CS4670 / 5670: Computer Vision

Noah Snavely

Segmentation

From Sandlot Science

Consumer face application: Apple iPhoto (Picasa, Facebook, etc.)

http://www.apple.com/ilife/iphoto/

Consumer application: Apple iPhoto

Can be trained to recognize pets!

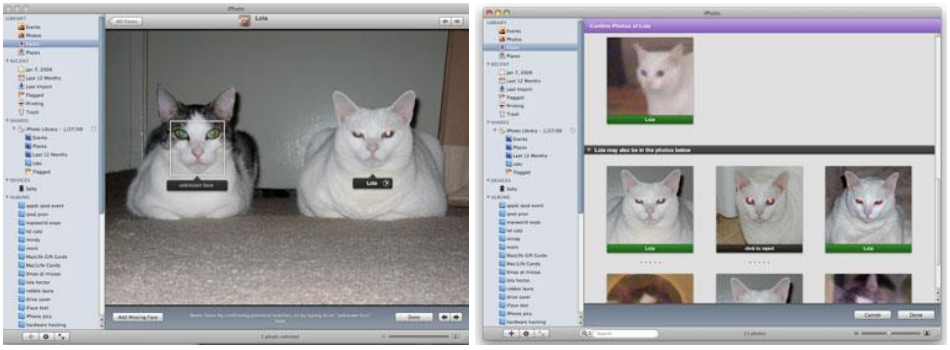

http://www.maclife.com/article/news/iphotos_faces_recognizes_cats

Consumer application: Apple iPhoto

Things iPhoto thinks are faces

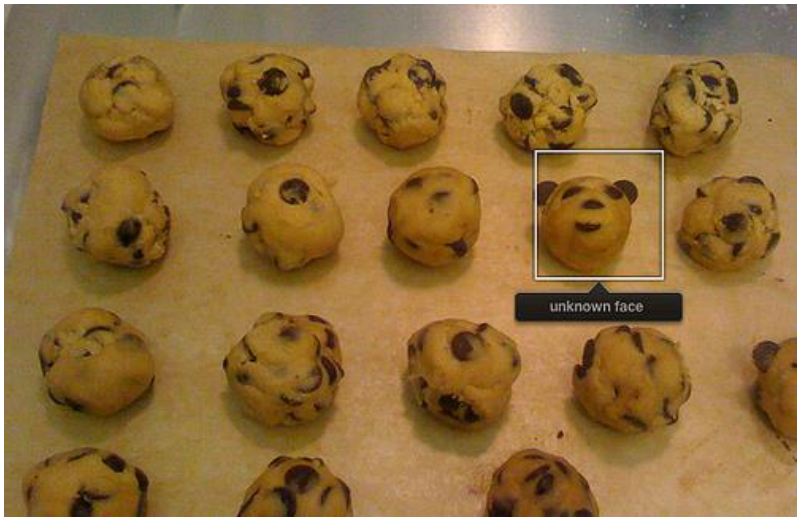

Funny Nikon ads

"The Nikon S60 detects up to 12 faces."

Lambda Labs

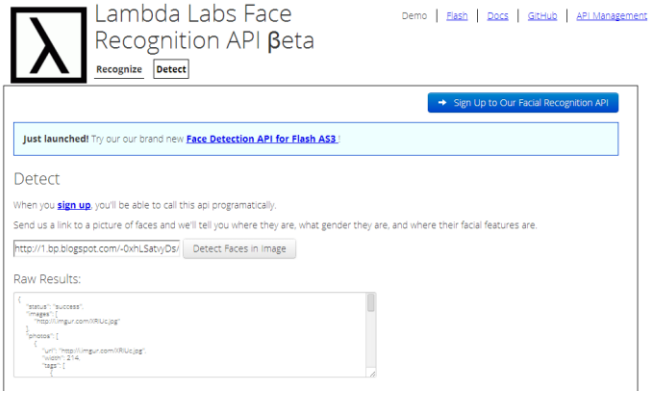

Image segmentation

The goals of segmentation

- Group together similar-looking pixels for efficiency of further processing
  - "Bottom-up" process

"superpixels"

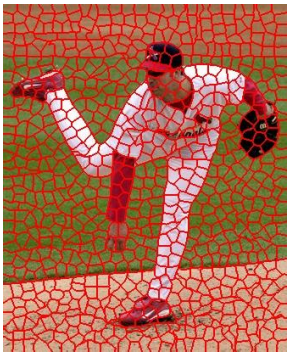

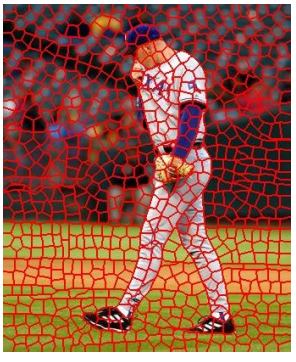

X. Ren and J. Malik. **Learning a classification model for segmentation.** ICCV 2003.

The goals of segmentation

- Separate image into coherent “objects”
  - “Bottom-up” or “top-down” process?

image

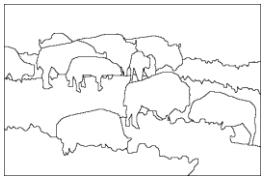

human segmentation

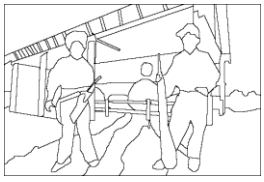

Berkeley segmentation database:
http://www.eecs.berkeley.edu/Research/Projects/CS/vision/grouping/segbench/

The Gestalt school

- Elements in a collection can have properties that result from relationships
  - “The whole is greater than the sum of its parts”

Illusory contours

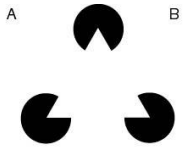

occlusion

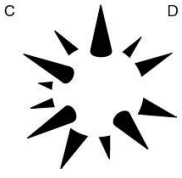

familiar configuration

http://en.wikipedia.org/wiki/Gestalt_psychology

Emergence

http://en.wikipedia.org/wiki/Gestalt_psychology

Gestalt factors

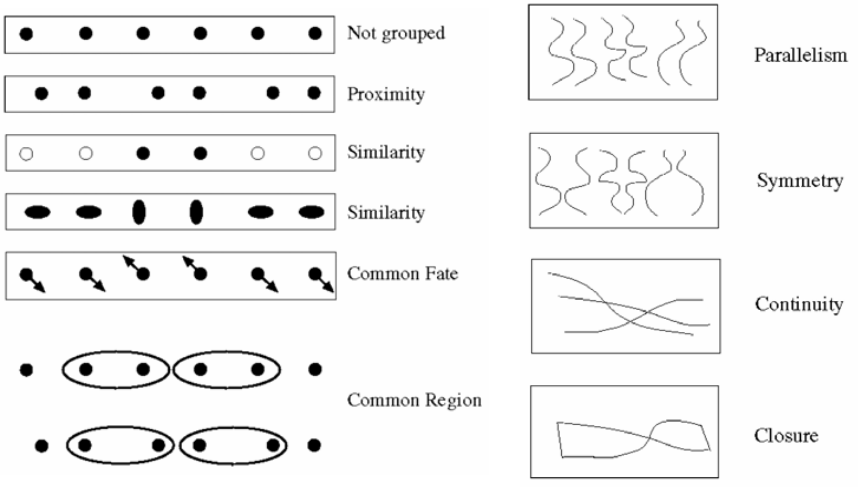

- These factors make intuitive sense, but are very difficult to translate into algorithms

Segmentation as clustering

- Cluster similar pixels (features) together

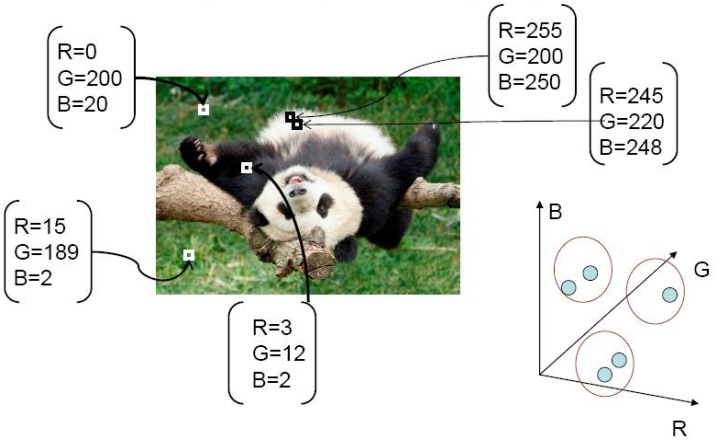

Source: K. Grauman

Segmentation as clustering

- K-means clustering based on intensity or color is essentially vector quantization of the image attributes
  - Clusters don't have to be spatially coherent

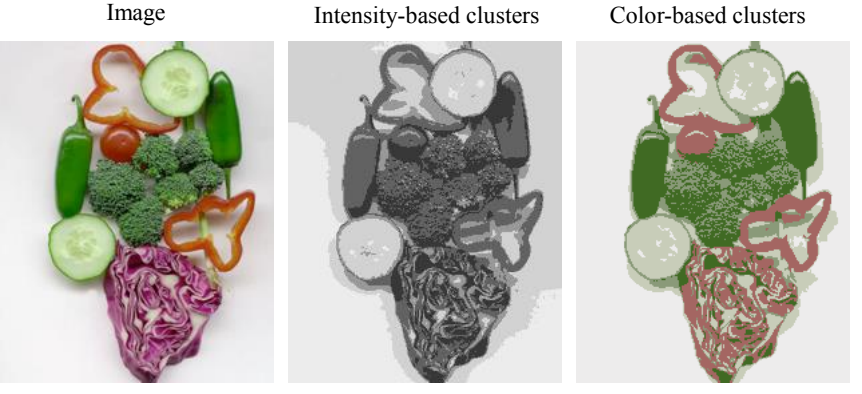

Segmentation as clustering

- Cluster similar pixels (features) together

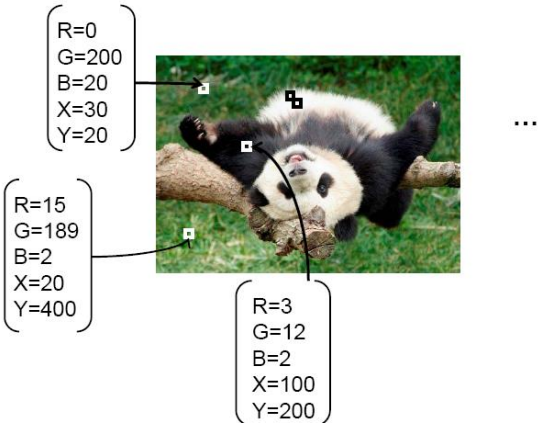

Source: K. Grauman

Segmentation as clustering

- Clustering based on (r,g,b,x,y) values enforces more spatial coherence

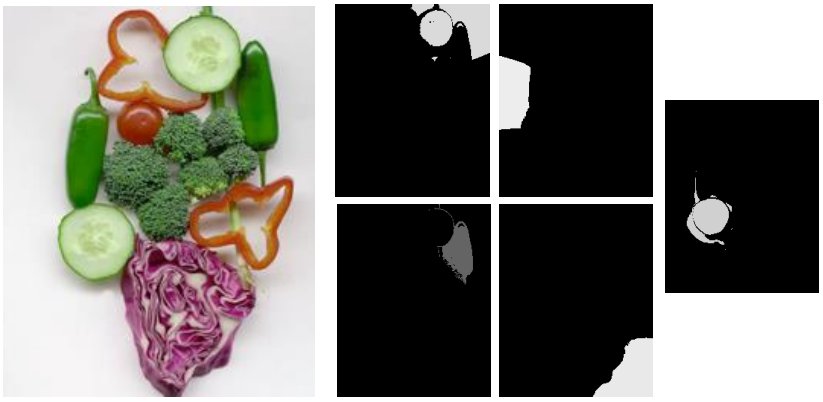

K-Means for segmentation

- Pros
  - Very simple method
  - Converges to a local minimum of the error function
- Cons
  - Memory-intensive
  - Need to pick K
  - Sensitive to initialization
  - Sensitive to outliers
  - Only finds “spherical” clusters

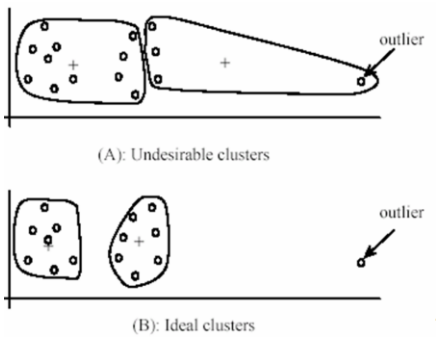

Mean shift clustering and segmentation

- An advanced and versatile technique for clustering-based segmentation

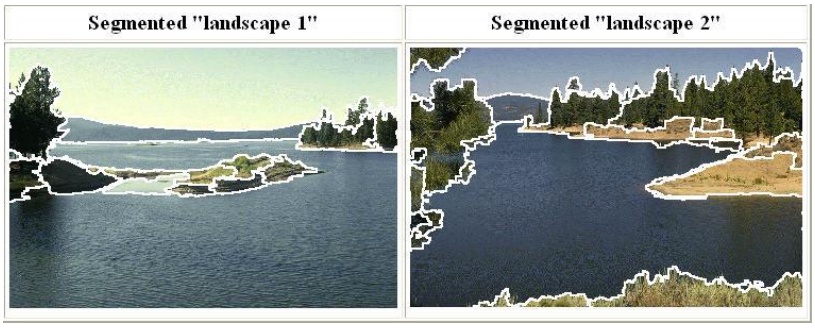

http://www.caip.rutgers.edu/~comanici/MSPAMI/msPamiResults.html

D. Comaniciu and P. Meer, Mean Shift: A Robust Approach toward Feature Space Analysis, PAMI 2002.

Mean shift algorithm

- The mean shift algorithm seeks *modes* or local maxima of density in the feature space

image

Feature space (L*u*v* color values)

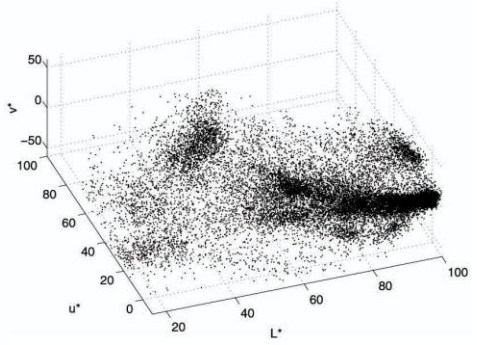

Mean shift

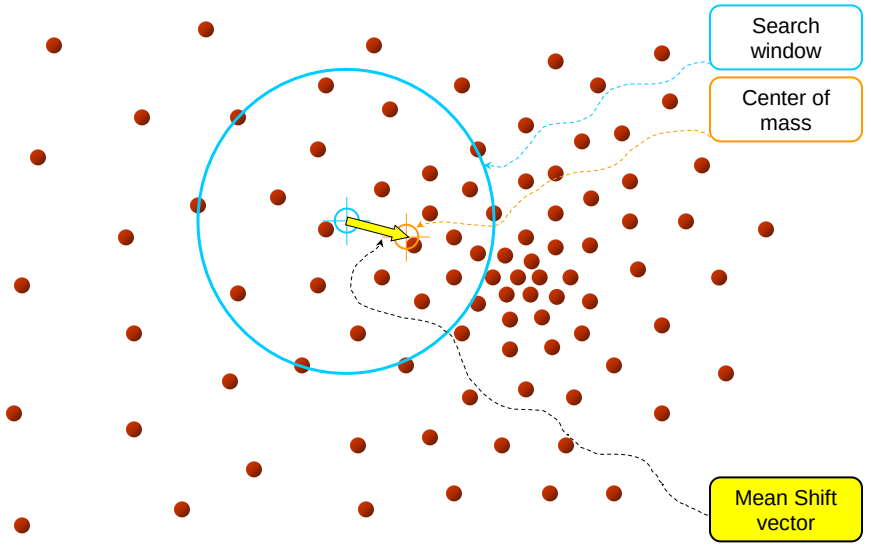

Slide by Y. Ukrainitz & B. Sarel

Mean shift

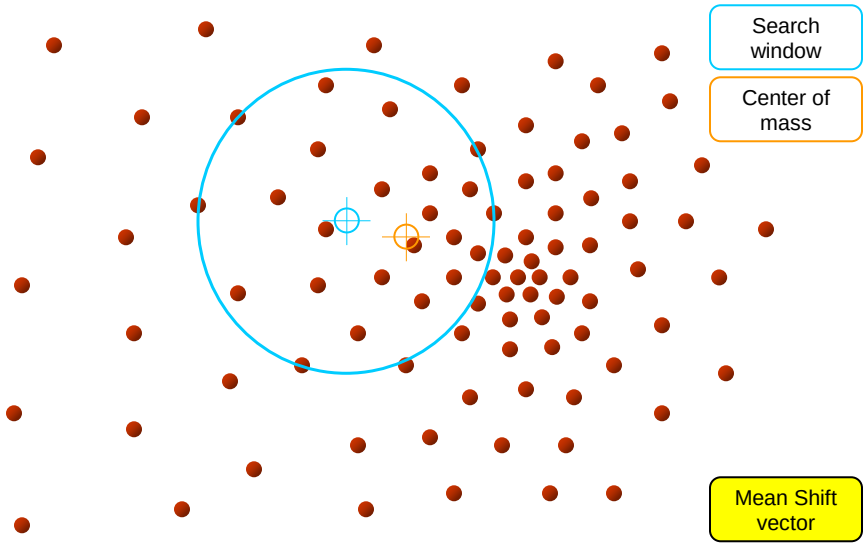

Slide by Y. Ukrainitz & B. Sarel

Mean shift

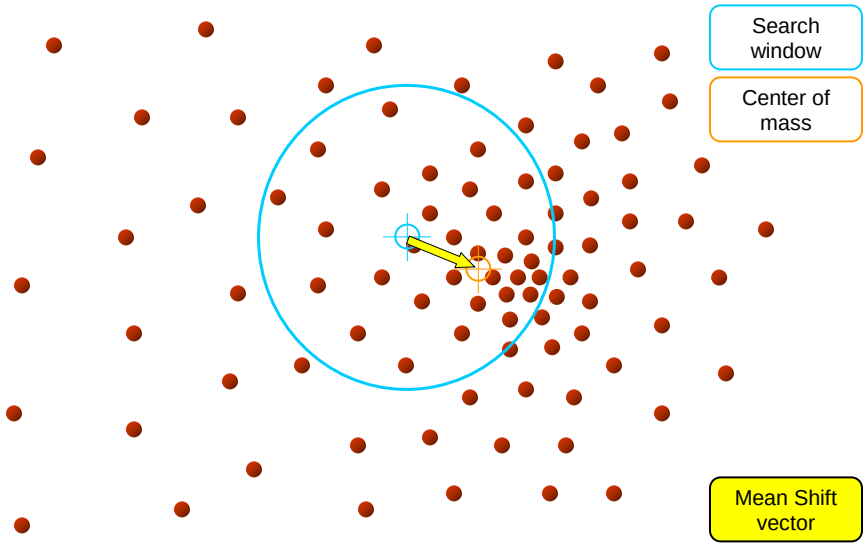

Slide by Y. Ukrainitz & B. Sarel

Mean shift

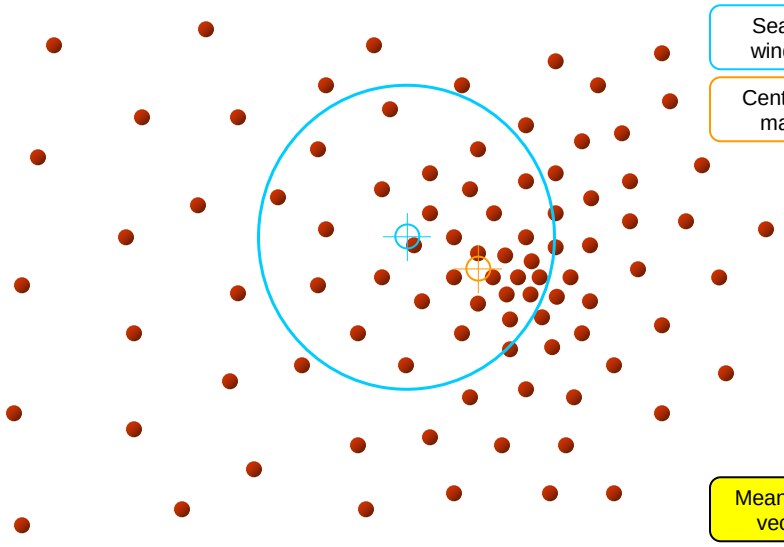

Slide by Y. Ukrainitz & B. Sarel

Mean shift

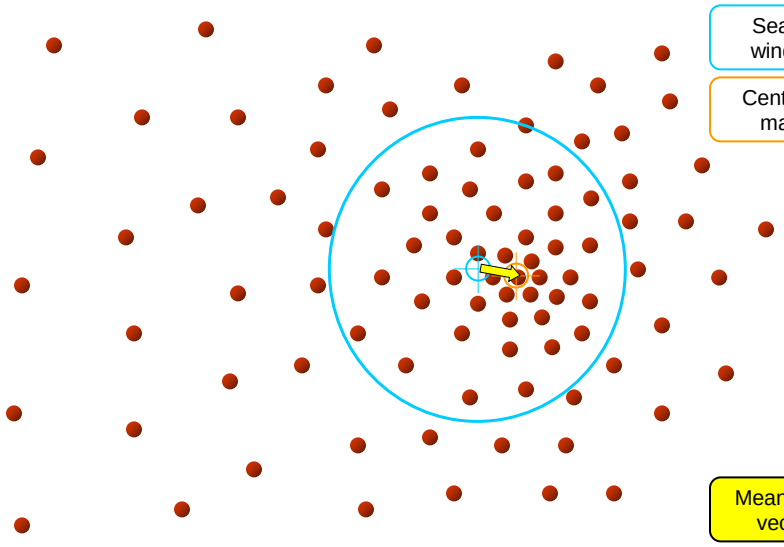

Slide by Y. Ukrainitz & B. Sarel

Mean shift

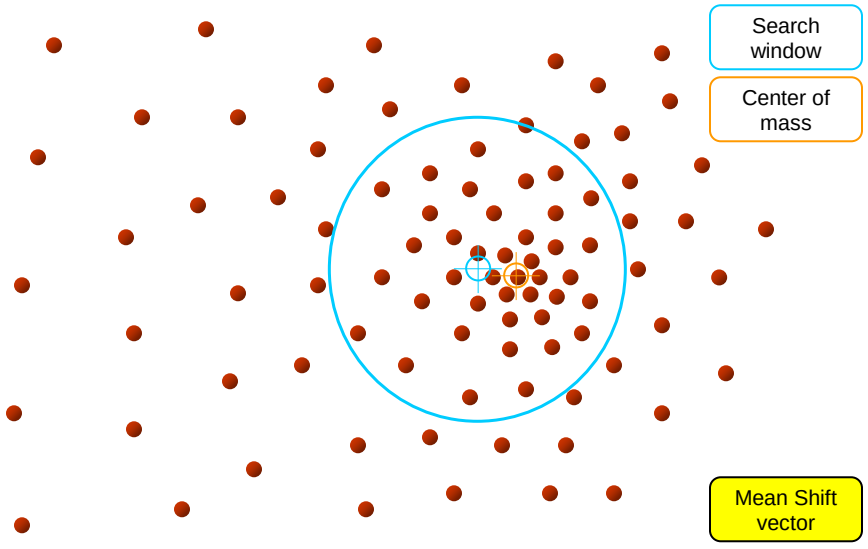

Slide by Y. Ukrainitz & B. Sarel

Mean shift

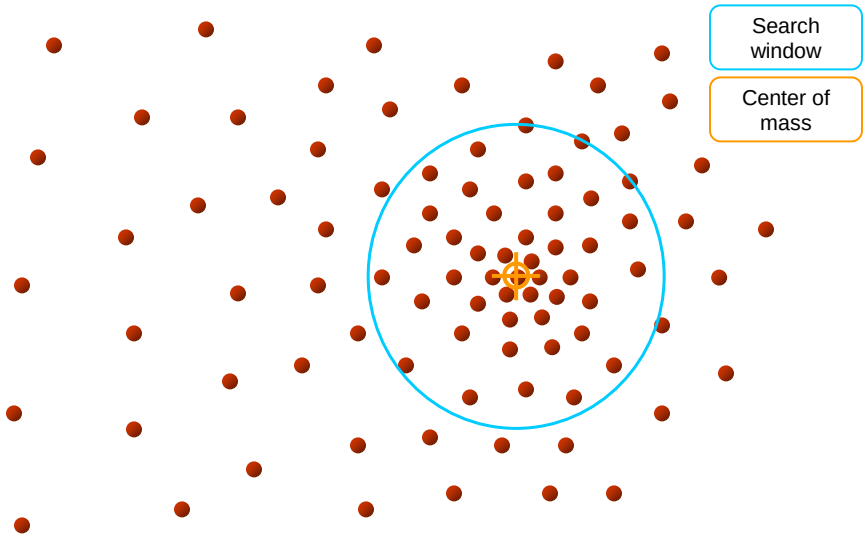

Slide by Y. Ukrainitz & B. Sarel

Mean shift clustering

- Cluster: all data points in the attraction basin of a mode
- Attraction basin: the region for which all trajectories lead to the same mode

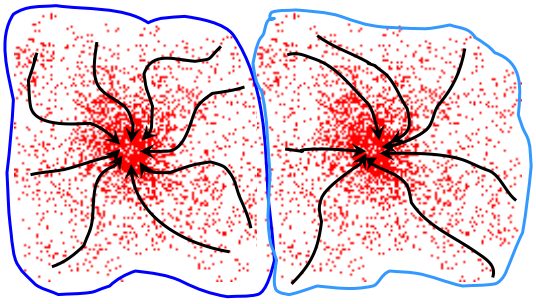

Slide by Y. Ukrainitz & B. Sarel

Mean shift clustering/segmentation

- Find features (color, gradients, texture, etc)
- Initialize windows at individual feature points
- Perform mean shift for each window until convergence
- Merge windows that end up near the same “peak” or mode

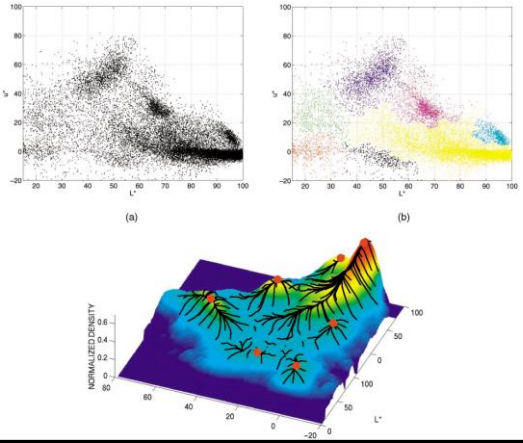

Mean shift segmentation results

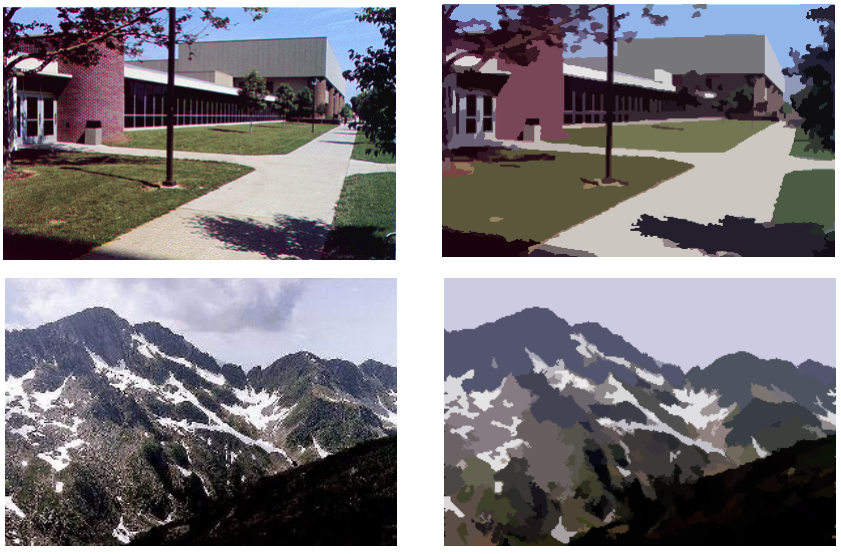

http://www.caip.rutgers.edu/~comanici/MSPAMI/msPamiResults.html

More results

More results

Mean shift pros and cons

- Pros
  - Does not assume spherical clusters
  - Just a single parameter (window size)
  - Finds variable number of modes
  - Robust to outliers
- Cons
  - Output depends on window size
  - Computationally expensive
  - Does not scale well with dimension of feature space