFT, GIST, CNN, etc.**
  - **Invariant to different transformations**

Image Retrieval

- **At pre-processing, build a database for efficient retrieval at runtime**

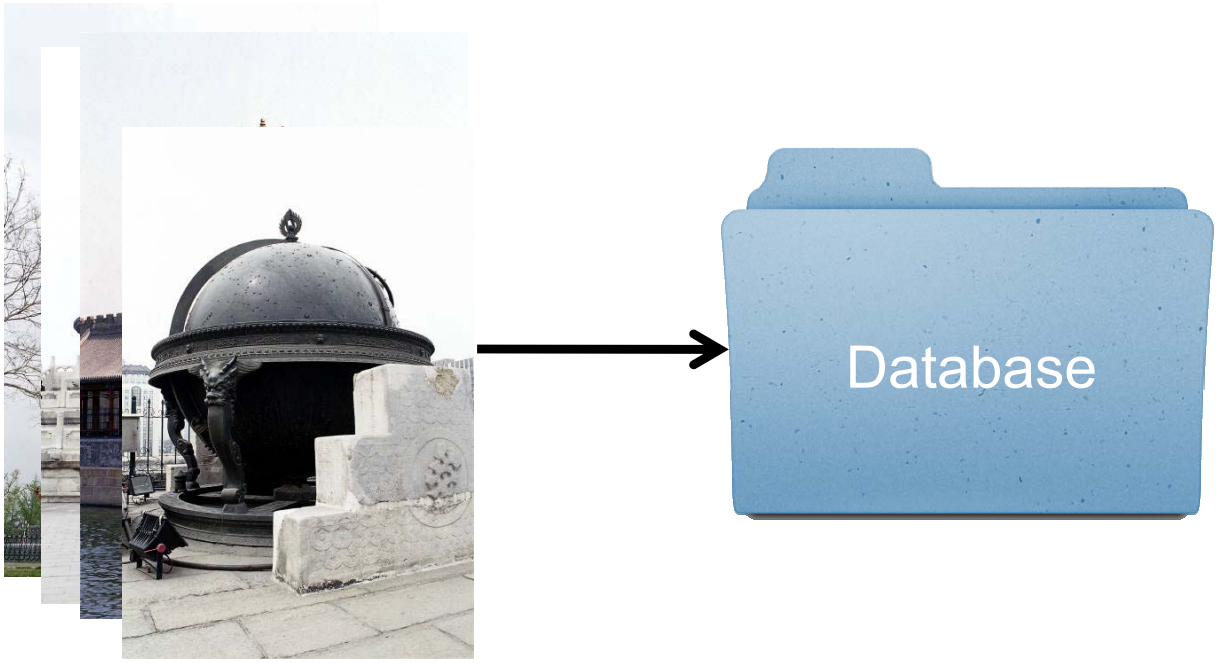

KAIST

Image Retrieval

- **At pre-processing, build a database for efficient retrieval at runtime**

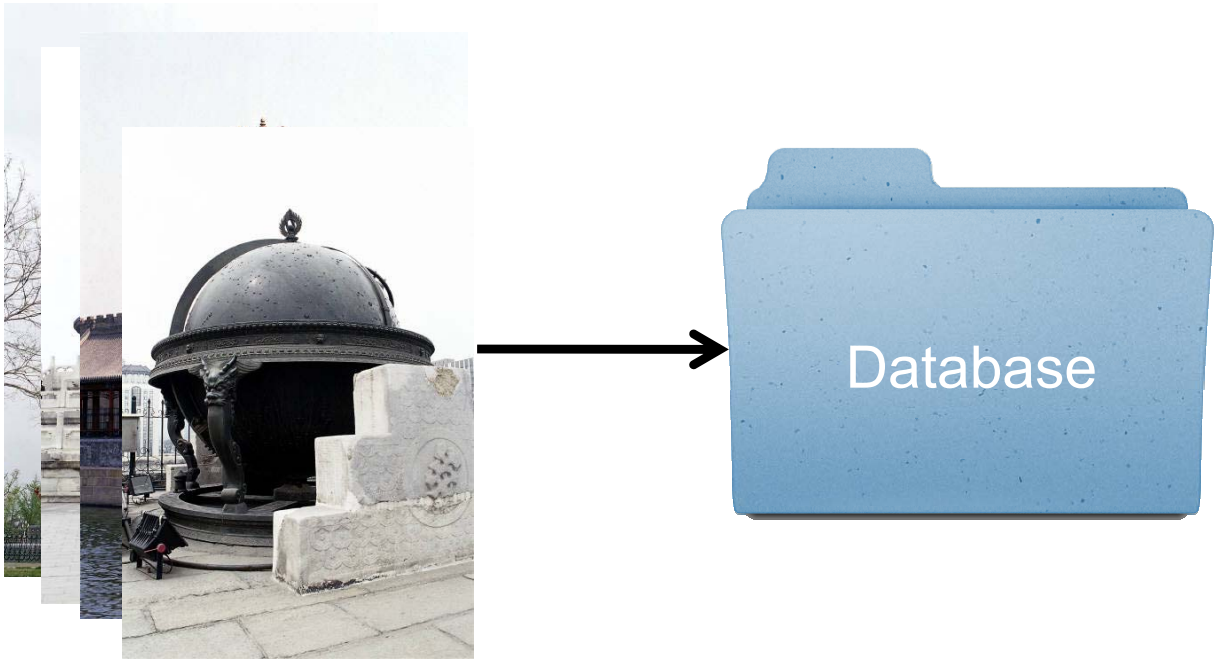

Index schemes: vocabulary trees, hashing, and inverted files

KAIST

Image Retrieval: Runtime Procedure

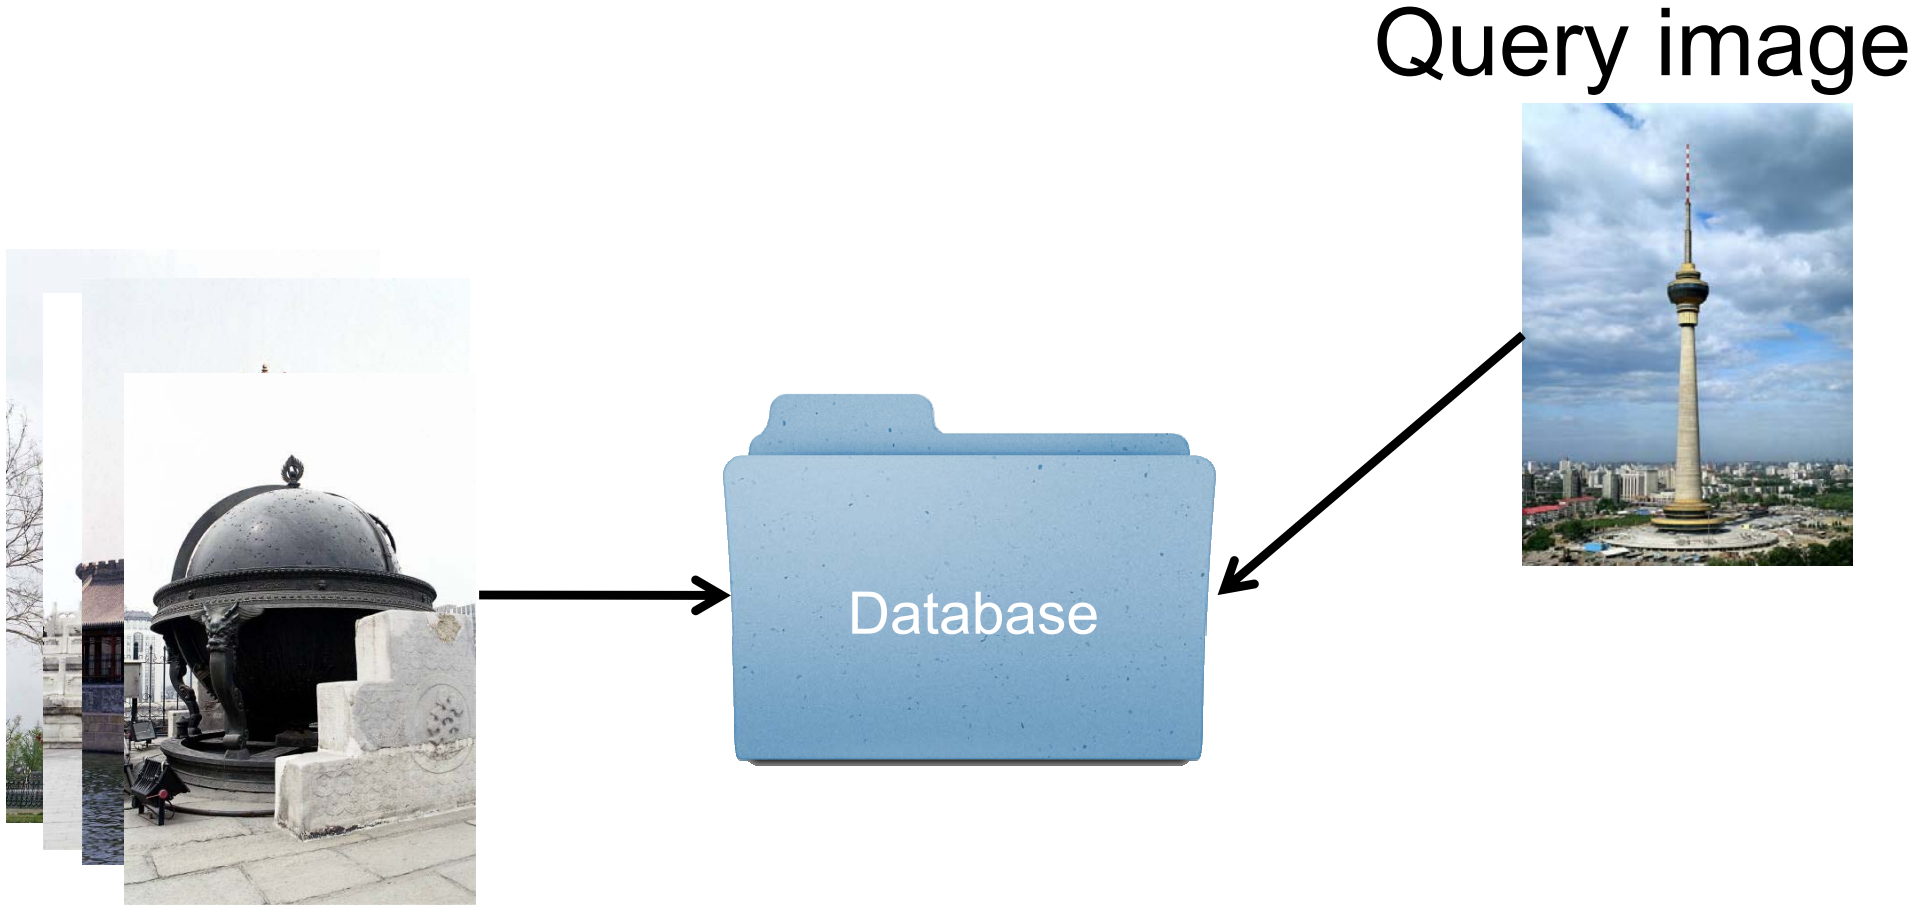

KAIST

Image Retrieval: Runtime Procedure

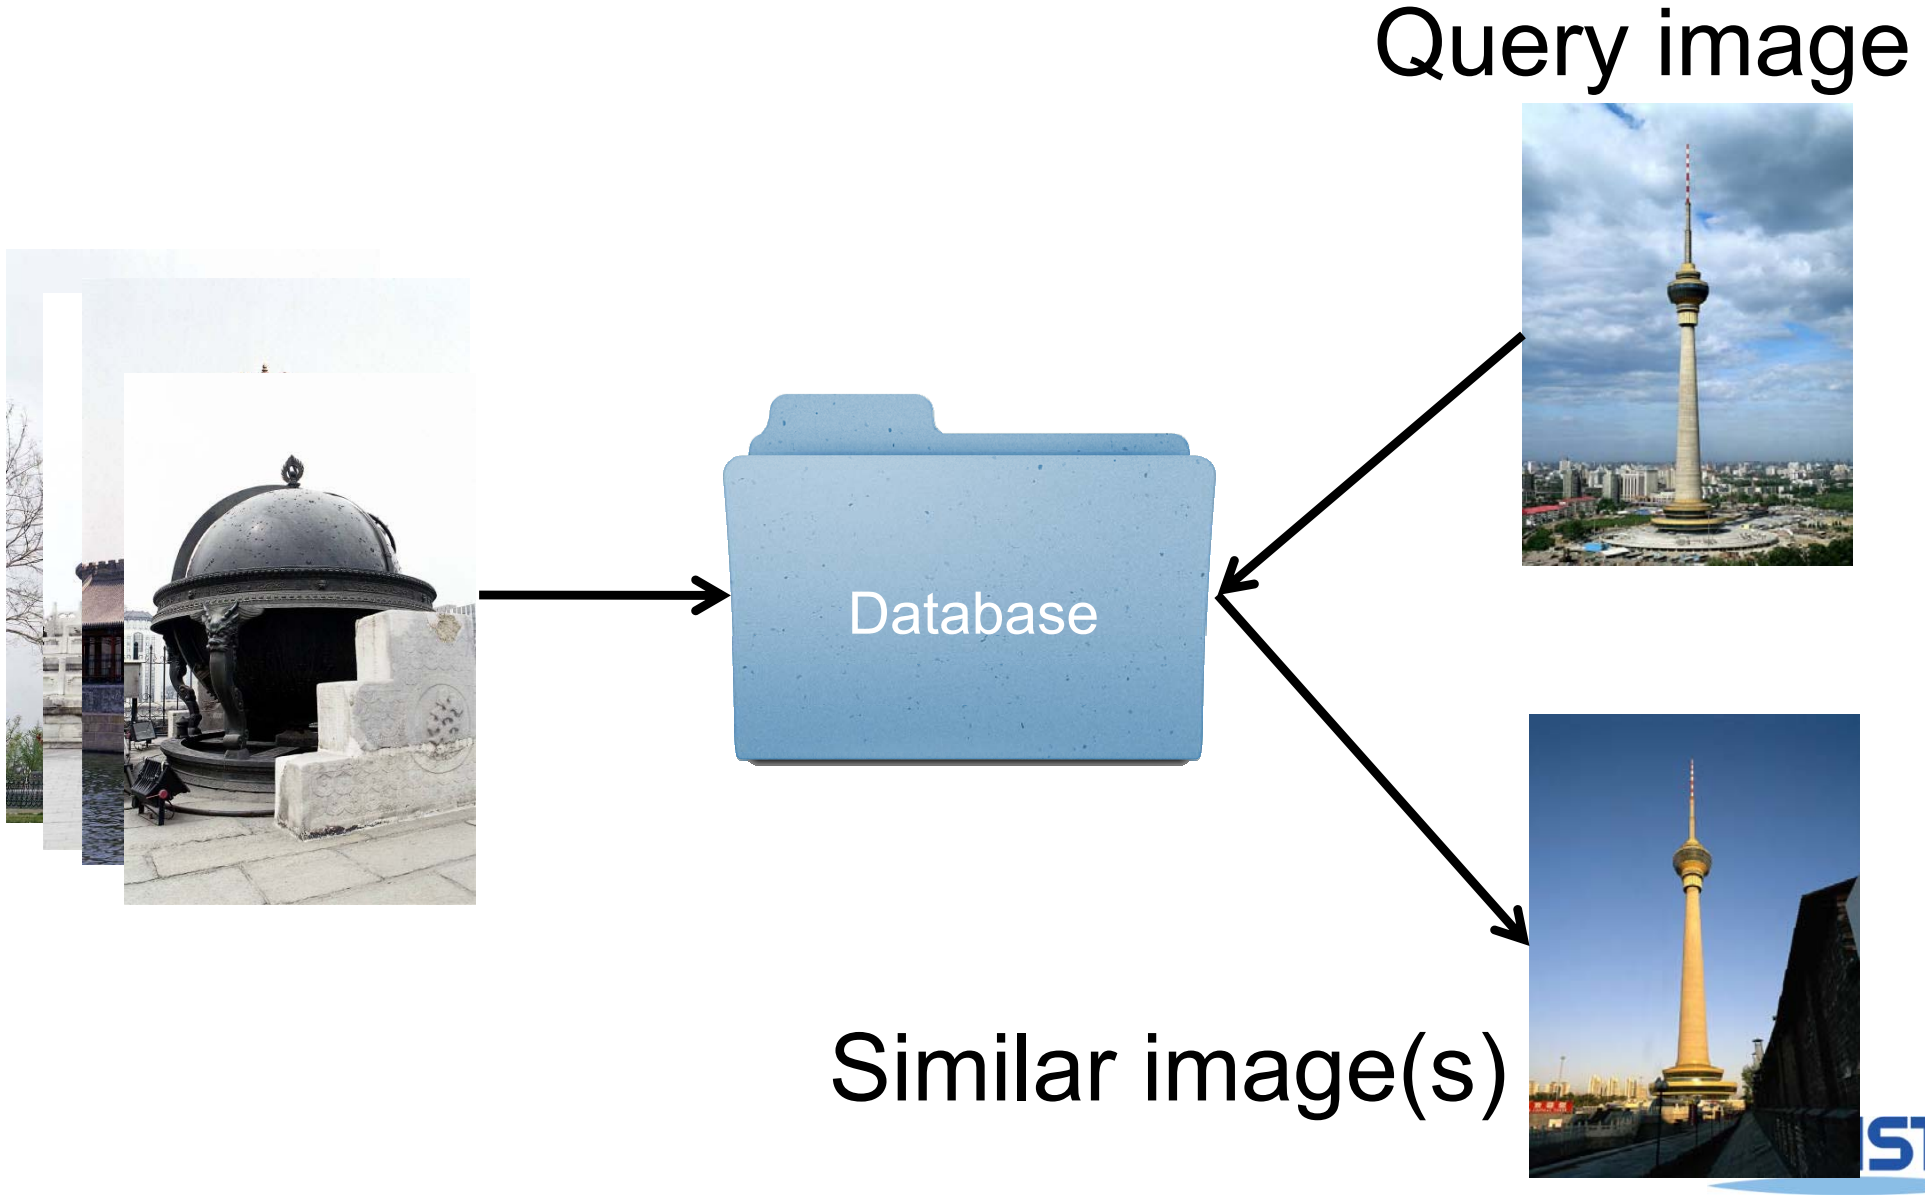

Post-Processing

Query image

Shortlist (e.g., 100 images)



..

Re-ranking
(spatial verification)



..

Motivation for Image Descriptors

- **RGB images are not robust for various changes (e.g., geometric and photometric transformations)**

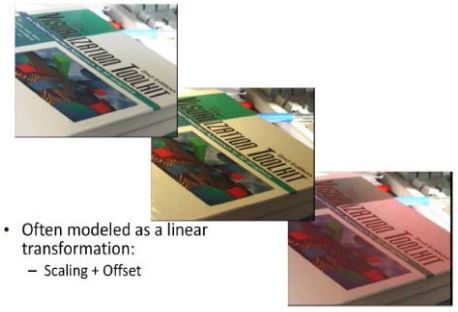

- **Image descriptors are classified into:**
  - **Global feature encoding the overall context**
  - **Local features encoding different parts of objects**
- **Global and local features are useful, but we focus on local features for now**
  - **More robust to various changes**

KAIST

Challenges: viewpoint variation

Michelangelo 1475-1564

Challenges: illumination

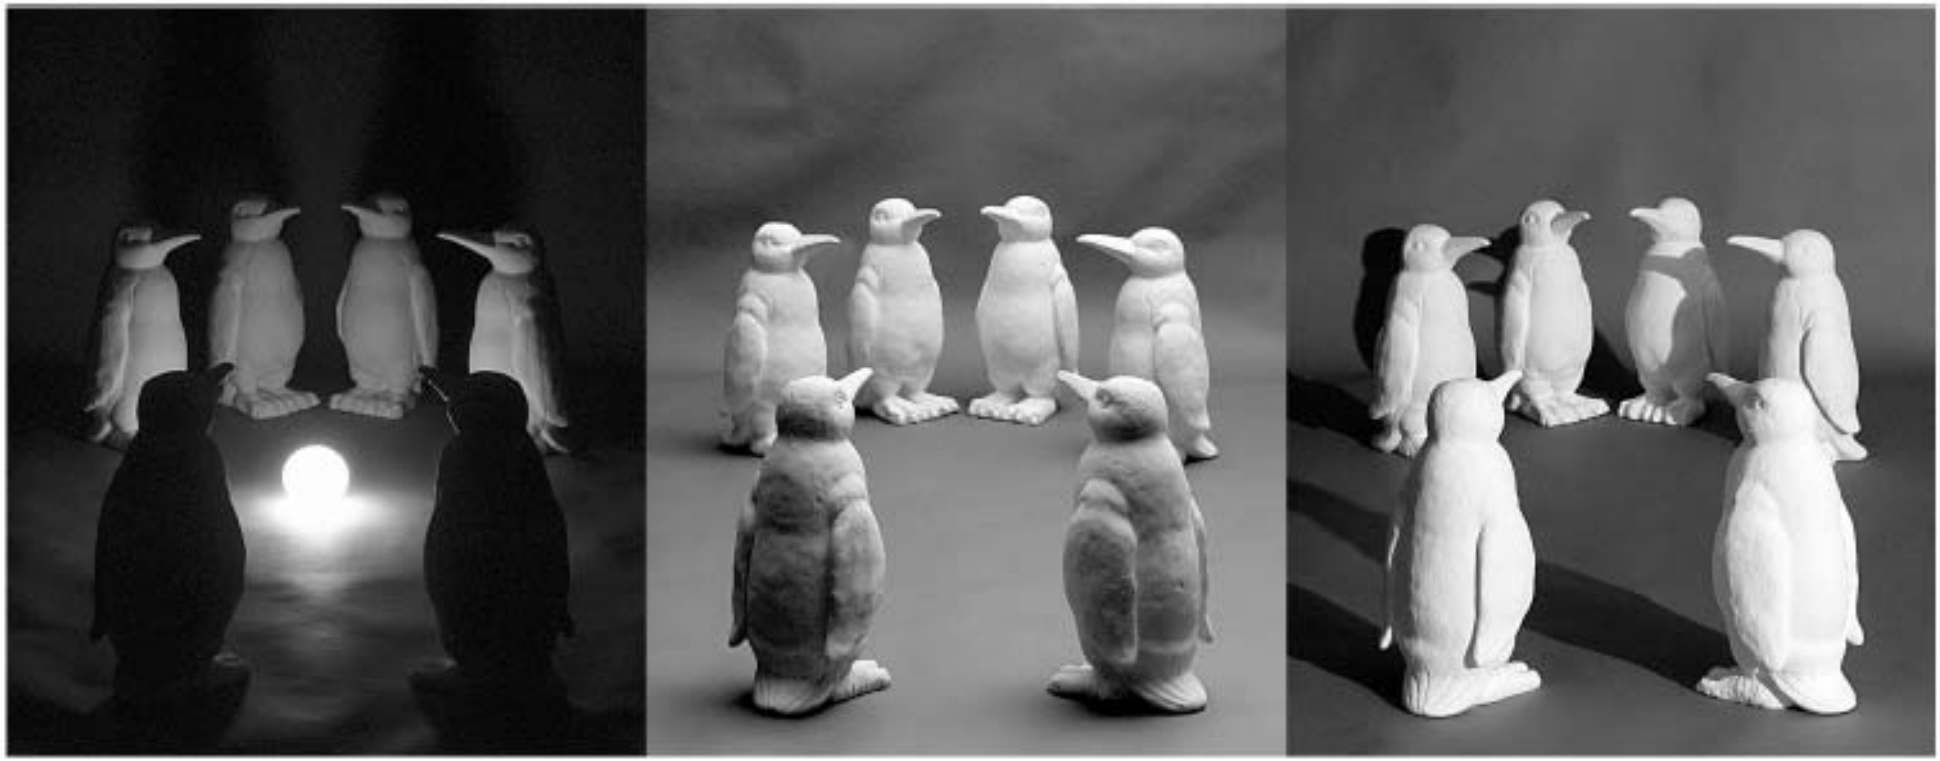

image credit: J. Koenderink

Challenges: scale

Challenges: deformation

Challenges: occlusion

Magritte, 1957

Challenges: background clutter

Kilmeny Niland. 1995

Challenges: intra-class variation

Application: Image Matching

by Diva Sian

by swashford

Slide credit: Steve Seitz

Harder Case

by Diva Sian

by scgbt

Slide credit: Steve Seitz

Application: Image Stitching

Slide credit: Darya Frolova, Denis Simakov

Application: Image Stitching

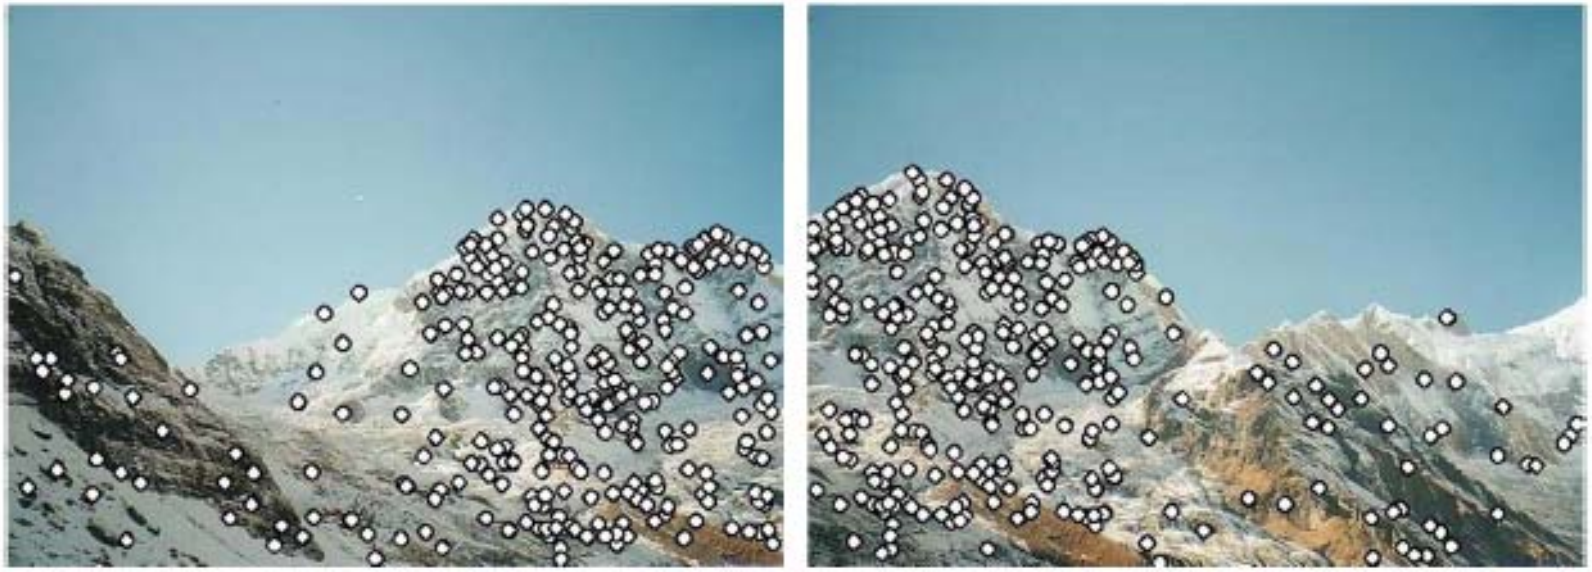

- Procedure:
  - Detect feature points in both images

Slide credit: Darya Frolova, Denis Simakov

Application: Image Stitching

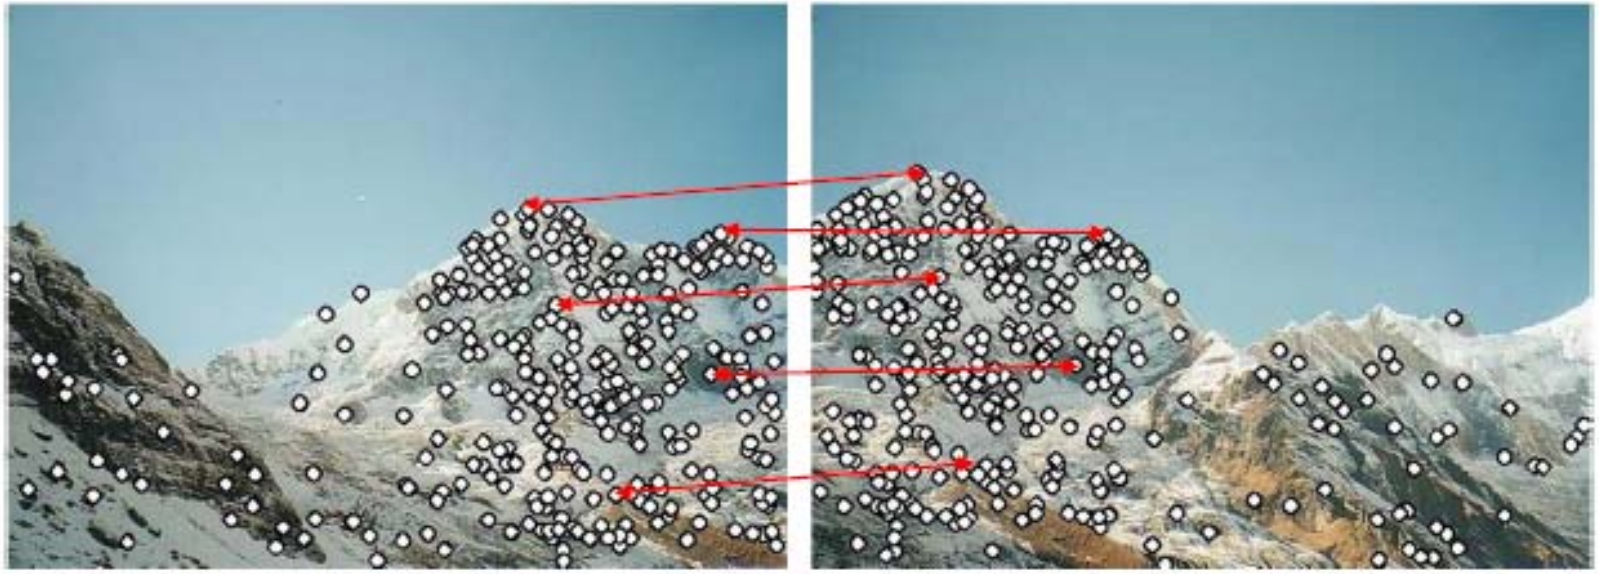

- Procedure:
  - Detect feature points in both images
  - Find corresponding pairs

Slide credit: Darya Frolova, Denis Simakov

Application: Image Stitching

- Procedure:
  - Detect feature points in both images
  - Find corresponding pairs
  - Use these pairs to align the images

Slide credit: Darya Frolova, Denis Simakov

Common Requirements

- Problem 1:
  - Detect the same point *independently* in both images

No chance to match!

This lecture

We need a repeatable detector!

Slide credit: Darya Frolova, Denis Simakov

Common Requirements

- Problem 1:
  - Detect the same point *independently* in both images
- Problem 2:
  - For each point correctly recognize the corresponding one

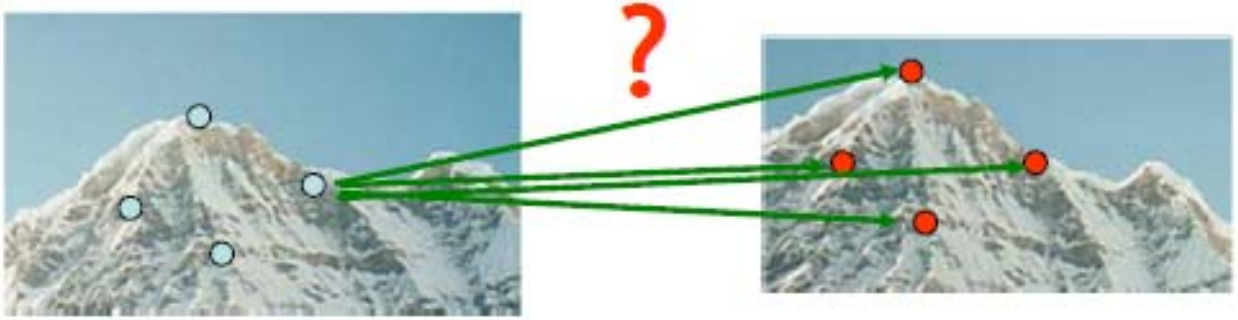

Next lecture

We need a reliable and distinctive descriptor!

Slide credit: Darya Frolova, Denis Simakov

Two Different Directions

- **Classical approaches**
  - **Manually designed in image processing and computer vision fields**
- **Deep learning approaches**
  - **Learned approaches, but are inspired by many prior (manually crafted) approaches**
- **In this class**
  - **We first talk about the classical approaches, followed by deep learning approaches**

Many Existing Detectors Available

- Hessian & Harris [Beaudet '78], [Harris '88]
- Laplacian, DoG [Lindeberg '98], [Lowe '99]
- Harris-/Hessian-Laplace [Mikolajczyk & Schmid '01]
- Harris-/Hessian-Affine [Mikolajczyk & Schmid '04]
- EBR and IBR [Tuytelaars & Van Gool '04]
- MSER [Matas '02]
- Salient Regions [Kadir & Brady '01]
- Others…

- *Those detectors have become a basic building block for many recent applications in Computer Vision.*

Slide credit: Bastian Leibe

Keypoint Localization

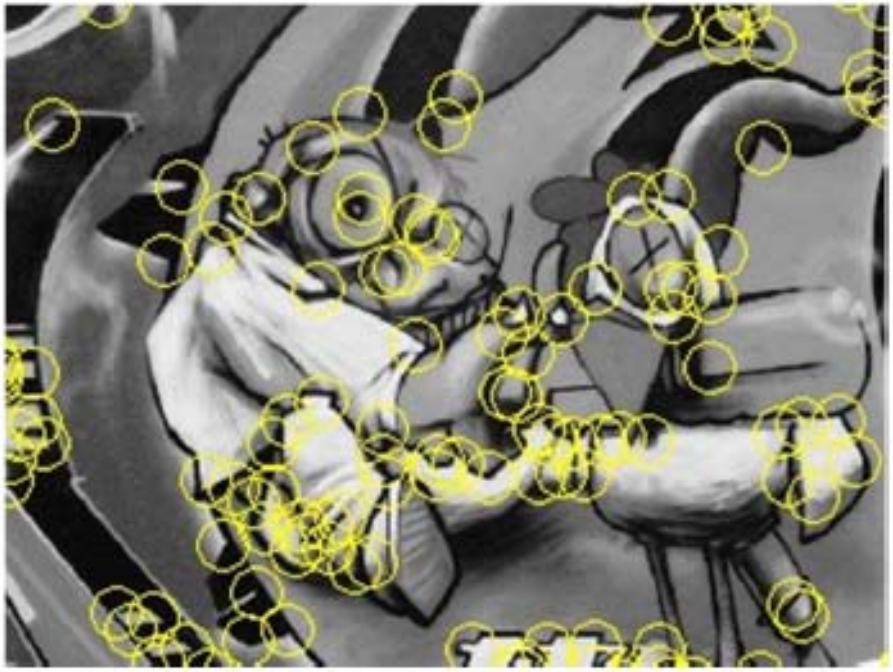

- Goals:
  - Repeatable detection
  - Precise localization
  - Interesting content

  $\Rightarrow$ *Look for two-dimensional signal changes*

Slide credit: Bastian Leibe

Finding Corners

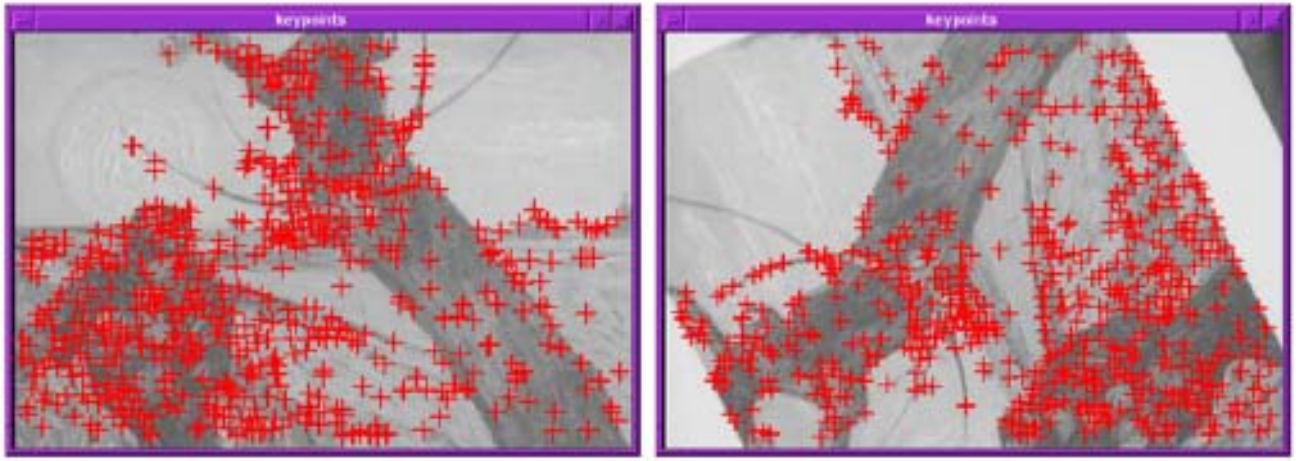

- Key property:
  - In the region around a corner, image gradient has two or more dominant directions
- Corners are *repeatable* and *distinctive*

C.Harris and M.Stephens. "A Combined Corner and Edge Detector."
Proceedings of the 4th Alvey Vision Conference, 1988.

Slide credit: Svetlana Lazebnik

Corners as Distinctive Interest Points

- Design criteria
  - We should easily recognize the point by looking through a small window (*locality*)
  - Shifting the window in *any direction* should give *a large change* in intensity (*good localization*)

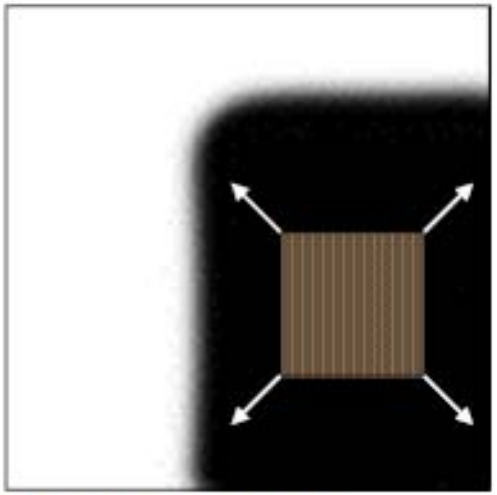

"flat" region:
no change in all directions

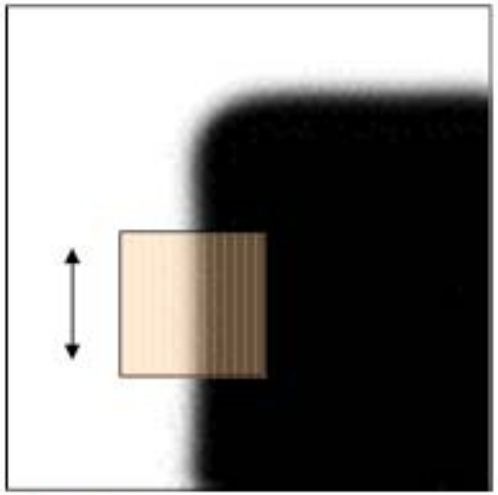

"edge":
no change along the edge direction

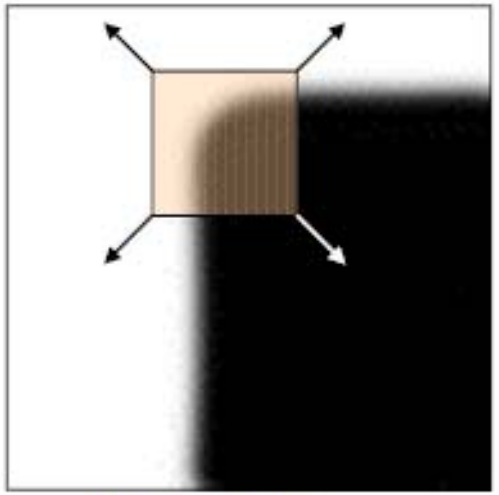

"corner":
significant change in all directions

Slide credit: Alyosha Efros

Harris Detector Formulation

- Change of intensity for the shift [u,v]:

$$E(u,v) = \sum_{x,y} w(x,y)\left[I(x+u, y+v) - I(x,y)\right]^2$$

Window function

Shifted intensity

Intensity

Window function $w(x,y) =$ 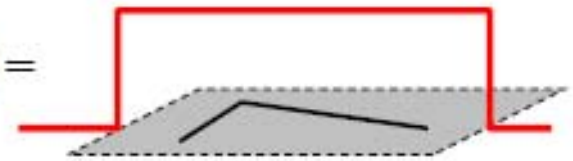 or 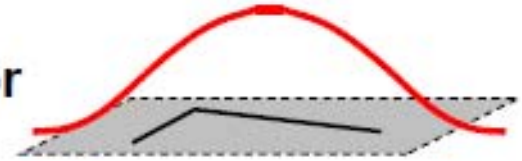

1 in window, 0 outside

Gaussian

Slide credit: Rick Szeliski

Main Intuition of Harris Detector

- **Approximated into the following:**

$$E(u,v) \approx \begin{bmatrix} u & v \end{bmatrix} M \begin{bmatrix} u \\ v \end{bmatrix}$$

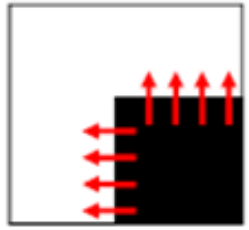

- **In the case of axis-aligned corner:**

$$M = \sum_{(x,y)\in P} \begin{bmatrix} I_x^2 & I_x I_y \\ I_x I_y & I_y^2 \end{bmatrix} = \begin{bmatrix} \lambda_1 & 0 \\ 0 & \lambda_2 \end{bmatrix}$$

  - λ **corresponds gradient directions**
  - **When both** $\lambda_1 \& \lambda_2$ **are non-zero, it is corner!**
  - **Can be extended to rotated corners**

Interpreting the Eigenvalues

- Classification of image points using eigenvalues of M:

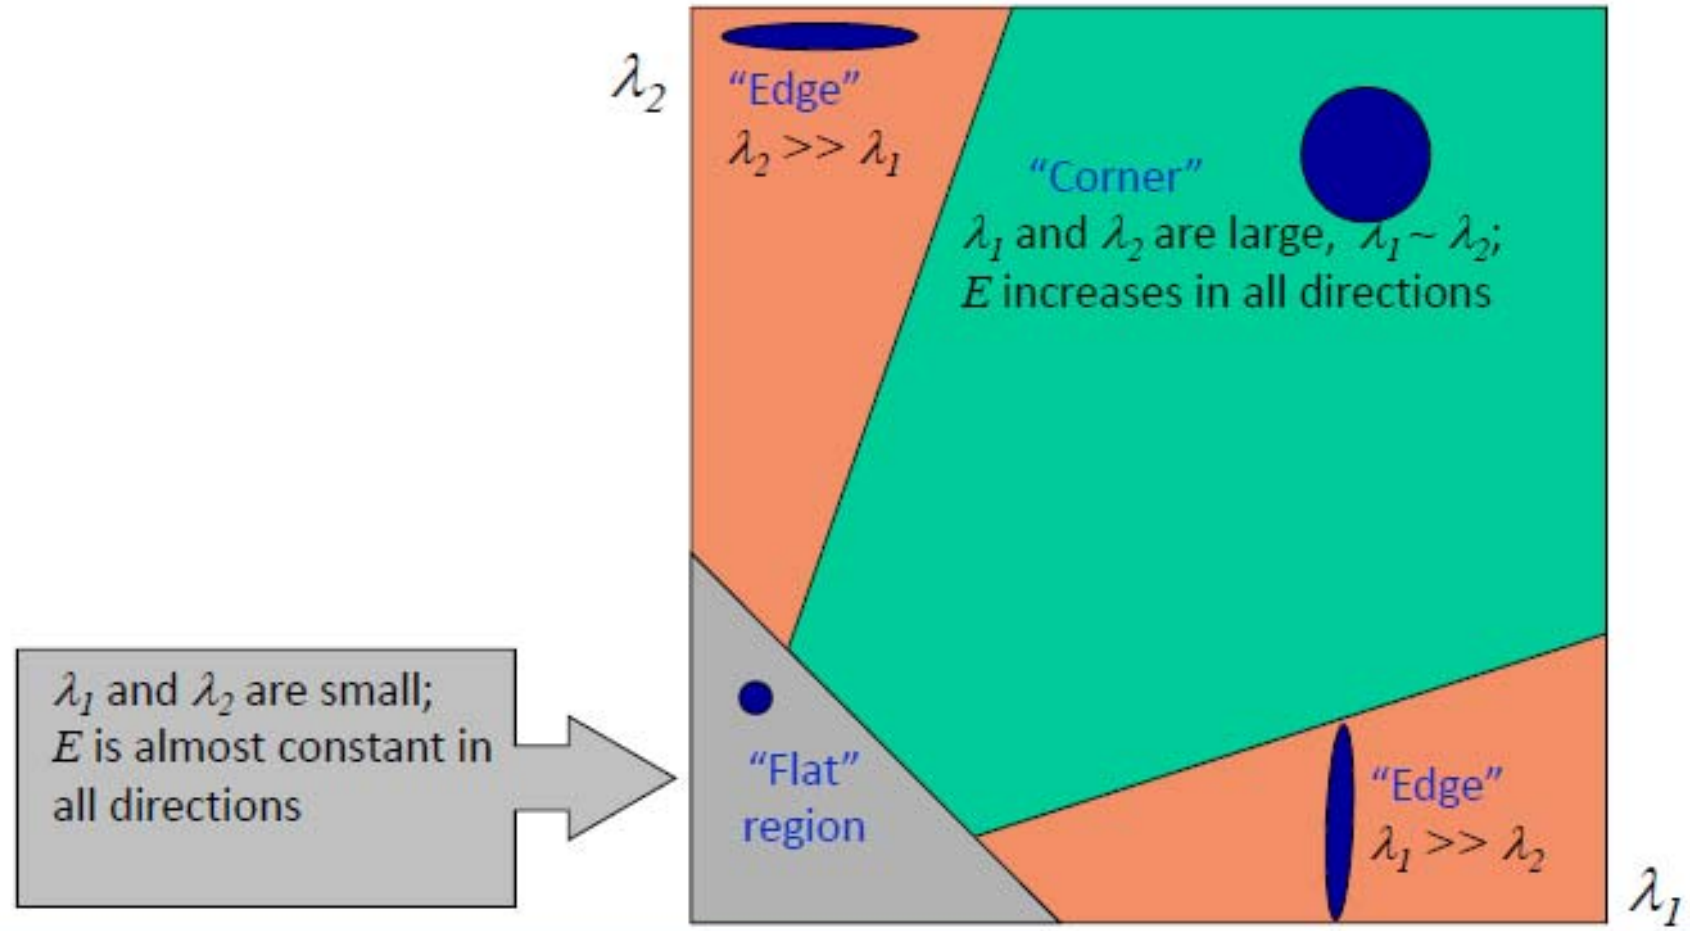

Slide credit: Kristen Grauman

Harris Detector: Workflow

Slide adapted from Darya Frolova, Denis Simakov

Harris Detector: Workflow

- computer corner responses R

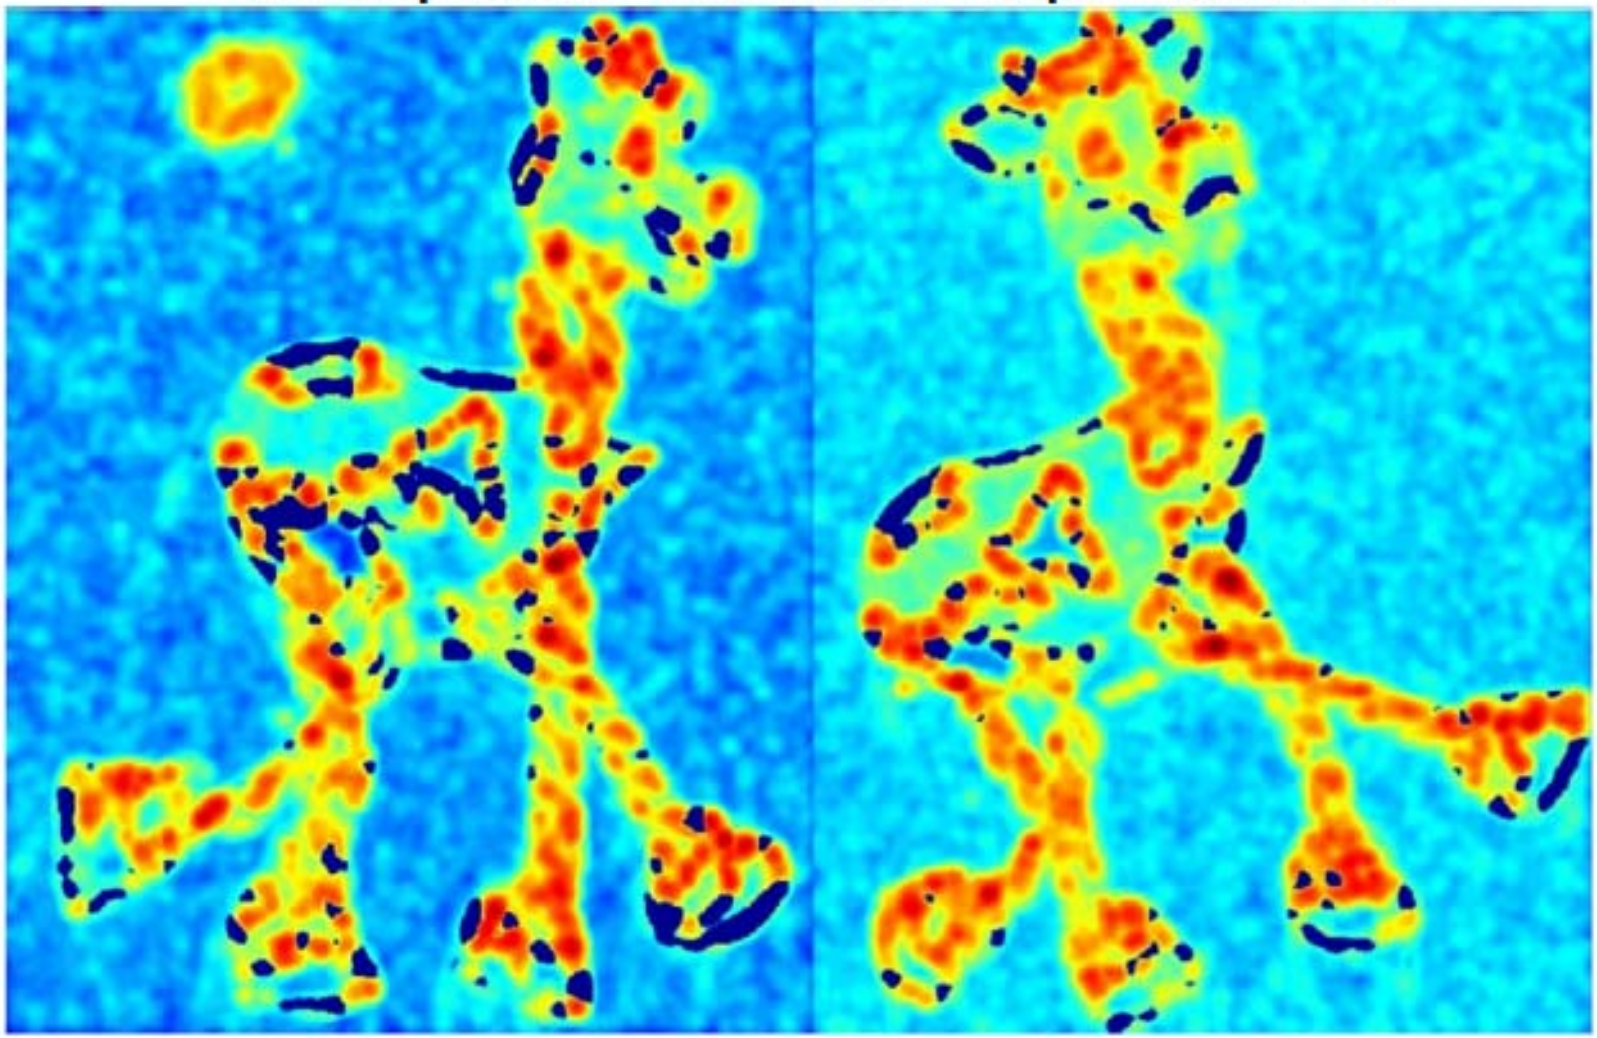

Slide adapted from Darya Frolova, Denis Simakov

Harris Detector: Workflow

- Take only the local maxima of R, where R>threshold

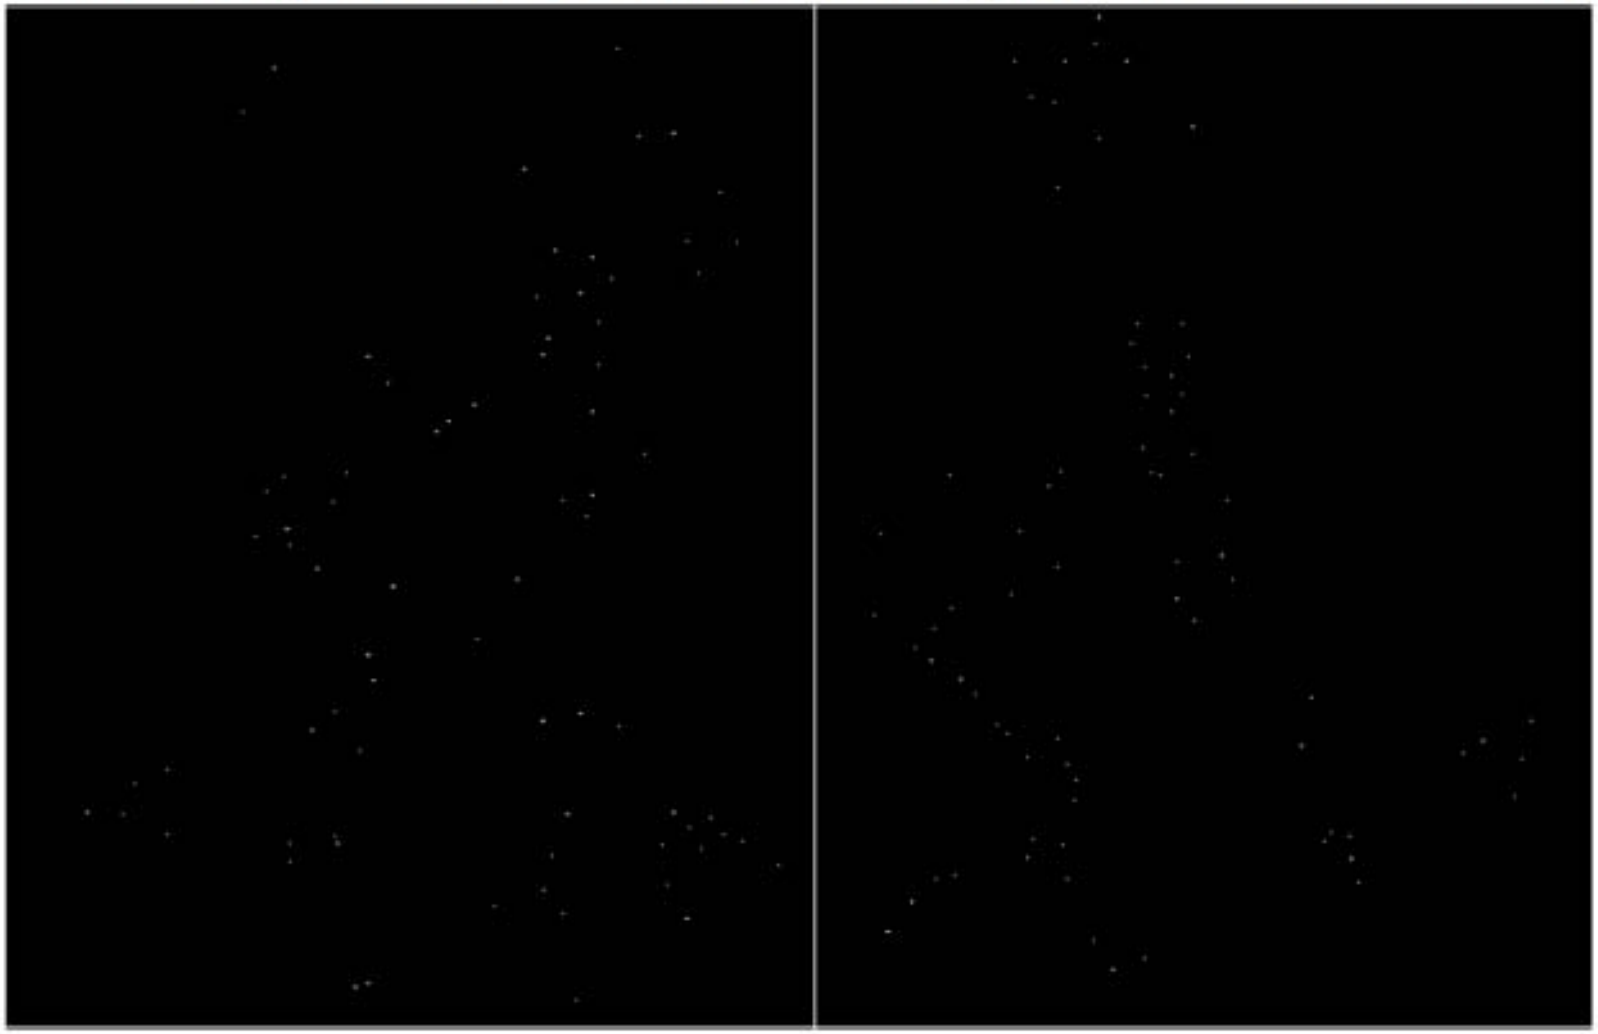

Slide adapted from Darya Frolova, Denis Simakov

Harris Detector: Workflow

- Resulting Harris points

Slide adapted from Darya Frolova, Denis Simakov

Harris Detector – Responses [Harris88]

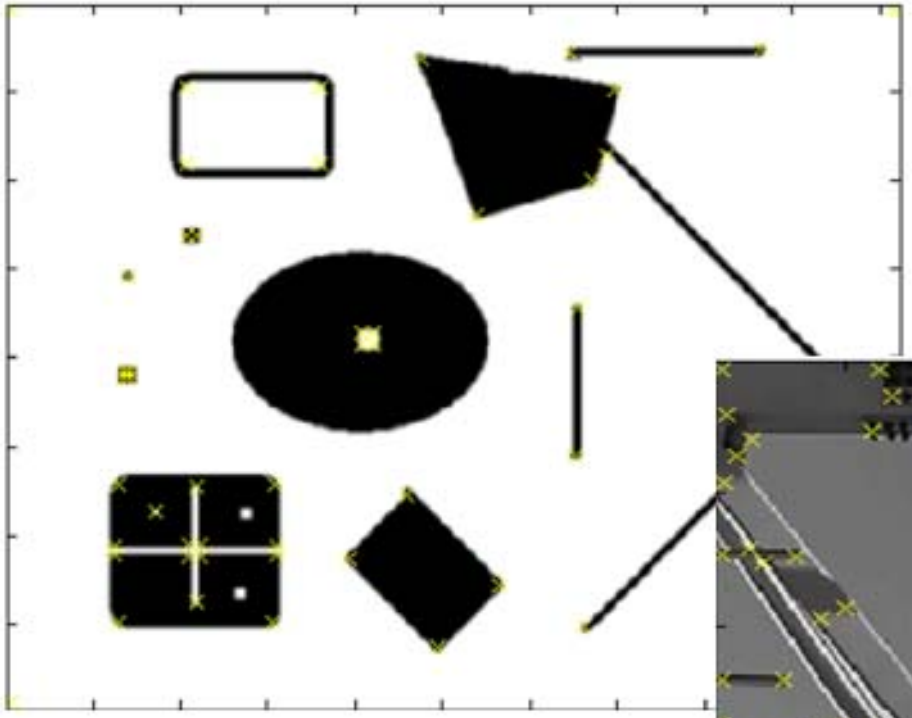

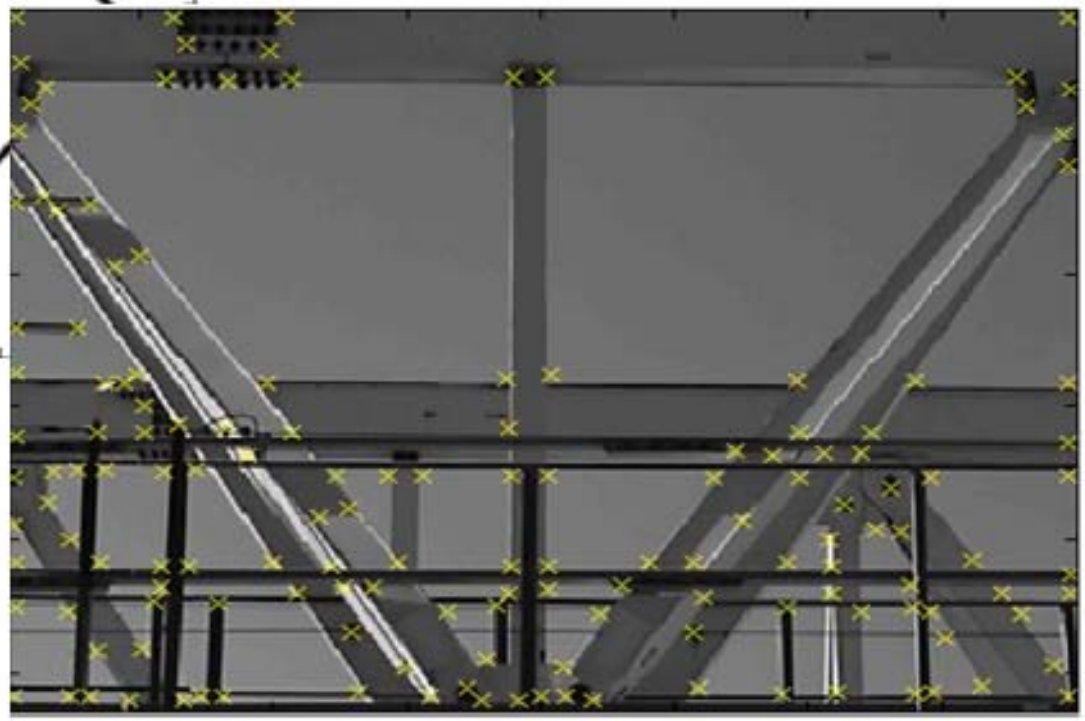

Effect: A very precise corner detector.

Slide credit: Krystian Mikolajczyk

Harris Detector – Responses [Harris88]

Slide credit: Krystian Mikolajczyk

Harris Detector – Responses [Harris88]

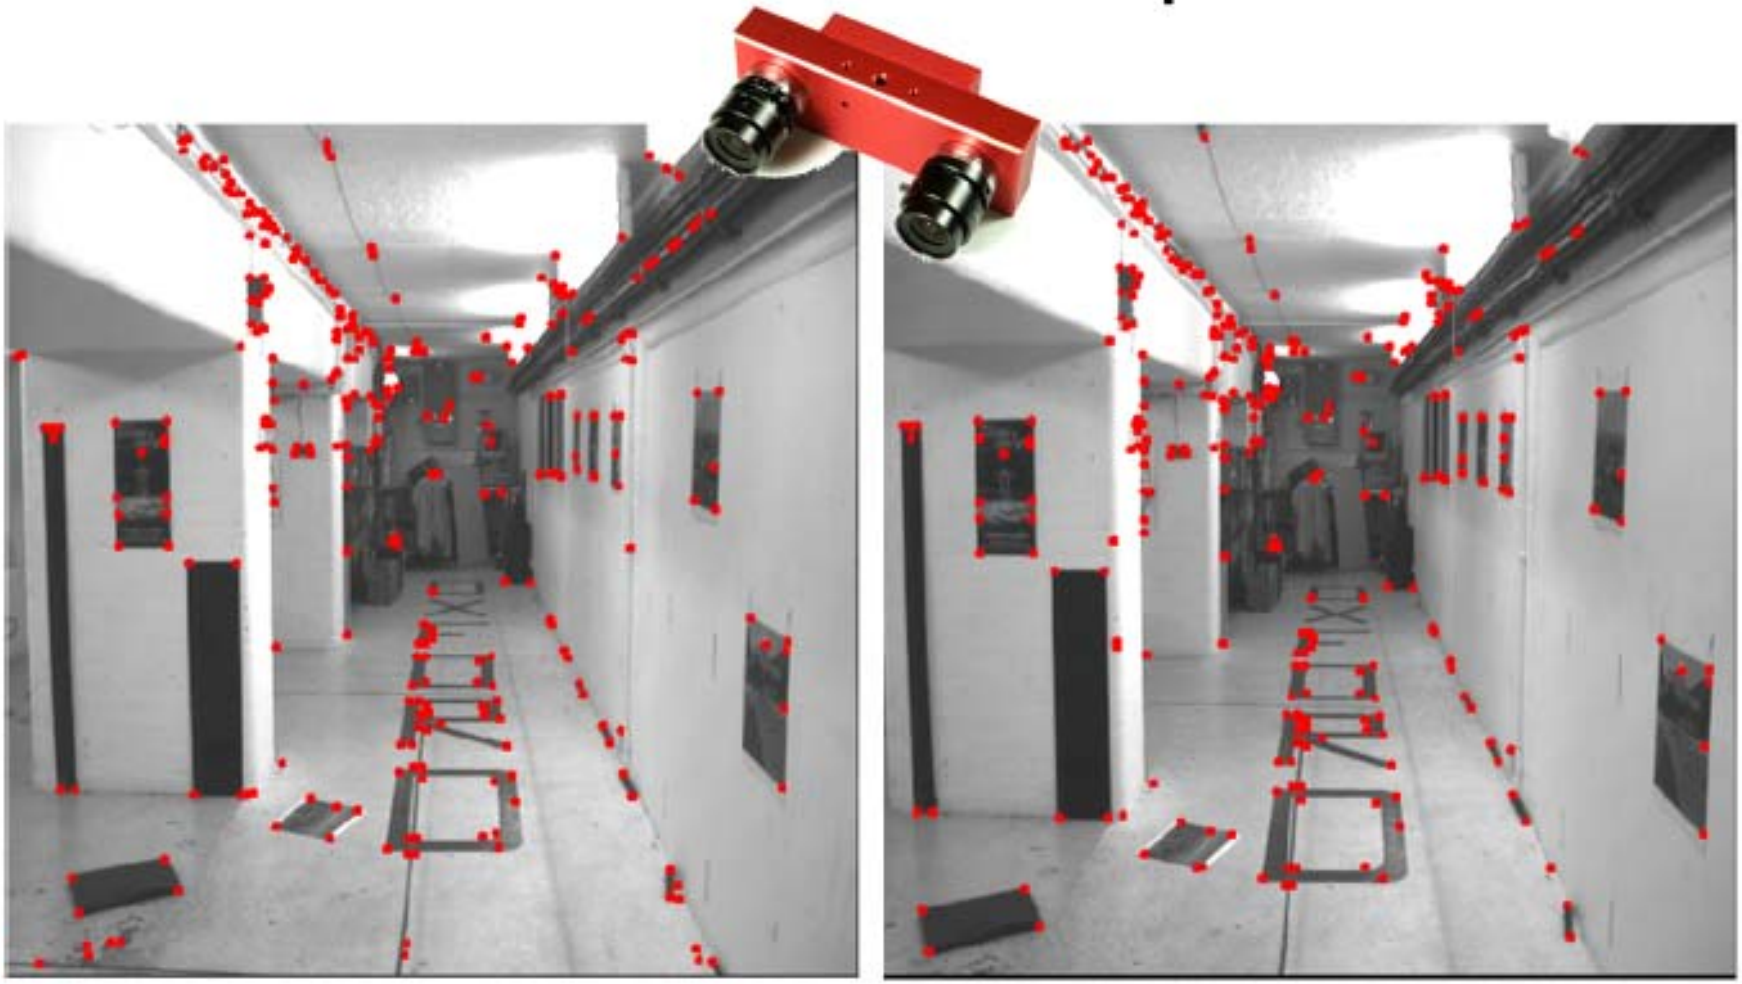

- Results are well suited for finding stereo correspondences

Slide credit: Kristen Grauman

Harris Detector: Properties

- Rotation invariance?

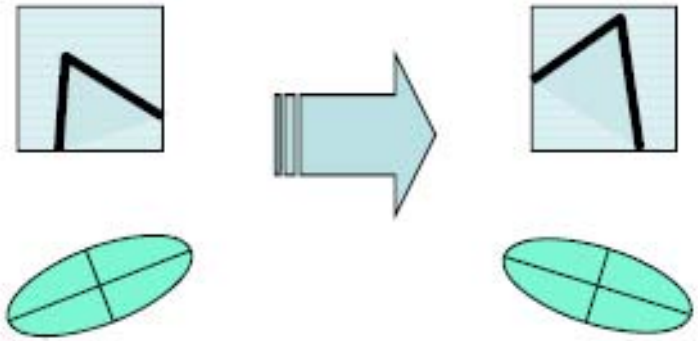

Ellipse rotates but its shape (i.e. eigenvalues) remains the same

Corner response R is invariant to image rotation

Slide credit: Kristen Grauman

Harris Detector: Properties

- Rotation invariance
- Scale invariance?

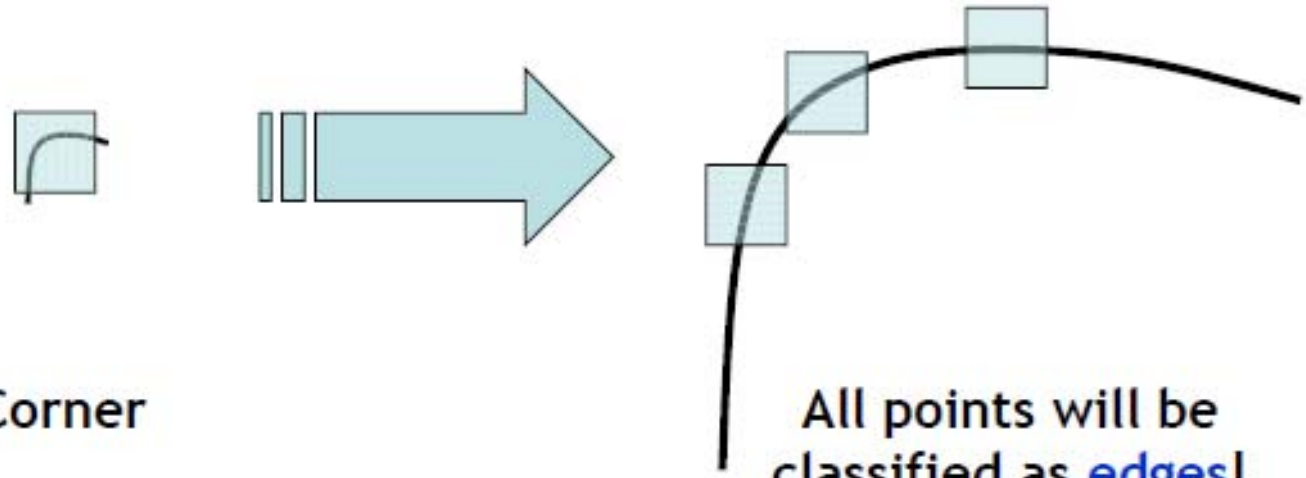

Not invariant to image scale!

Slide credit: Kristen Grauman

Class Objective were:

- Knew related conferences to our course theme
- Understand locally invariant features
  - Key point localization
  - Harris detector: manually designed detector → automatically learned detector using deep learning

Next Time..

- **Scale invariant region selection**

Homework for Every Class

- **Go over the next lecture slides**
- **Come up with one question on what we have discussed today**
  - **https://forms.gle/7vqvJFAcBsebaQs68**
- **Go over recent papers on image search, and submit their summary**
  - **Just one or two (Korean or English) paragraphs are okay**
  - **Do not copy the abstract of the paper**
  - **https://forms.gle/yq19VqqLXwW7TyvZ9**