Recognizing objects

Prof. Noah Snavely
CS1114
http://cs1114.cs.cornell.edu

Cornell University
Computer Science

Administrivia

- Assignment 4 due on Friday
  - The first problem is tricky – please get started early!
- Quiz 4 next Tuesday, 3/31
- Prelim 2 in two weeks, 4/7 (in class)
  - Covers everything since Prelim 1
- Midterm course evaluations

Convex hull

- The convex hull of a set of points is the *smallest* convex polygon that contains all of the points

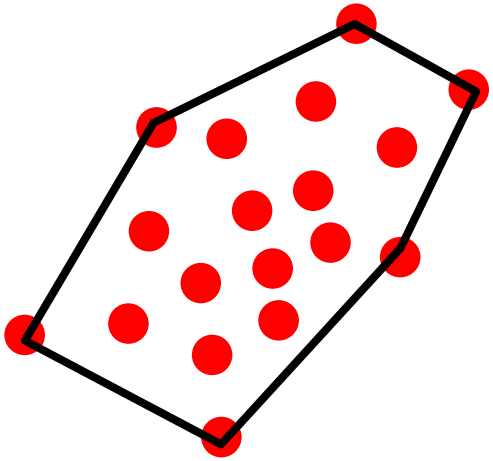

Inverse mapping and bilinear interpolation

- Given an image I_{in} and a 2D transform T, compute the transformed image I_{out} using inverse mapping

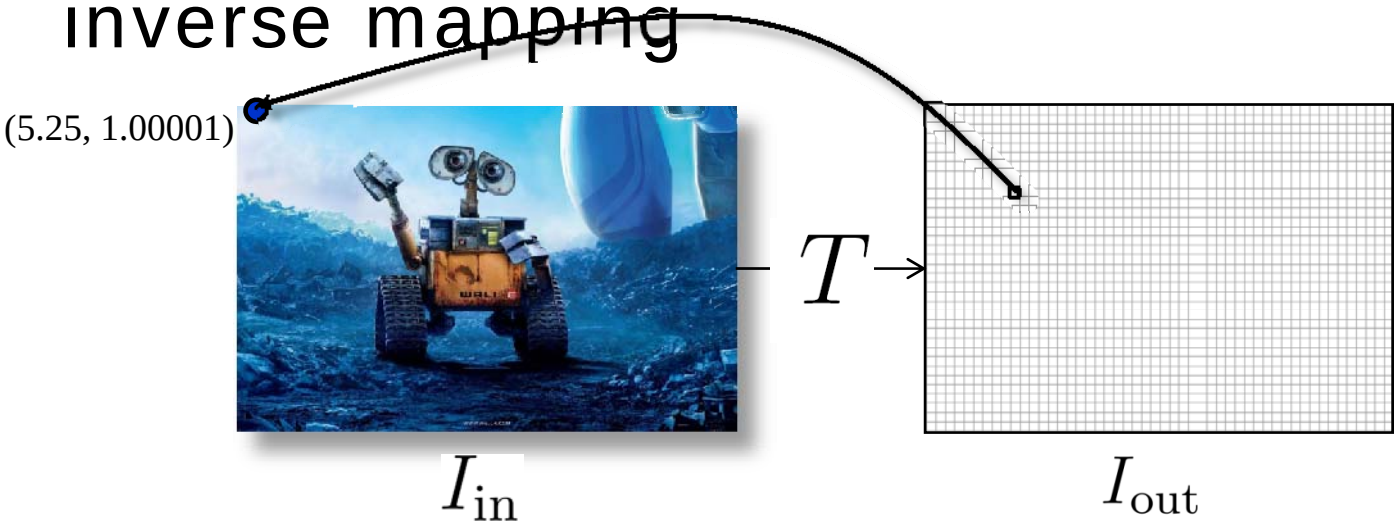

- Given: $T^{-1}\begin{bmatrix} 10 \\ 10 \end{bmatrix} = \begin{bmatrix} 5.25 \\ 1.00001 \end{bmatrix}$

Inverse mapping and bilinear interpolation

- Given: $T^{-1}\begin{bmatrix} 10 \\ 10 \end{bmatrix} = \begin{bmatrix} 5.25 \\ 1.00001 \end{bmatrix}$

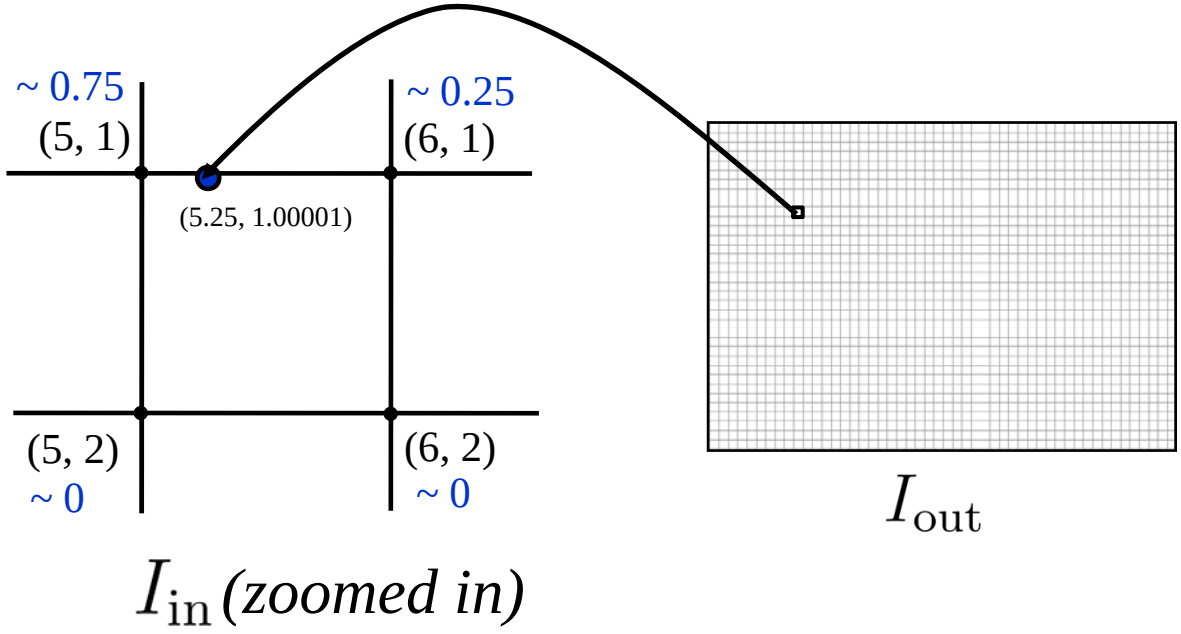

Object recognition

- Our robots can now recognize this:

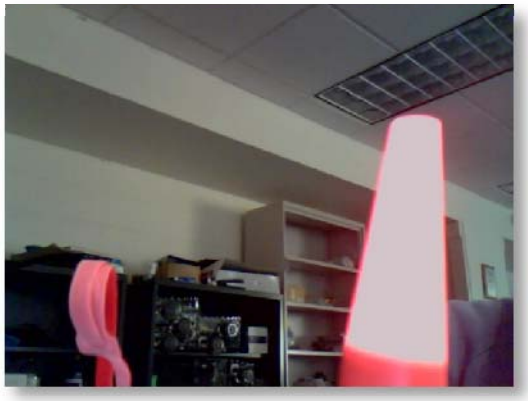

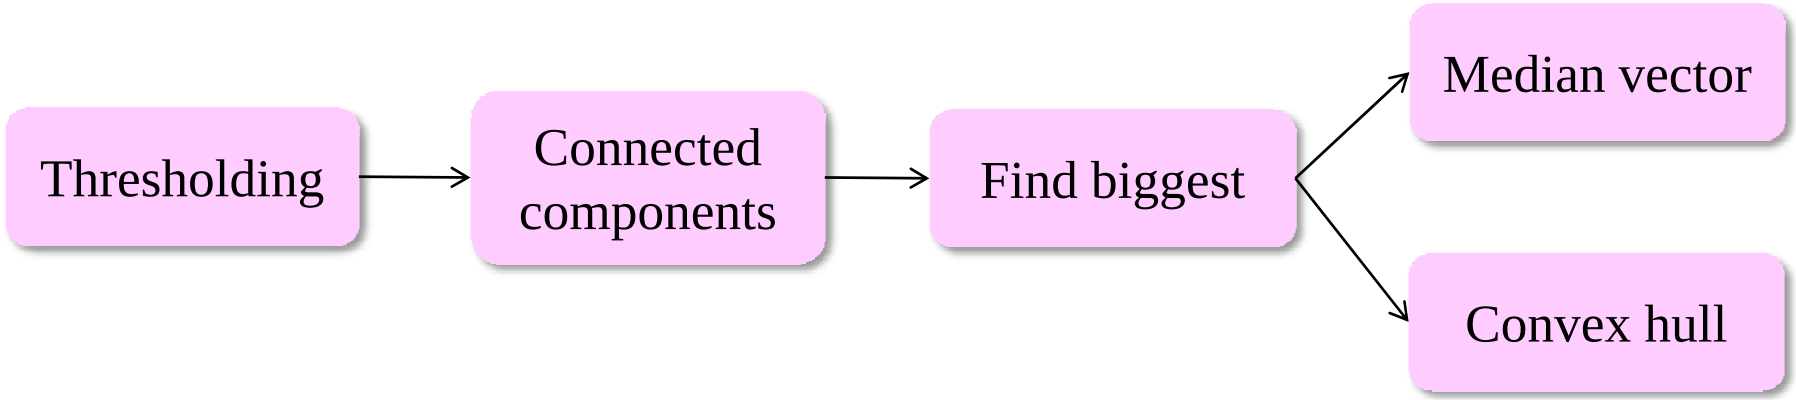

Object recognition

- What else can we recognize with this algorithm?

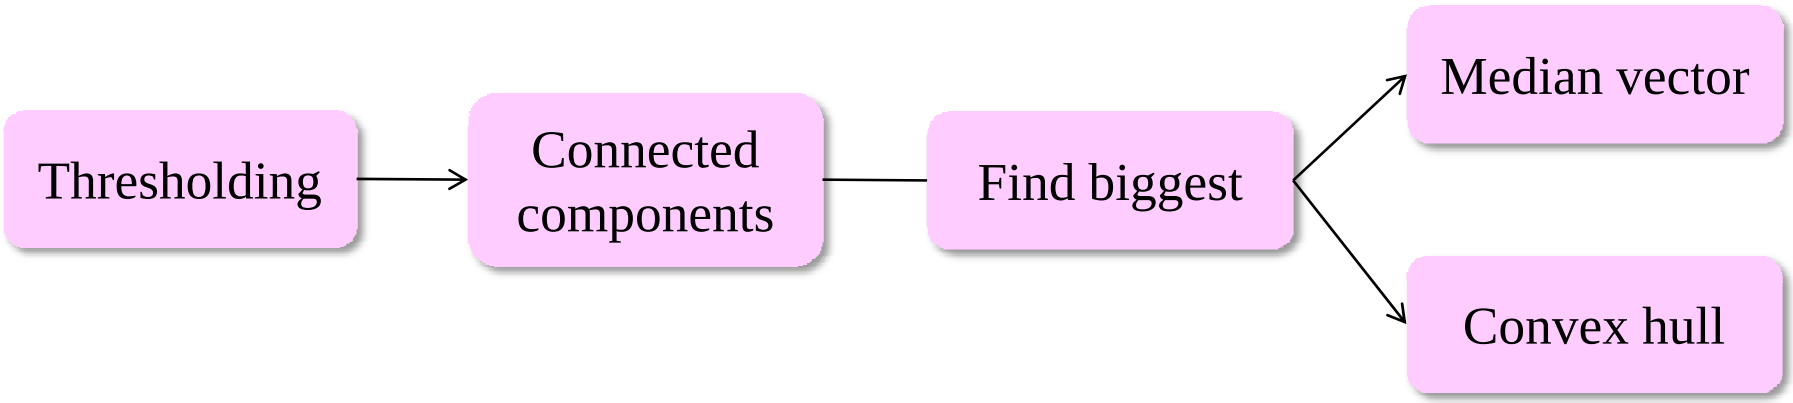

Object recognition

- Can we tell these objects apart?

*All trademarks are properties of their respective owners

Object recognition

- How about these objects?

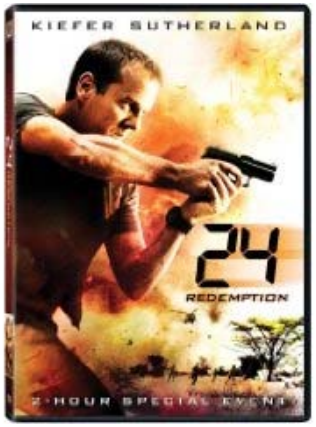

Object recognition

- For some objects **color** is a good feature to use for recognition
  - For instance, a bright red lightstick

- What other features might we use?

Object recognition

- The most general form involves recognizing all of the different *types* of objects in an image

- This problem is extremely difficult for computers

Easier problem

- Given an image of an object, can we find that object in another image?

template

?

new image

- How might we solve this problem?

Finding specific objects

- What problems might come up?

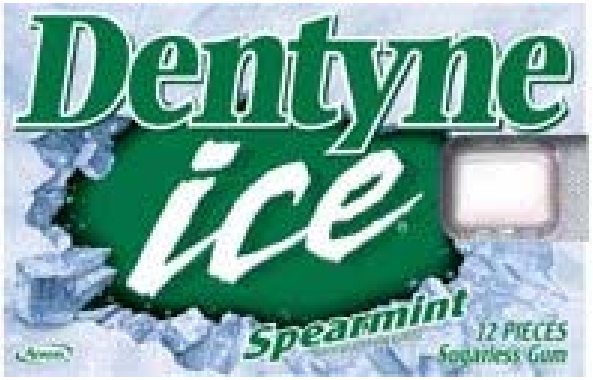

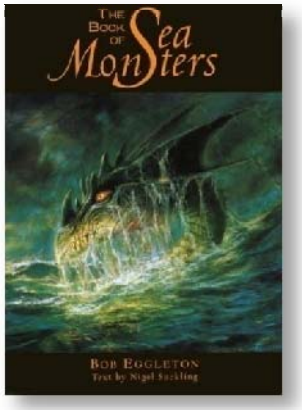

Feature-based image matching

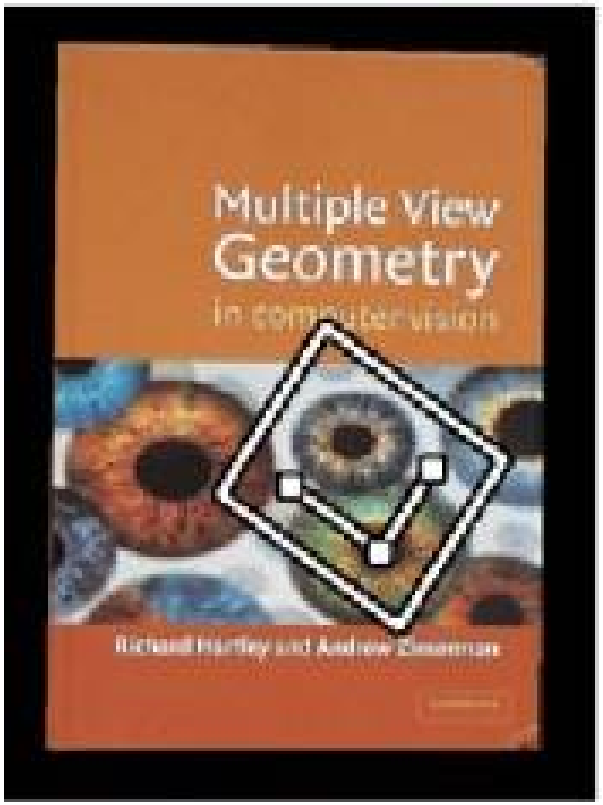

(Slides courtesy Steve Seitz)

Feature-based image matching

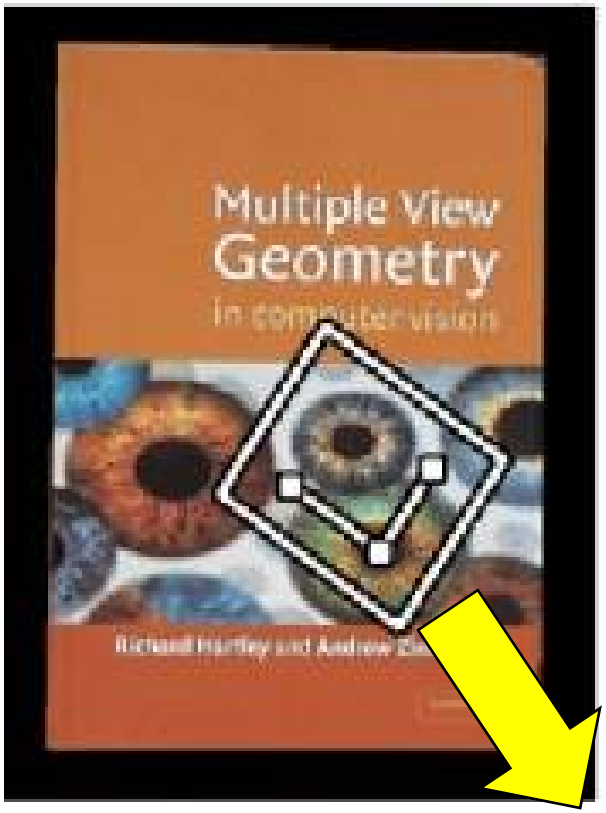

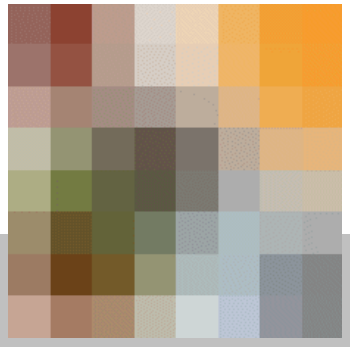

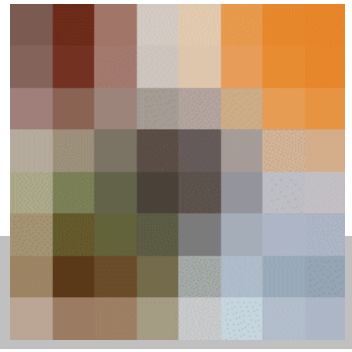

What makes a good feature?

Snoop demo

(Slides courtesy Steve Seitz)

Want uniqueness

- Look for image regions that are unusual
  - Lead to unambiguous matches in other images

- How to define “unusual”?

(Slides courtesy Steve Seitz)

Local measures of uniqueness

- Suppose we only consider a small window of pixels
  - What defines whether a feature is a good or bad candidate?

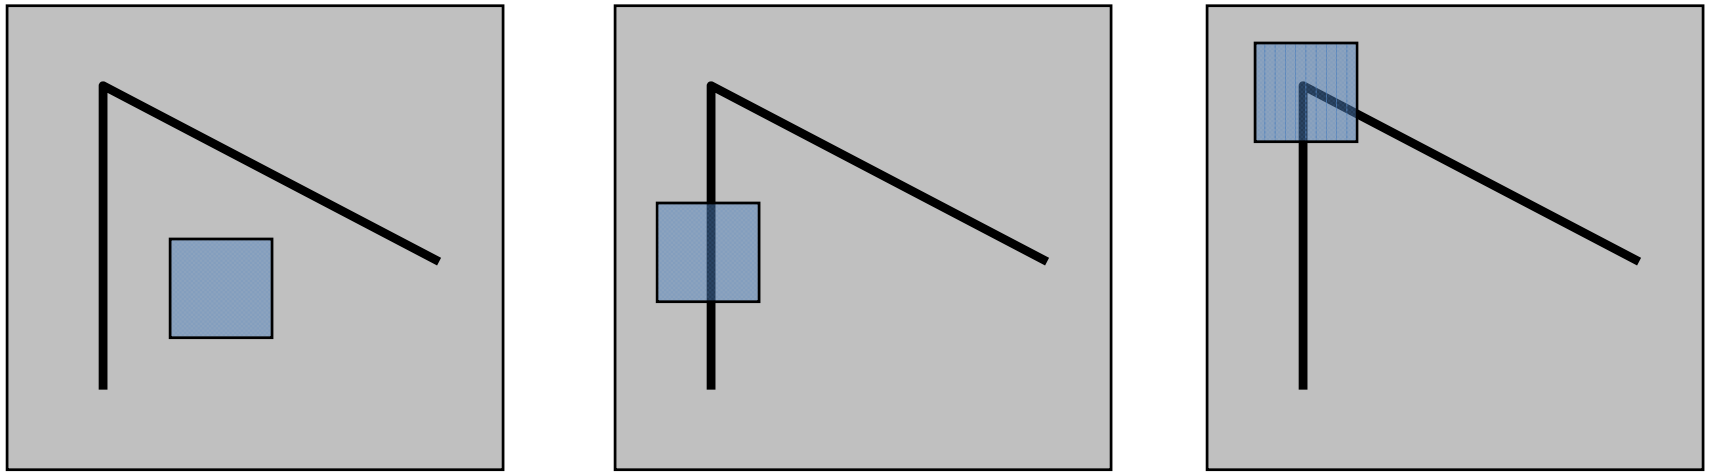

Slide adapted from Darya Frolova, Denis Simakov, Weizmann Institute.

Feature detection

- Local measure of feature uniqueness
  - How does the window change when you shift it?
  - Shifting the window in *any direction* causes a *big change*

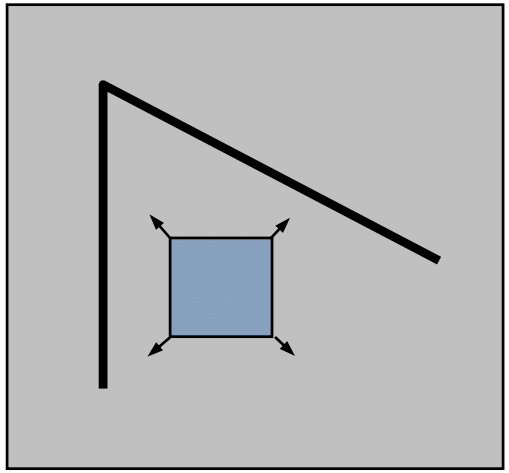

"flat" region:
no change in
all directions

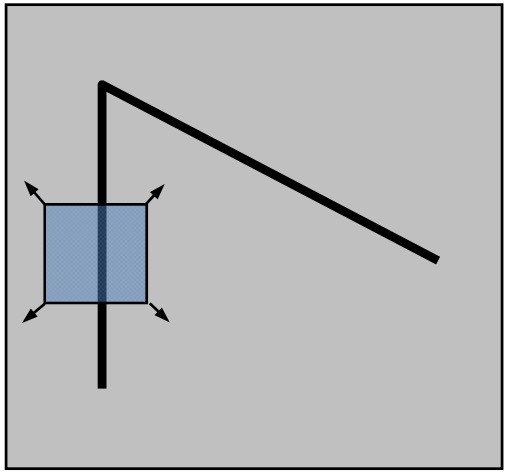

"edge":
no change along
the edge direction

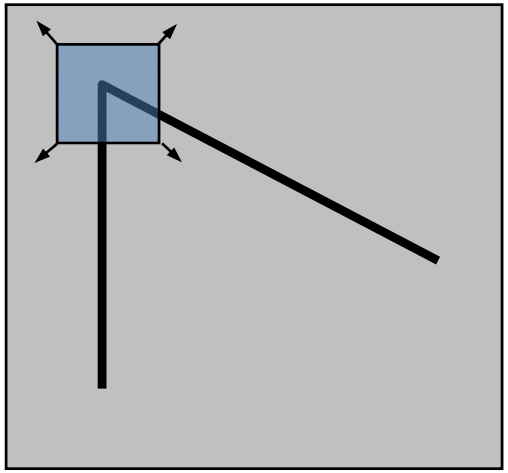

"corner":
significant change
in all directions

Corner detector example

f value (red high, blue low)

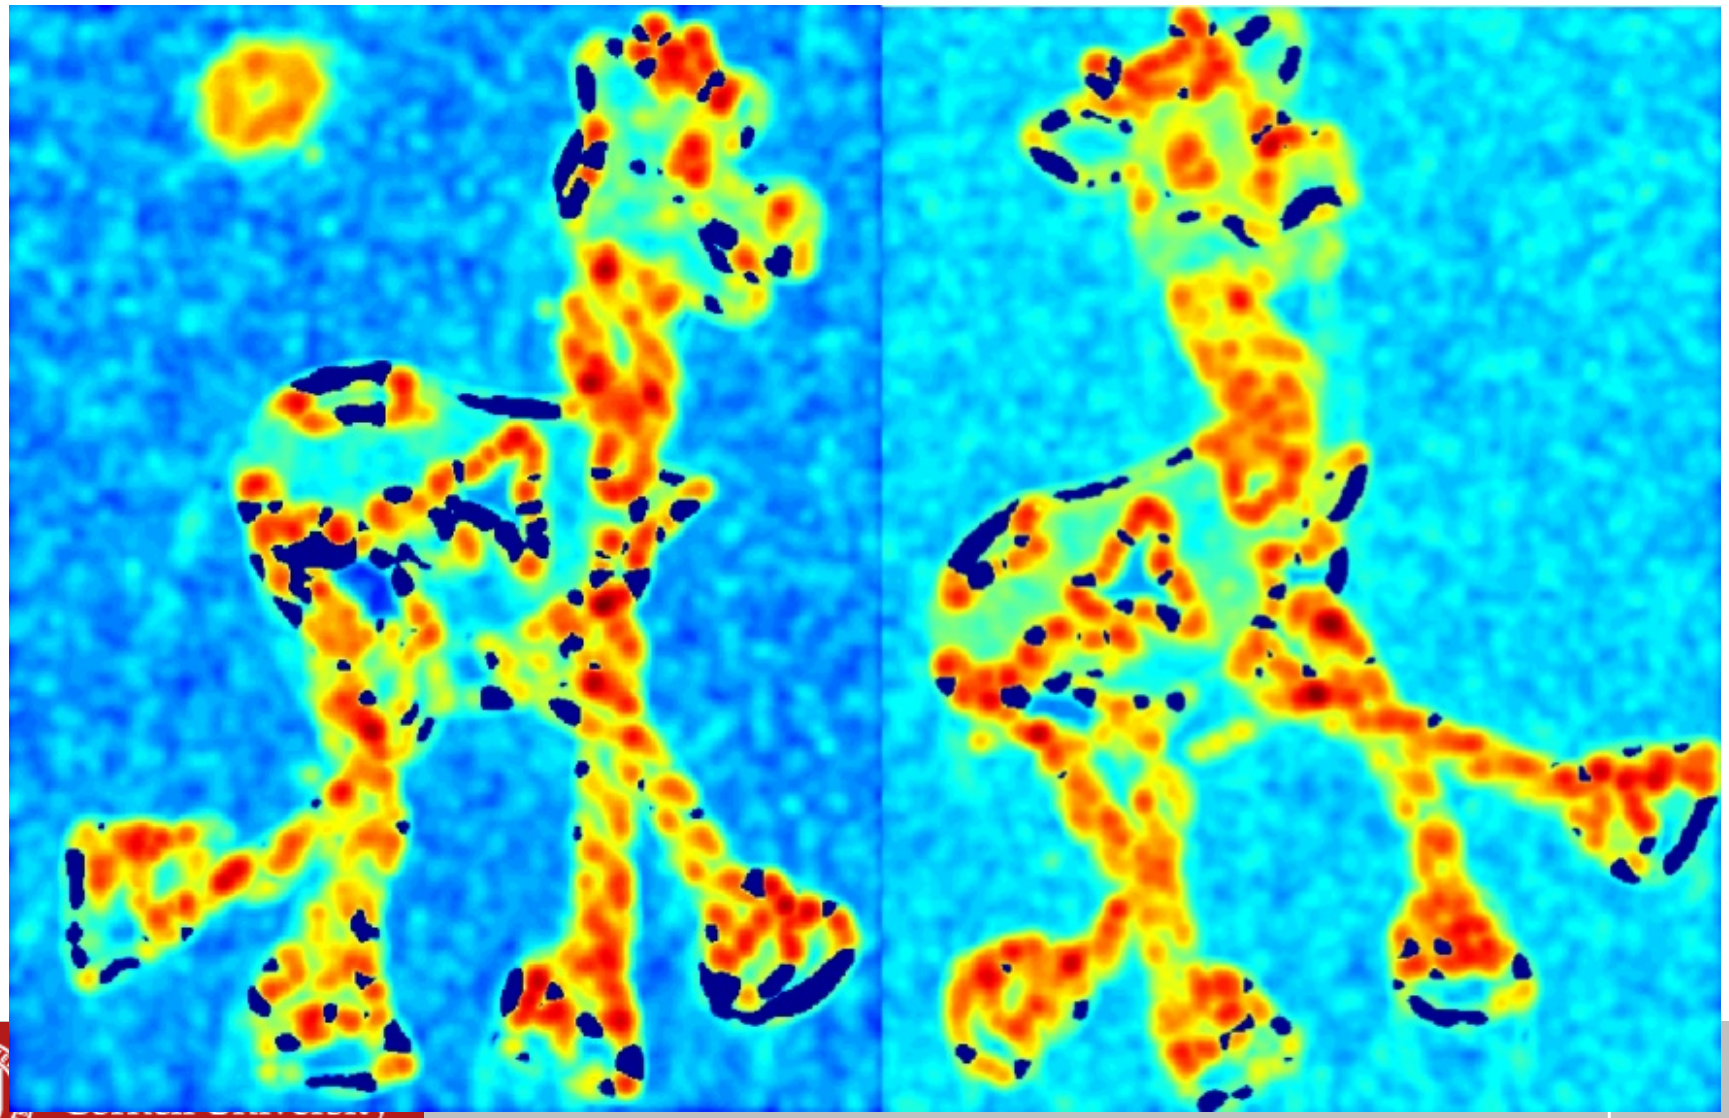

Threshold (f > value)

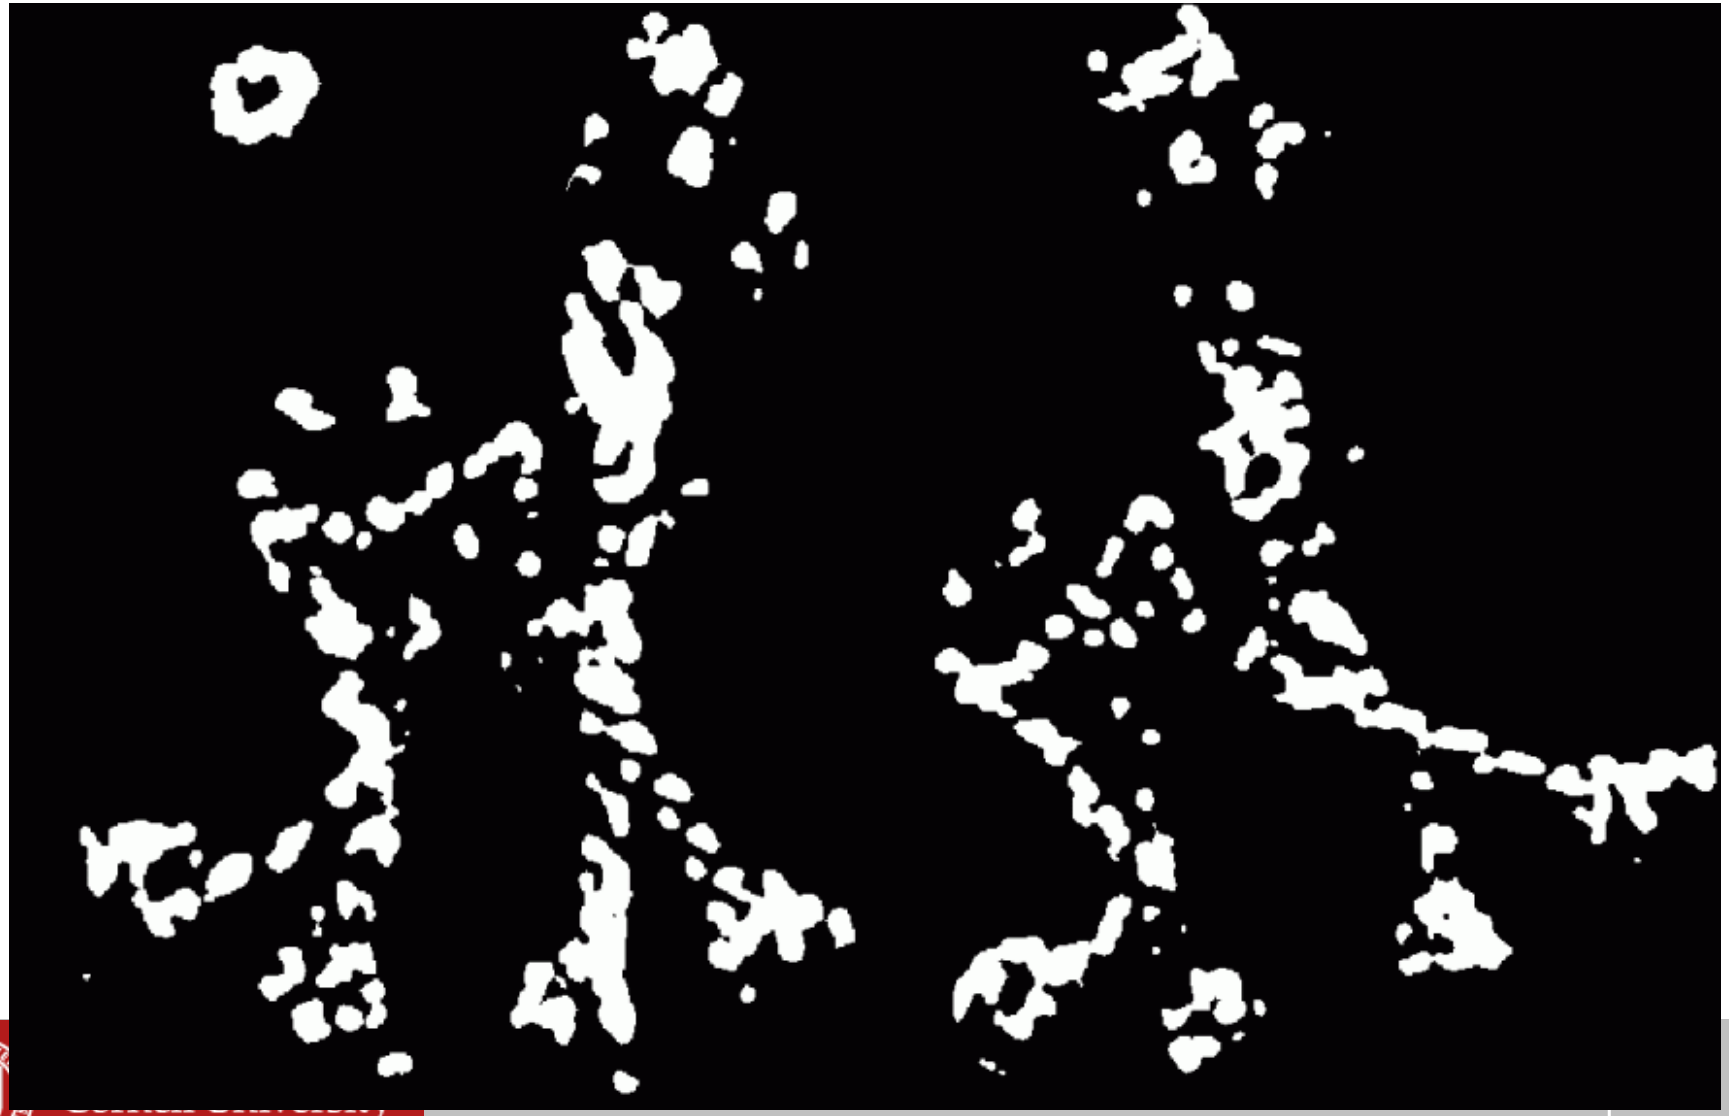

Find local maxima of f

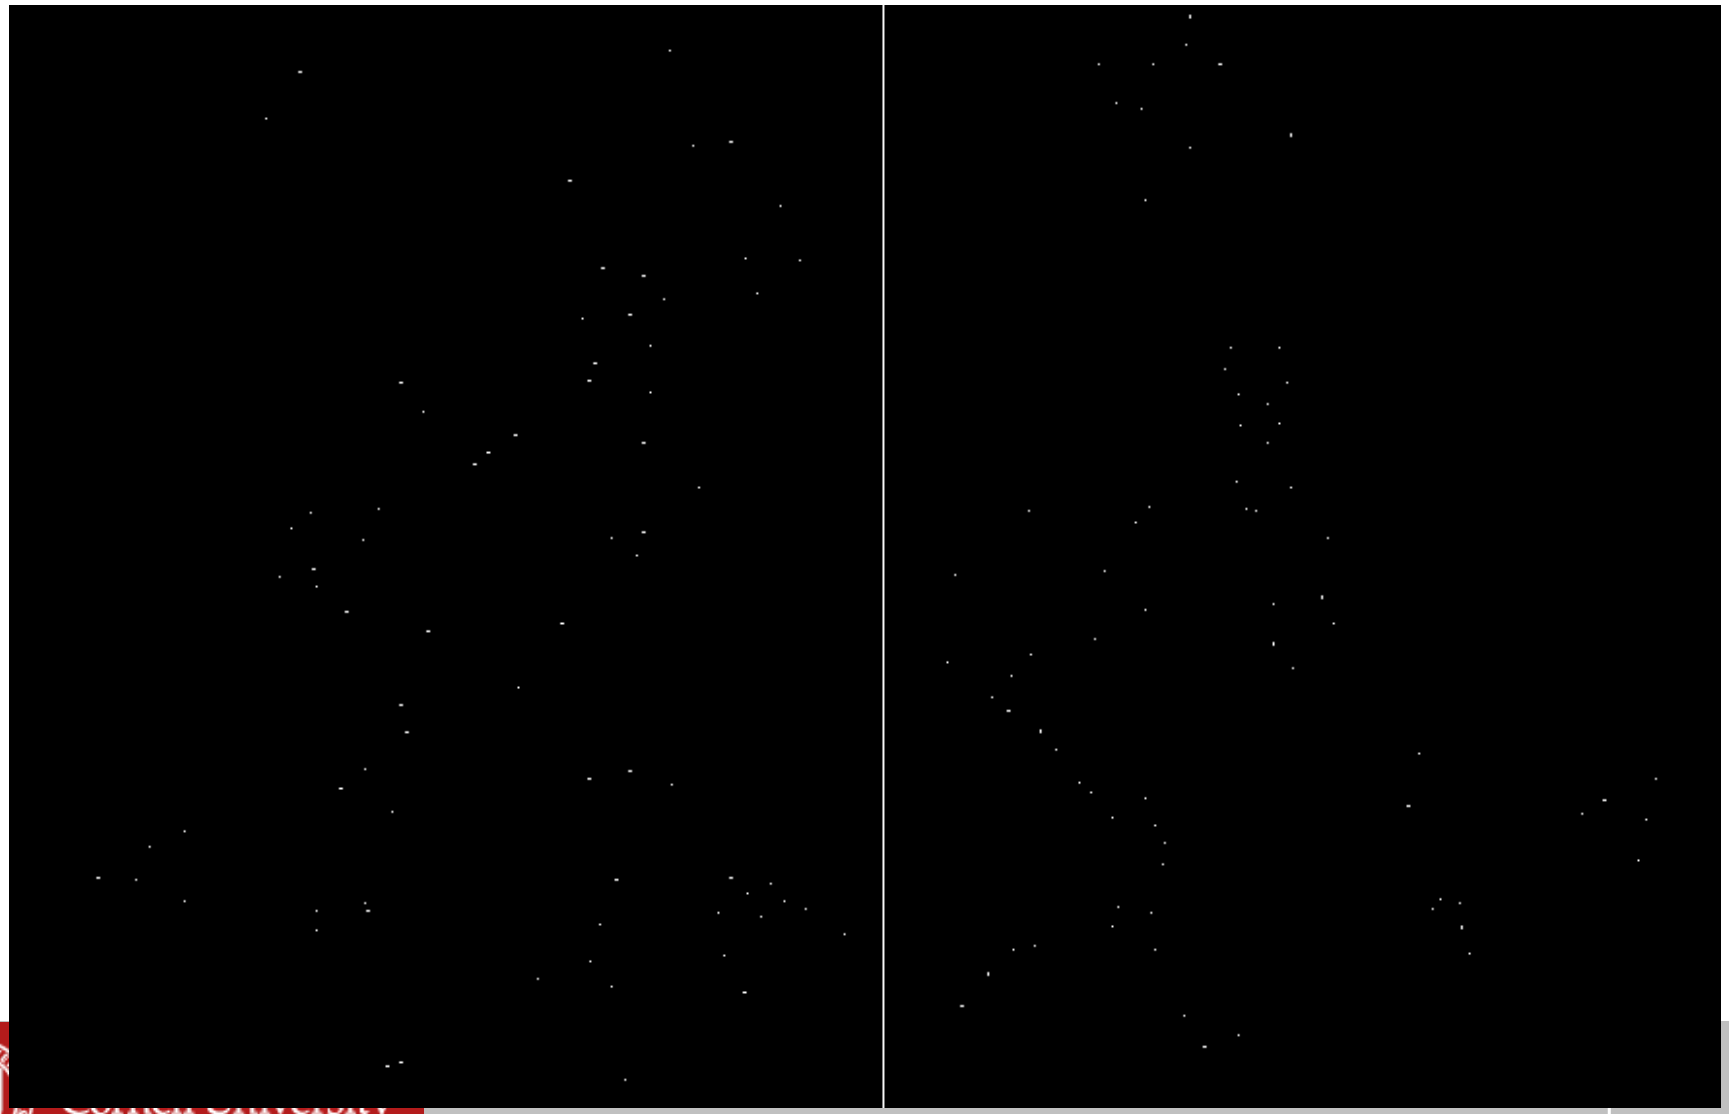

Feature descriptors

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

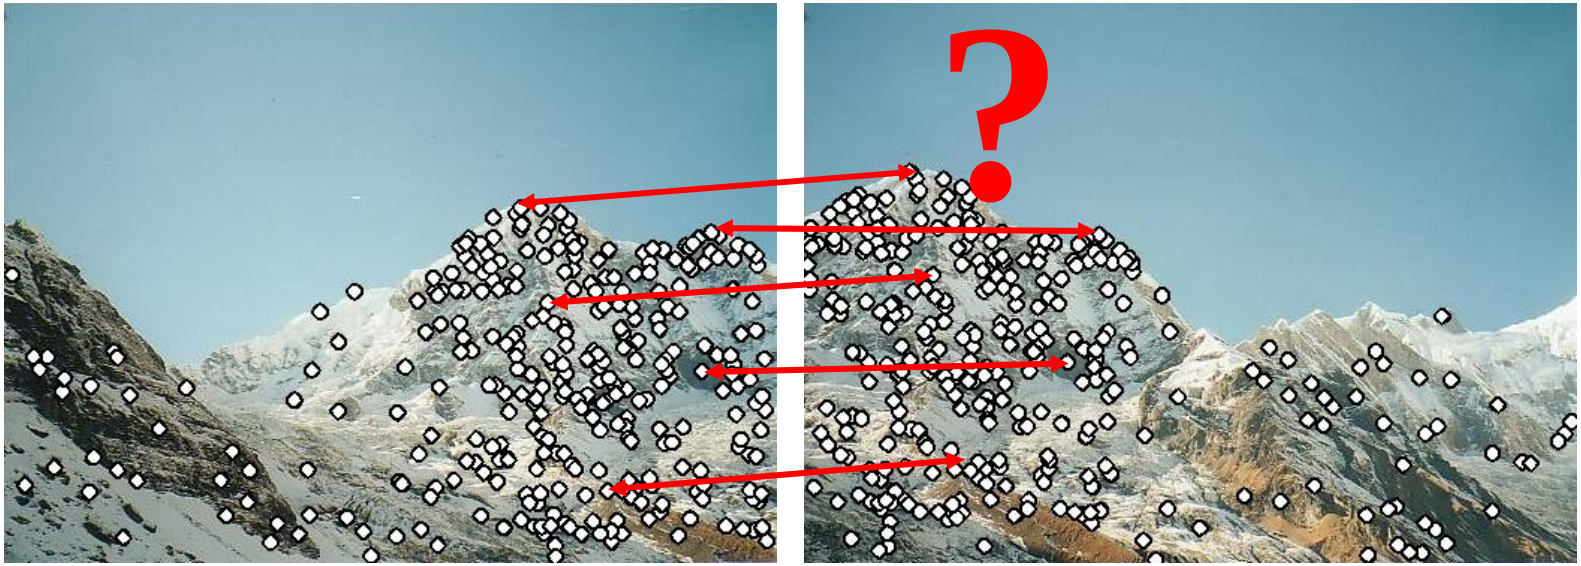

Matching feature points

- How do we match features?
- Lots of possibilities
  - Simple option: match square windows around the point

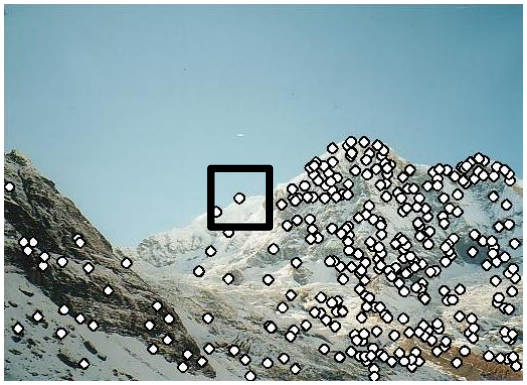

  - State of the art approach: SIFT
    - David Lowe, UBC http://www.cs.ubc.ca/~lowe/keypoints/

Invariant local features

- Find features that are invariant to transformations
  - geometric invariance: translation, rotation, scale
  - photometric invariance: brightness, exposure, …

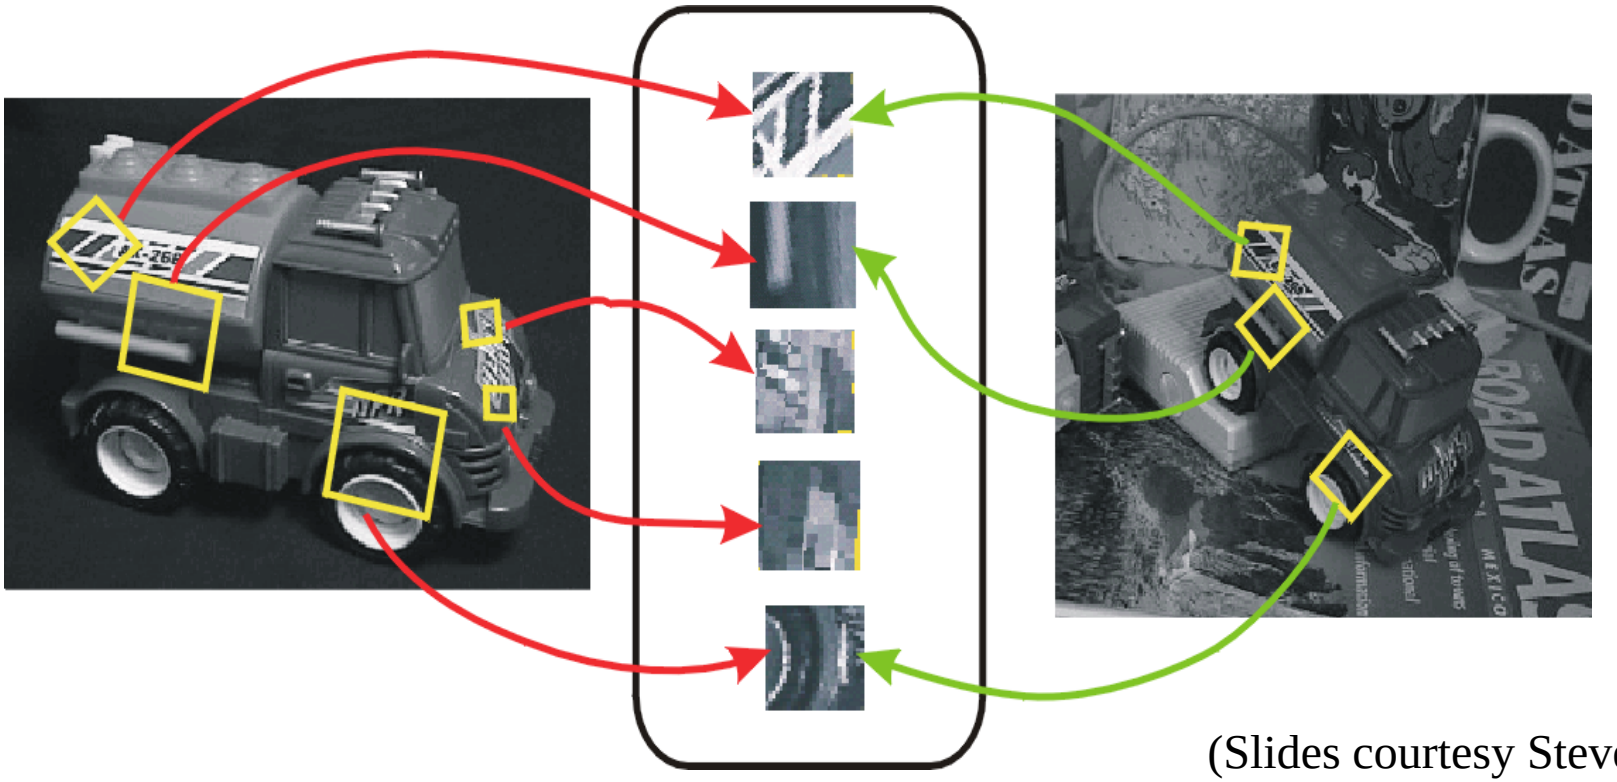

(Slides courtesy Steve Seitz)

SIFT descriptor

- Very complicated, but very powerful
- (The details aren't all that important for this class.)
- 128 dimensional descriptor

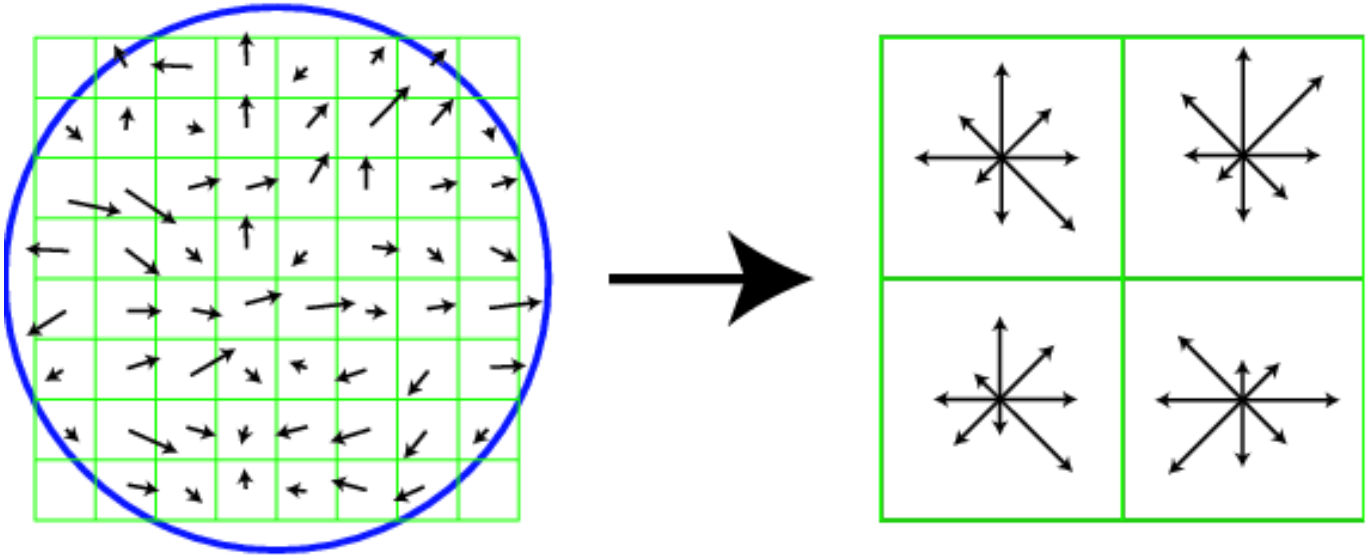

Adapted from a slide by David Lowe

SIFT demo

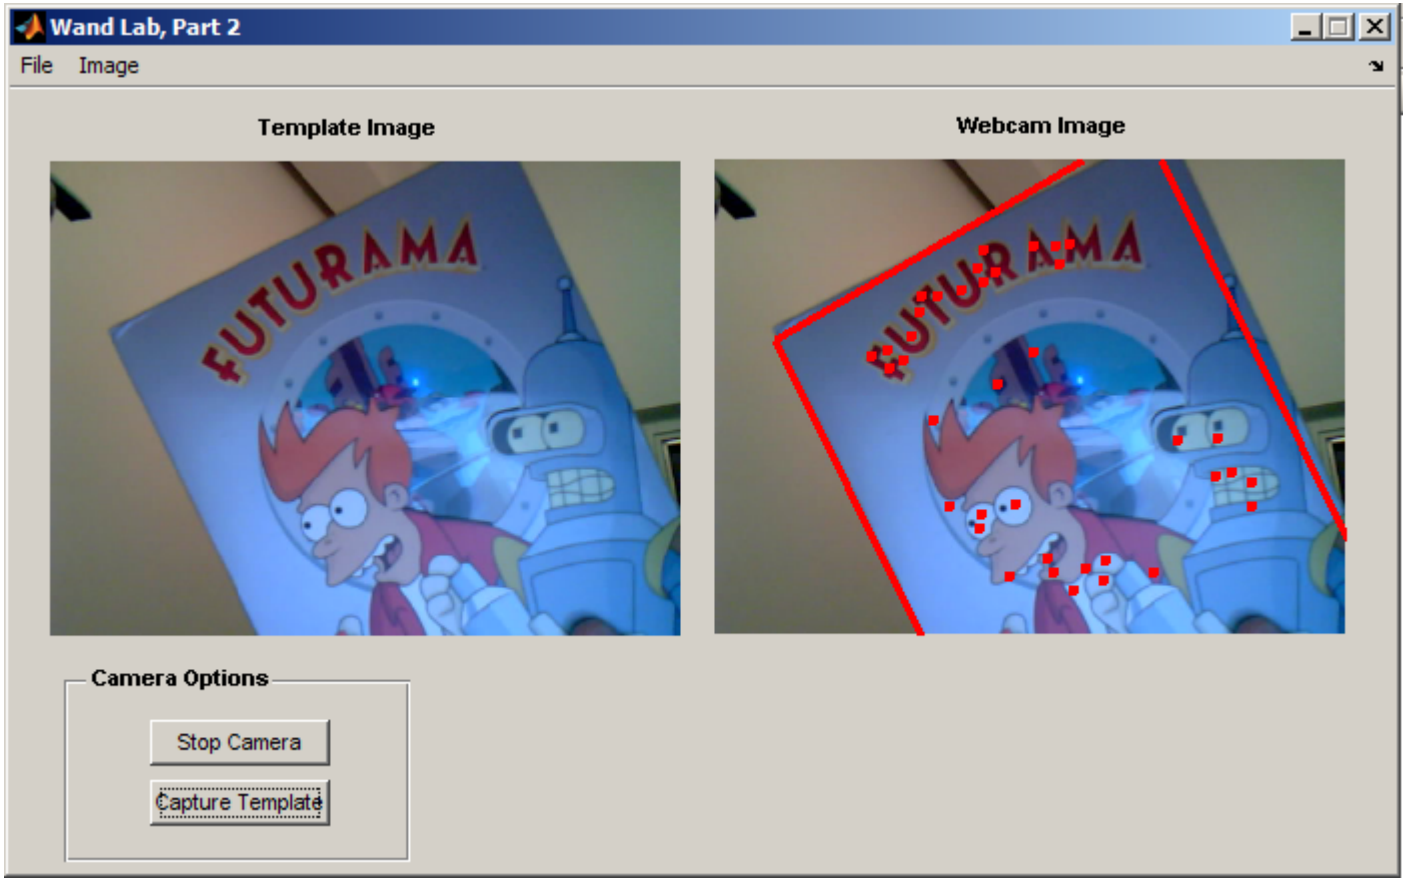
Wand Lab, Part 2
File
Image
Template Image
Webcam Image
FUTURAMA
FUTURAMA
Camera Options
Stop Camera
Capture Template