


Local invariant feature detection

Thurs Feb 23
Kristen Grauman
UT Austin

Last time

- Fitting an arbitrary shape with “active” deformable contours



Deformable contours

a.k.a. active contours, snakes

Given: initial contour (model) near desired object

Goal: evolve the contour to fit exact object boundary



Main idea: elastic band is iteratively adjusted so as to

- be near image positions with high gradients, **and**
- satisfy shape “preferences” or contour priors

[Snakes: Active contour models, Kass, Witkin, & Terzopoulos, ICCV1987]

Figure credit: Yuri Boykov

Deformable contours: intuition





Image from http://www.healthline.com/blogs/exercise_fitness/uploaded_images/HandBand2-795868.JPG

Slide credit: Kristen Grauman

Limitations

- May over-smooth the boundary



- Cannot follow topological changes of objects



Limitations

- External energy: snake does not really “see” object boundaries in the image unless it gets very close to it.





image gradients ∇I
are large only directly on the boundary



Distance transform


- External image can instead be taken from the **distance transform** of the edge image.




original



-gradient



distance transform



edges

Value at (x,y) tells how far that position is from the nearest edge point (or other binary image structure)

>> `help bwdist`

Slide credit: Kristen Grauman

Aspects we need to consider

- Representation of the contours
- Defining the energy functions
 - External
 - Internal
- Minimizing the energy function
- Extensions:
 - Tracking
 - Interactive segmentation

Tracking via deformable contours

1. Use final contour/model extracted at frame t as an initial solution for frame $t+1$
2. Evolve initial contour to fit exact object boundary at frame $t+1$
3. Repeat, initializing with most recent frame.



Tracking Heart Ventricles
(multiple frames)

Tracking via deformable contours




[Visual Dynamics Group](#), Dept. Engineering Science, University of Oxford.

Applications: Traffic monitoring
Human-computer interaction
Animation
Surveillance
Computer assisted diagnosis in medical imaging


<http://www.robots.ox.ac.uk/~vdg/>

3D active contours



Jörgen Ahlberg
<http://www.cvl.isy.liu.se/ScOut/Masters/Papers/Ex1708.pdf>

Interactive forces




How can we implement such an *interactive* force with deformable contours?

Slide credit: Kristen Grauman

Interactive forces

- An energy function can be altered online based on user input – use the cursor to push or pull the initial snake away from a point.
- Modify external energy term to include:



$$E_{push} = \sum_{i=0}^{n-1} \frac{r^2}{|v_i - p|^2}$$

Nearby points get pushed hardest

Intelligent scissors

Another form of interactive segmentation:

Compute optimal paths **from every point to the seed** based on edge-related costs.





Figure 2: Image demonstrating how the live-wire segment adapts and snaps to an object boundary as the free point moves (via cursor movement). The path of the free point is shown in white. Live-wire segments from previous free point positions (t_0 , t_1 , and t_2) are shown in green.

VIDEO

[Mortensen & Barrett, SIGGRAPH 1995, CVPR 1999]

Intelligent scissors



- <http://rivit.cs.byu.edu/Eric/Eric.html>

Deformable contours: pros and cons

Pros:

- Useful to track and fit non-rigid shapes
- Contour remains connected
- Possible to fill in “subjective” contours
- Flexibility in how energy function is defined, weighted.

Cons:

- Must have decent initialization near true boundary, may get stuck in local minimum
- Parameters of energy function must be set well based on prior information

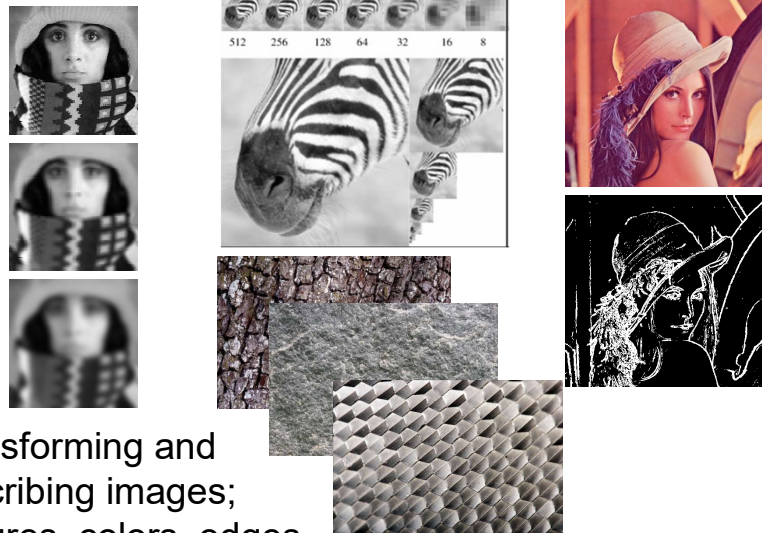
Slide credit: Kristen Grauman

Recap: Deformable contours

- Deformable shapes and active contours are useful for
 - Segmentation: fit or “snap” to boundary in image
 - Tracking: previous frame’s estimate serves to initialize the next
- Fitting active contours:
 - Define terms to encourage certain shapes, smoothness, low curvature, push/pulls, ...
 - Use weights to control relative influence of each component cost
 - Can optimize 2d snakes with Viterbi algorithm.
- Image structure (esp. gradients) can act as attraction force for *interactive* segmentation methods.

Slide credit: Kristen Grauman


Previously: Features and filters




Transforming and
describing images;
textures, colors, edges

Slide credit: Kristen Grauman


Previously: Grouping & fitting




Parallelism




Symmetry



Continuity





Closure



[fig from Shi et al]

Clustering,
segmentation,
fitting; what parts
belong together?




Slide credit: Kristen Grauman


Now: Multiple views



Matching, invariant features,
stereo vision, instance
recognition



Hartley and Zisserman




Lowe




Fei-Fei Li

Slide credit: Kristen Grauman

Important tool for multiple views: Local features



Multi-view matching relies on **local feature** correspondences.



How to detect *which local features* to match?


Local features: main components

1) Detection: Identify the interest points



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

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



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


3) Matching: Determine correspondence between descriptors in two views



Slide credit: Kristen Grauman

Goal: interest operator repeatability

- We want to detect (at least some of) the same points in both images.




No chance to find true matches!

- Yet we have to be able to run the detection procedure *independently* per image.

Goal: descriptor distinctiveness


- We want to be able to reliably determine which point goes with which.



- Must provide some invariance to geometric and photometric differences between the two views.

Local features: main components

1) Detection: Identify the interest points



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

3) Matching: Determine correspondence between descriptors in two views



• What points would you choose?

Slide credit: Kristen Grauman


Detecting corners



Slide credit: Kristen Grauman

Detecting corners

Compute “cornerness” response at every pixel.



Slide credit: Kristen Grauman

Detecting corners



Slide credit: Kristen Grauman


Detecting local invariant features

- Detection of interest points
 - Harris corner detection
 - Scale invariant blob detection: LoG
- (Next time: description of local patches)


Corners as distinctive interest points

We should easily recognize the point by looking through a small window


Shifting a window in *any direction* should give a *large change* in intensity



“flat” region:
no change in
all directions



“edge”:
no change along
the edge
direction





“corner”:
significant
change in all
directions

Slide credit: Alyosha Efros, Darya Frolova, Denis Simakov

Corners as distinctive interest points

$$M = \sum w(x, y) \begin{bmatrix} I_x I_x & I_x I_y \\ I_x I_y & I_y I_y \end{bmatrix}$$

2 x 2 matrix of image derivatives (averaged in neighborhood of a point).



Notation:


$$I_x \Leftrightarrow \frac{\partial I}{\partial x}$$

$$I_y \Leftrightarrow \frac{\partial I}{\partial y}$$

$$I_x I_y \Leftrightarrow \frac{\partial I}{\partial x} \frac{\partial I}{\partial y}$$

What does this matrix reveal?

First, consider an axis-aligned corner:



What does this matrix reveal?

First, consider an axis-aligned corner:

$$M = \sum \begin{bmatrix} I_x^2 & I_x I_y \\ I_x I_y & I_y^2 \end{bmatrix} = \begin{bmatrix} \lambda_1 & 0 \\ 0 & \lambda_2 \end{bmatrix}$$

This means dominant gradient directions align with x or y axis


Look for locations where **both** λ 's are large.

If either λ is close to 0, then this is **not** corner-like.

What if we have a corner that is not aligned with the image axes?

What does this matrix reveal?


Since M is symmetric, we have $M = X \begin{bmatrix} \lambda_1 & 0 \\ 0 & \lambda_2 \end{bmatrix} X^T$



$$Mx_i = \lambda_i x_i$$

The *eigenvalues* of M reveal the amount of intensity change in the two principal orthogonal gradient directions in the window.


Corner response function



“edge”:

$$\lambda_1 \gg \lambda_2$$


$$\lambda_2 \gg \lambda_1$$



“corner”:

λ_1 and λ_2 are large,

$$\lambda_1 \sim \lambda_2;$$



“flat” region

λ_1 and λ_2 are small;

Cornerness score
(other variants possible)

$$f = \frac{\lambda_1 \lambda_2}{\lambda_1 + \lambda_2}$$

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

Harris Detector: Steps




Harris Detector: Steps

Compute corner response f




Harris Detector: Steps

Find points with large corner response: $f > \text{threshold}$




Harris Detector: Steps

Take only the points of local maxima of f



Harris Detector: Steps



Properties of the Harris corner detector

Rotation invariant? Yes


$$M = X \begin{bmatrix} \lambda_1 & 0 \\ 0 & \lambda_2 \end{bmatrix} X^T$$

Scale invariant?

Properties of the Harris corner detector

Rotation invariant? Yes

Scale invariant? No



All points will be
classified as **edges**

Corner !

Scale invariant interest points

How can we independently select interest points in each image, such that the detections are repeatable across different scales?



Automatic Scale Selection




How to find corresponding patch sizes, with only one image in hand?

Automatic scale selection


Intuition:

- Find scale that gives local maxima of some function f in both position and scale.



Automatic Scale Selection


- Function responses for increasing scale (scale signature)





K. Grauman, B. Leibe

Automatic Scale Selection

- Function responses for increasing scale (scale signature)



$$f(I_{i_1..i_m}(x, \sigma))$$




$$f(I_{i_1..i_m}(x', \sigma))$$


K. Grauman, B. Leibe

Automatic Scale Selection

- Function responses for increasing scale (scale signature)



$$f(I_{i_1..i_m}(x, \sigma))$$




$$f(I_{i_1..i_m}(x', \sigma))$$

K. Grauman, B. Leibe

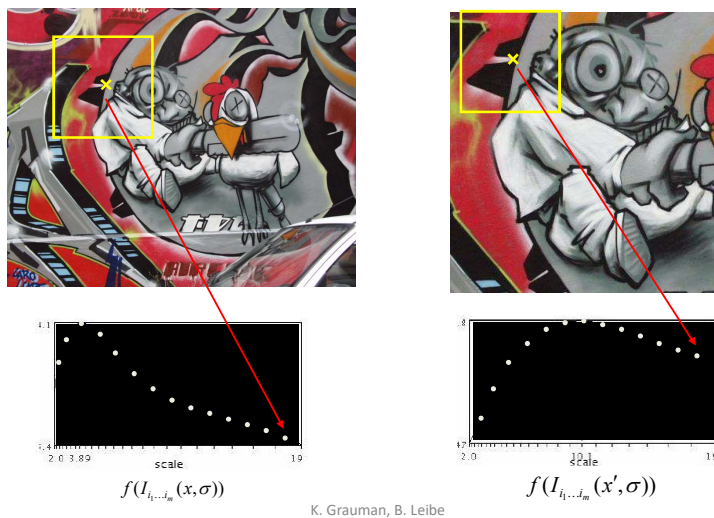
Automatic Scale Selection

- Function responses for increasing scale (scale signature)





Automatic Scale Selection

- Function responses for increasing scale (scale signature)




Automatic Scale Selection

- Function responses for increasing scale (scale signature)



$$f(I_{i_1...i_n}(x, \sigma))$$





$$f(I_{i_1...i_n}(x', \sigma'))$$

K. Grauman, B. Leibe

What can be the “signature” function?

Blob detection in 2D


Laplacian of Gaussian: Circularly symmetric operator for blob detection in 2D



$$\nabla^2 g = \frac{\partial^2 g}{\partial x^2} + \frac{\partial^2 g}{\partial y^2}$$


Blob detection in 2D: scale selection

Laplacian-of-Gaussian = "blob" detector $\nabla^2 g = \frac{\partial^2 g}{\partial x^2} + \frac{\partial^2 g}{\partial y^2}$



Blob detection in 2D

We define the *characteristic scale* as the scale that produces peak of Laplacian response



characteristic scale

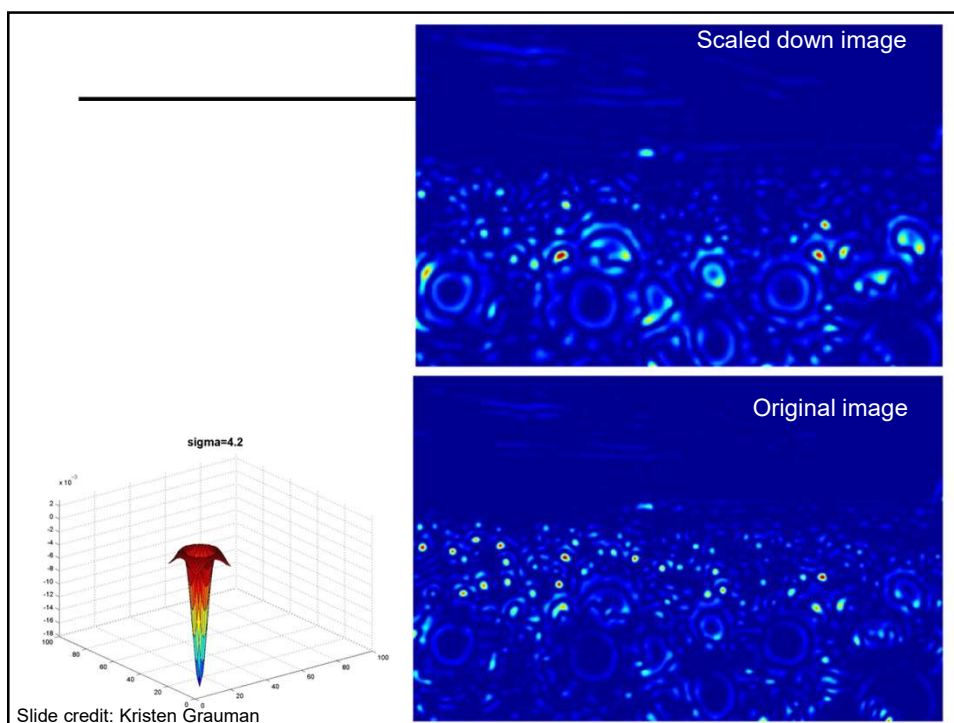
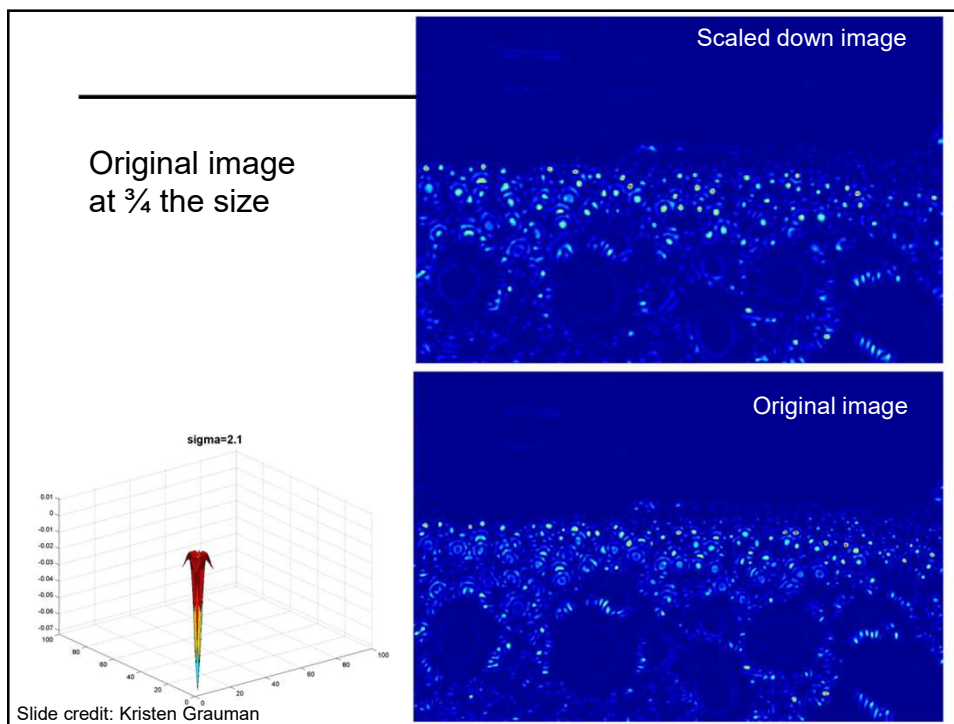
Slide credit: Lana Lazebnik

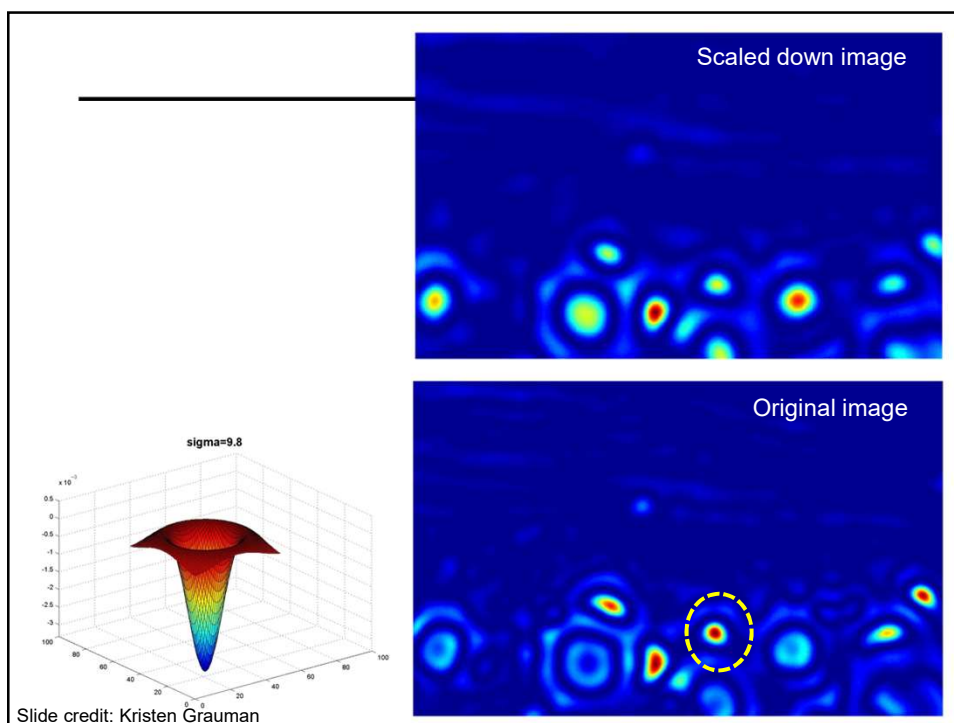
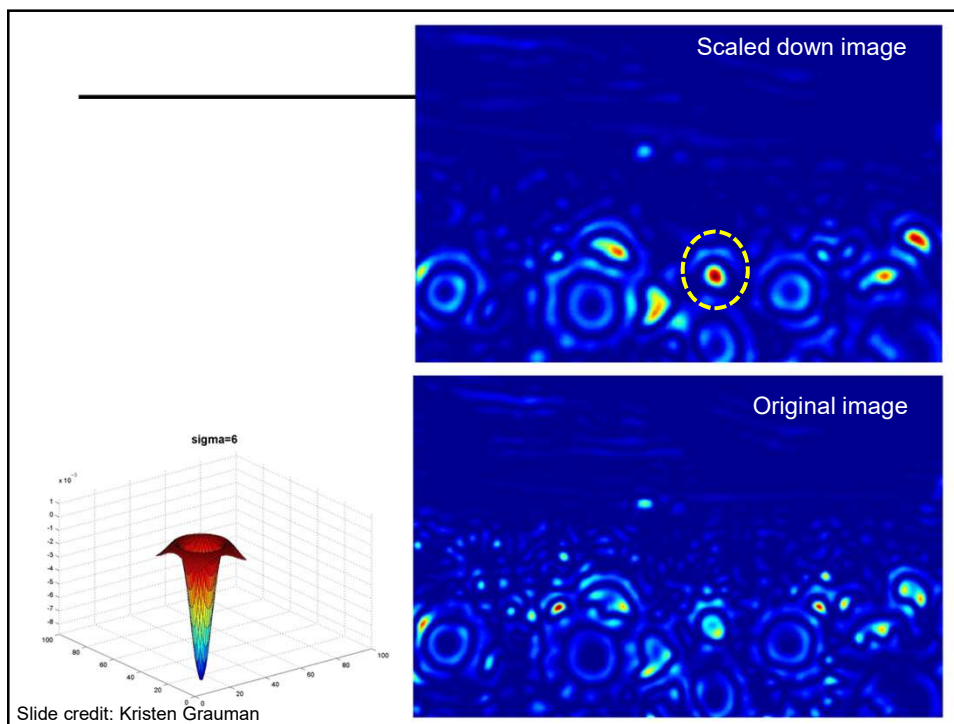
Example


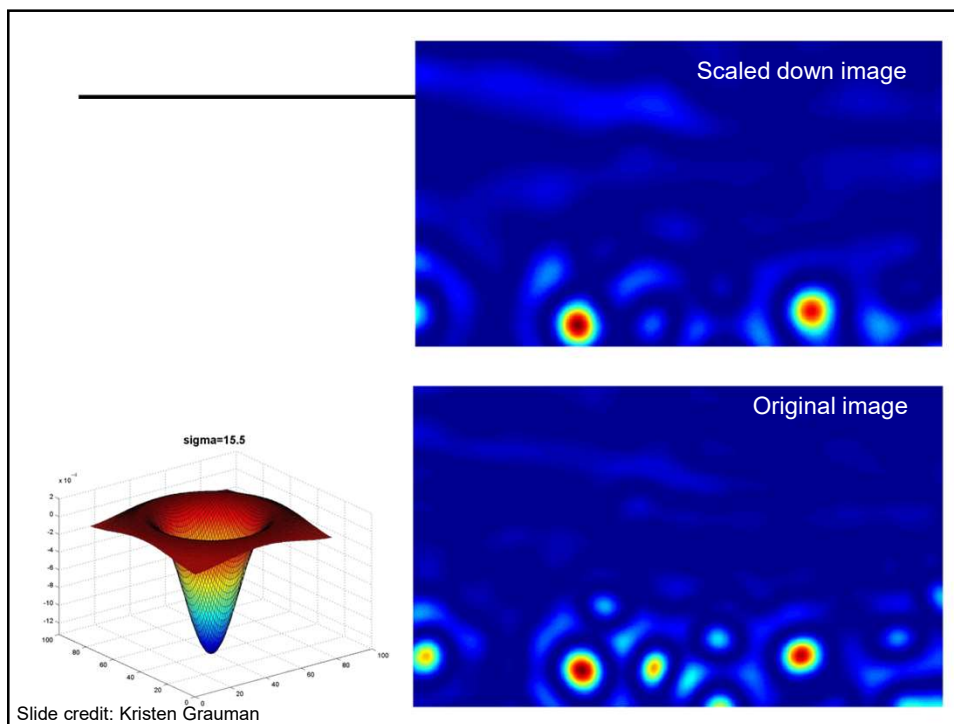
Original image
at $\frac{3}{4}$ the size



Slide credit: Kristen Grauman

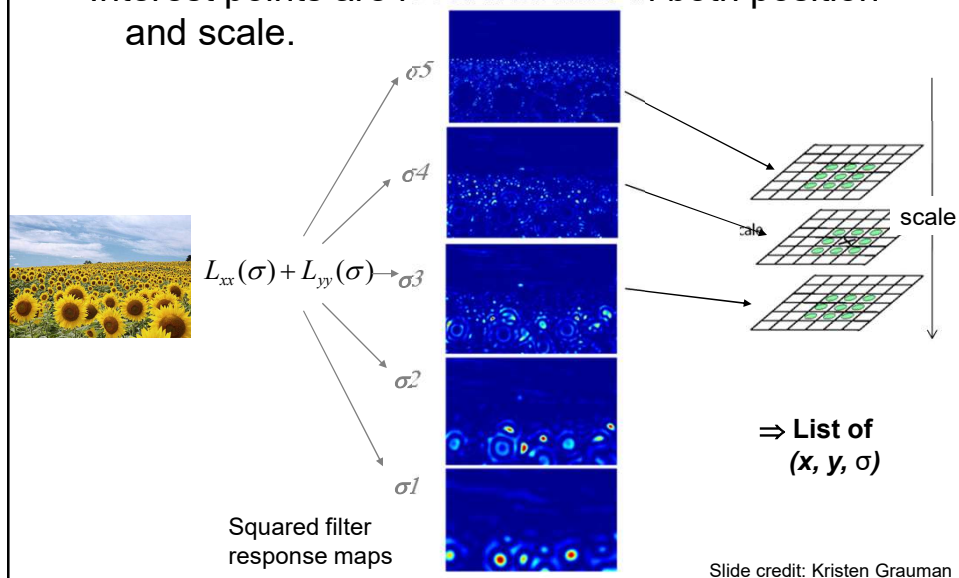






Scale invariant interest points

Interest points are local maxima in both position and scale.




Scale-space blob detector: Example

original image



scale-space maxima of $(\nabla_{norm}^2 L)^2$



T. Lindeberg. Feature detection with automatic scale selection. IJCV 1998.

Scale-space blob detector: Example



Image credit: Lana Lazebnik

Technical detail


We can approximate the Laplacian with a difference of Gaussians; more efficient to implement.

$$L = \sigma^2 (G_{xx}(x, y, \sigma) + G_{yy}(x, y, \sigma))$$

(Laplacian)

$$DoG = G(x, y, k\sigma) - G(x, y, \sigma)$$

(Difference of Gaussians)



Summary


- Desirable properties for local features for correspondence
- Basic matching pipeline
- Interest point detection
 - Harris corner detector
 - Laplacian of Gaussian, automatic scale selection

Local features: main components

- 1) Detection: Identify the interest points

NEXT TIME

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



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

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