Segmentation by Mean-Shift Clustering

Robert B. Fisher
School of Informatics
University of Edinburgh

Slides credit:Bob Fisher & Vittorio Ferrari

Mean-Shift Segmentation

- An advanced and versatile technique for clustering-based segmentation

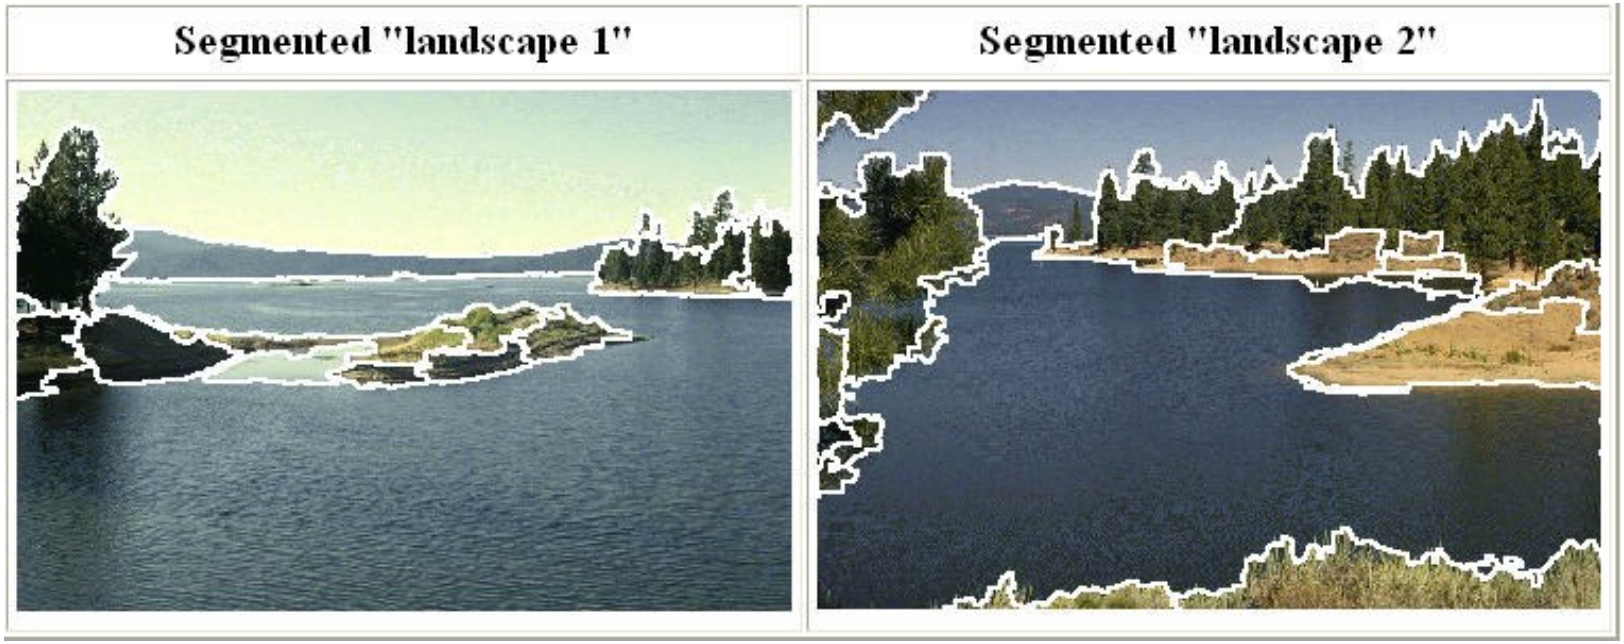

http://www.caip.rutgers.edu/~comanici/MSPAMI/msPamiResults.html

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

Slide credit: Svetlana Lazebnik

Finding Modes in a Histogram

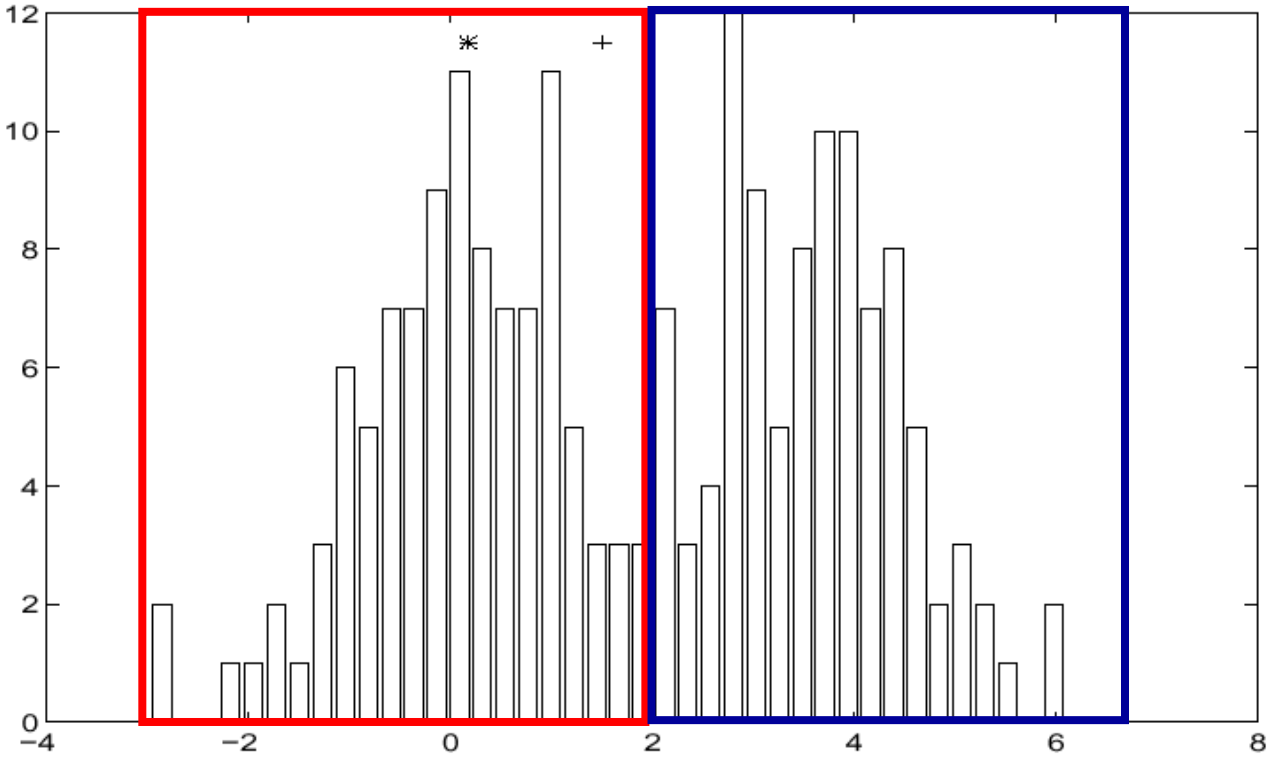

- How many modes are there?
  - *Mode* = local maximum of a given distribution
  - Easy to see, hard to compute

Slide adapted from Steve Seitz

Mean-Shift Algorithm

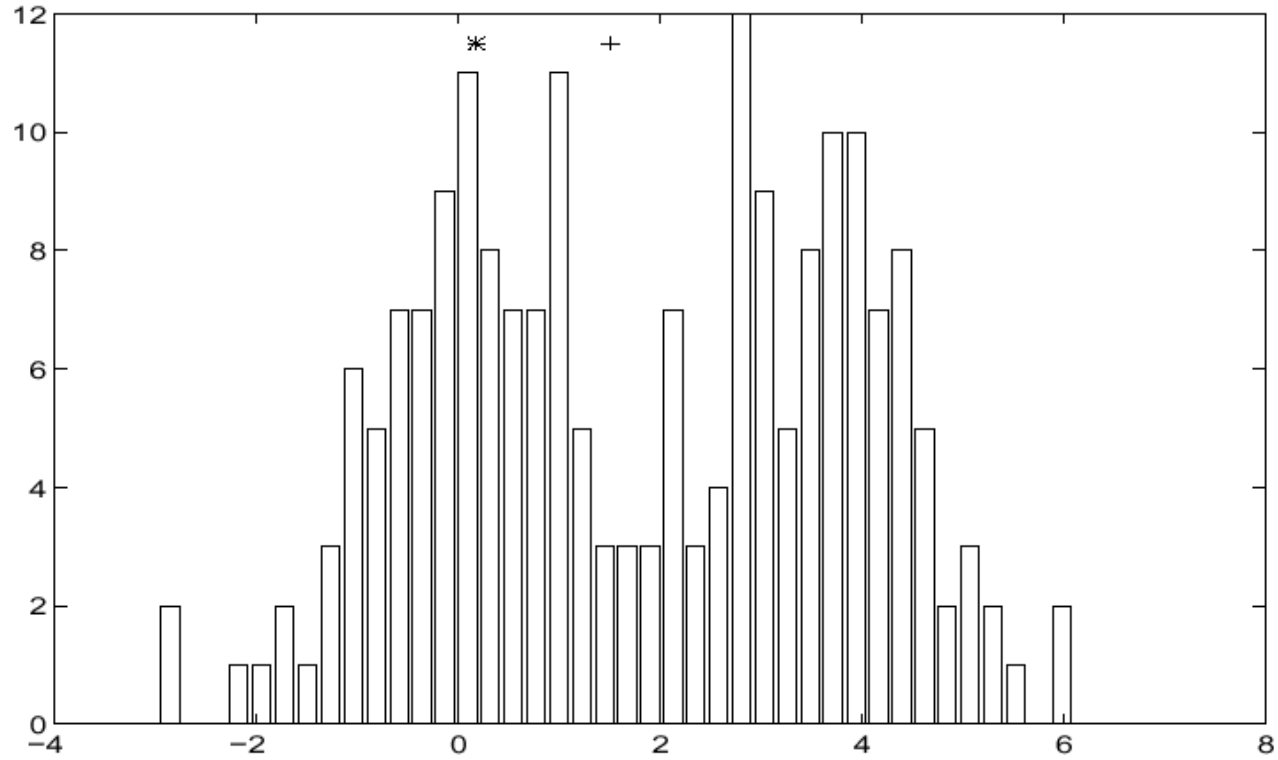

- **Iterative Mode Search**
  1. Initialize random seed center and window W
  2. Calculate center of gravity (the "mean") of W: $\sum_{x \in W} xH(x)$
  3. Shift the search window to the mean
  4. Repeat steps 2+3 until convergence

Slide adapted from Steve Seitz

Mean-Shift Example 1

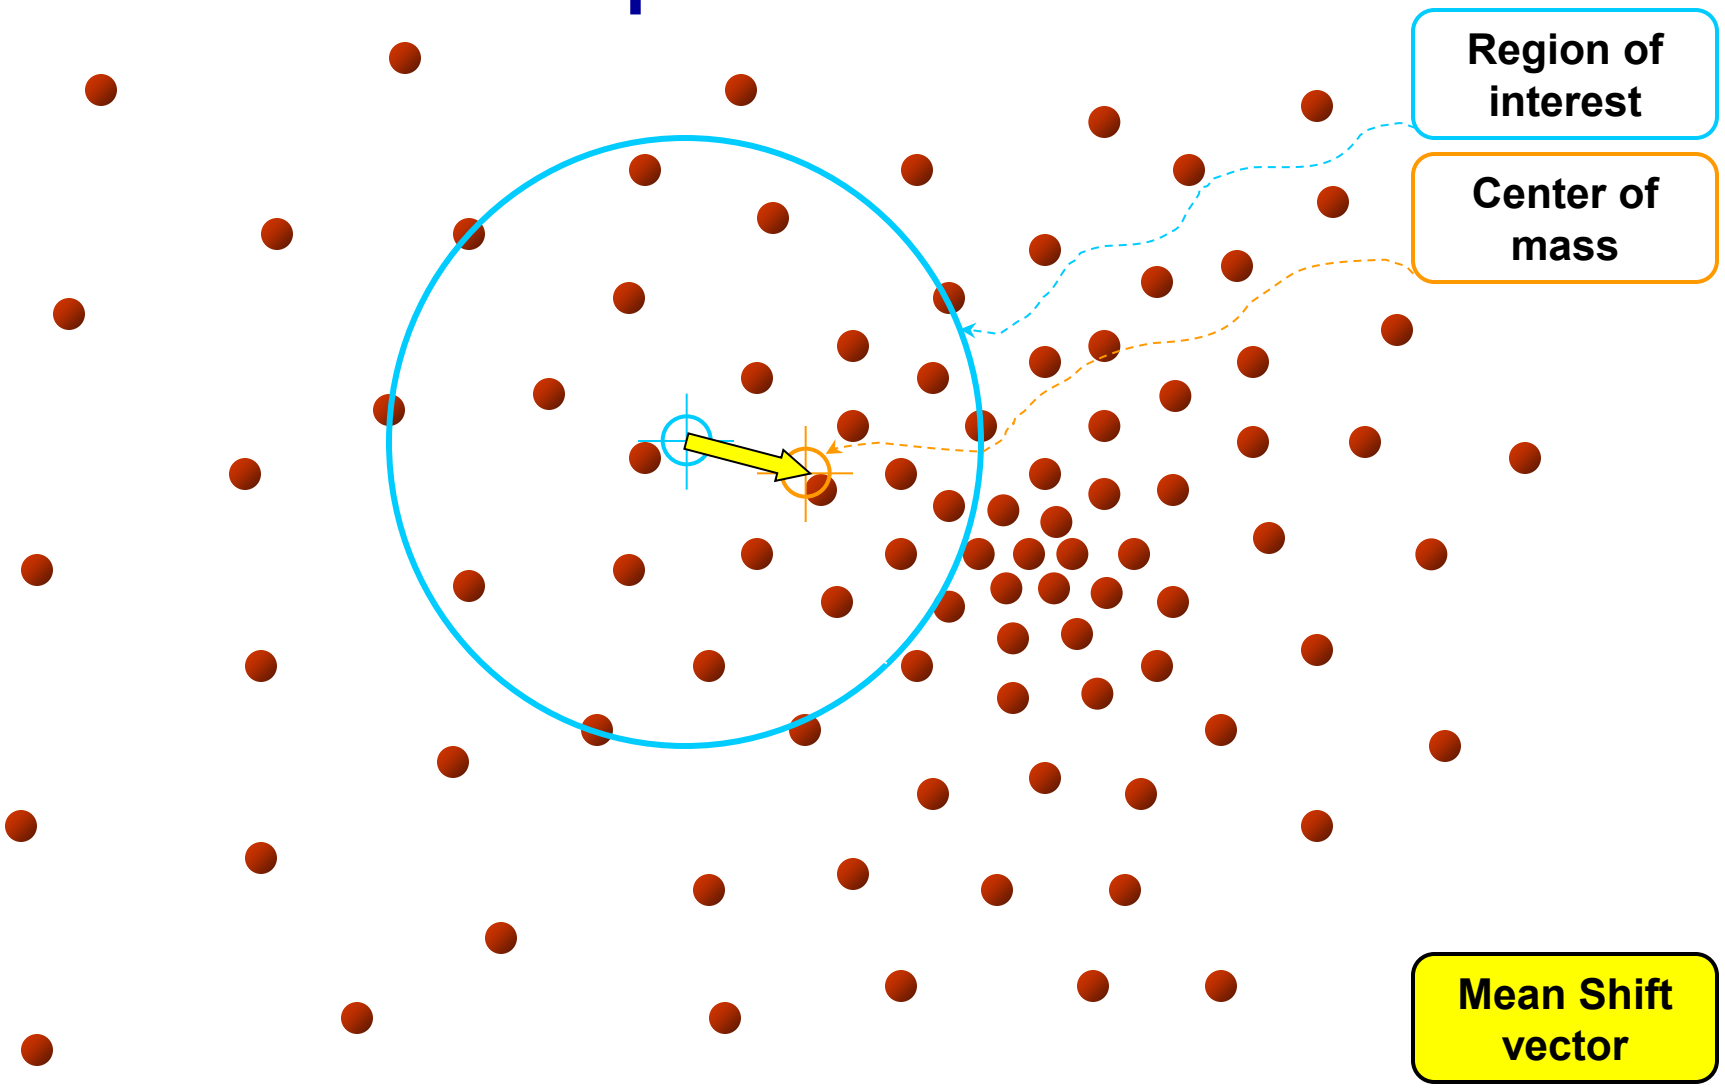

Slide by Y. Ukrainitz & B. Sarel

Mean-Shift Example 2

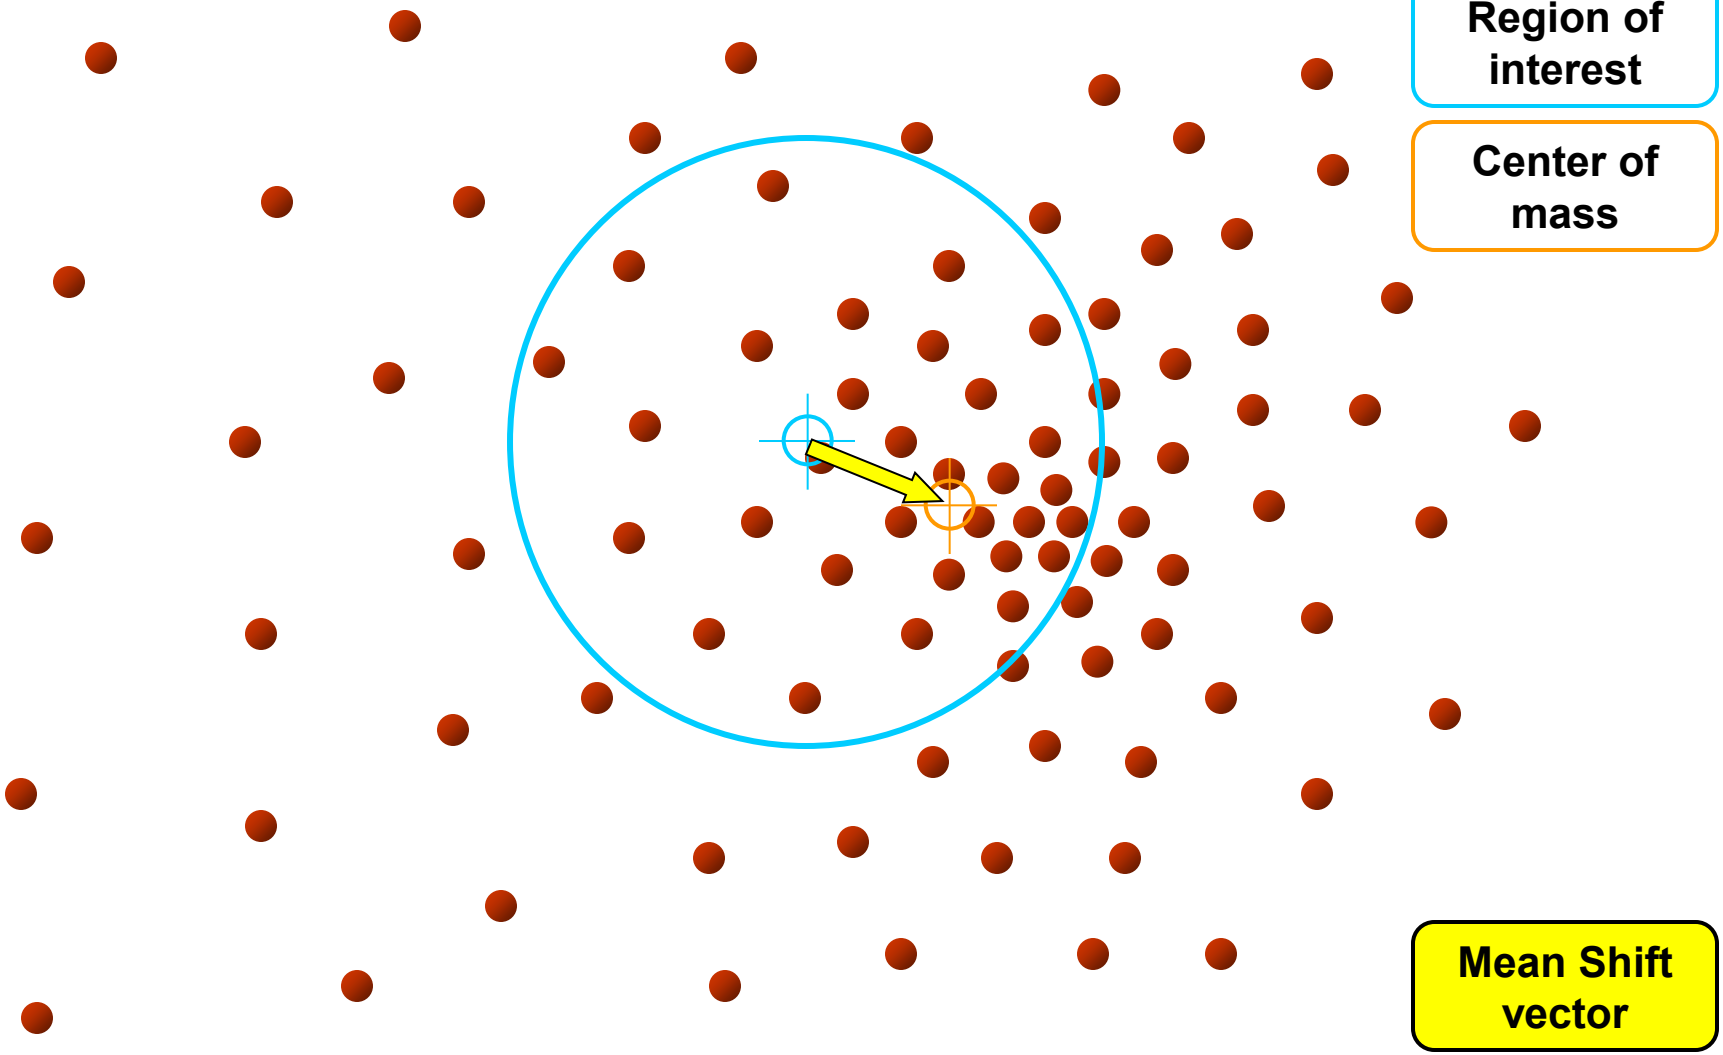

Slide by Y. Ukrainitz & B. Sarel

Mean-Shift Example 3

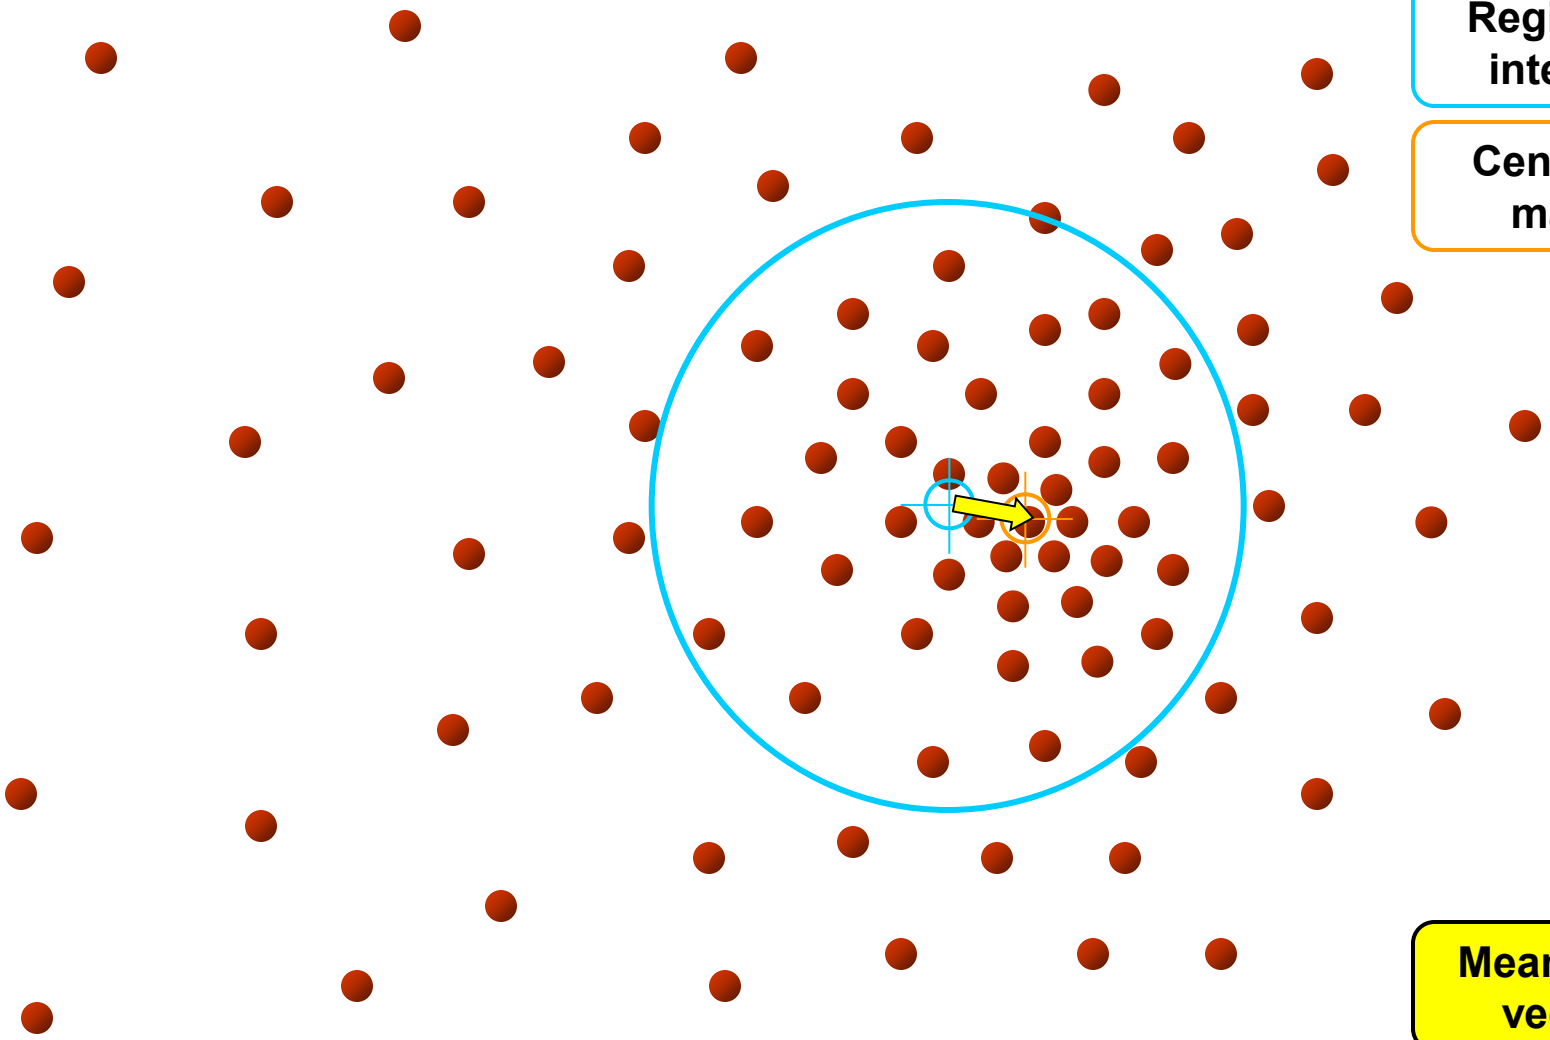

Slide by Y. Ukrainitz & B. Sarel

Mean-Shift Example 4

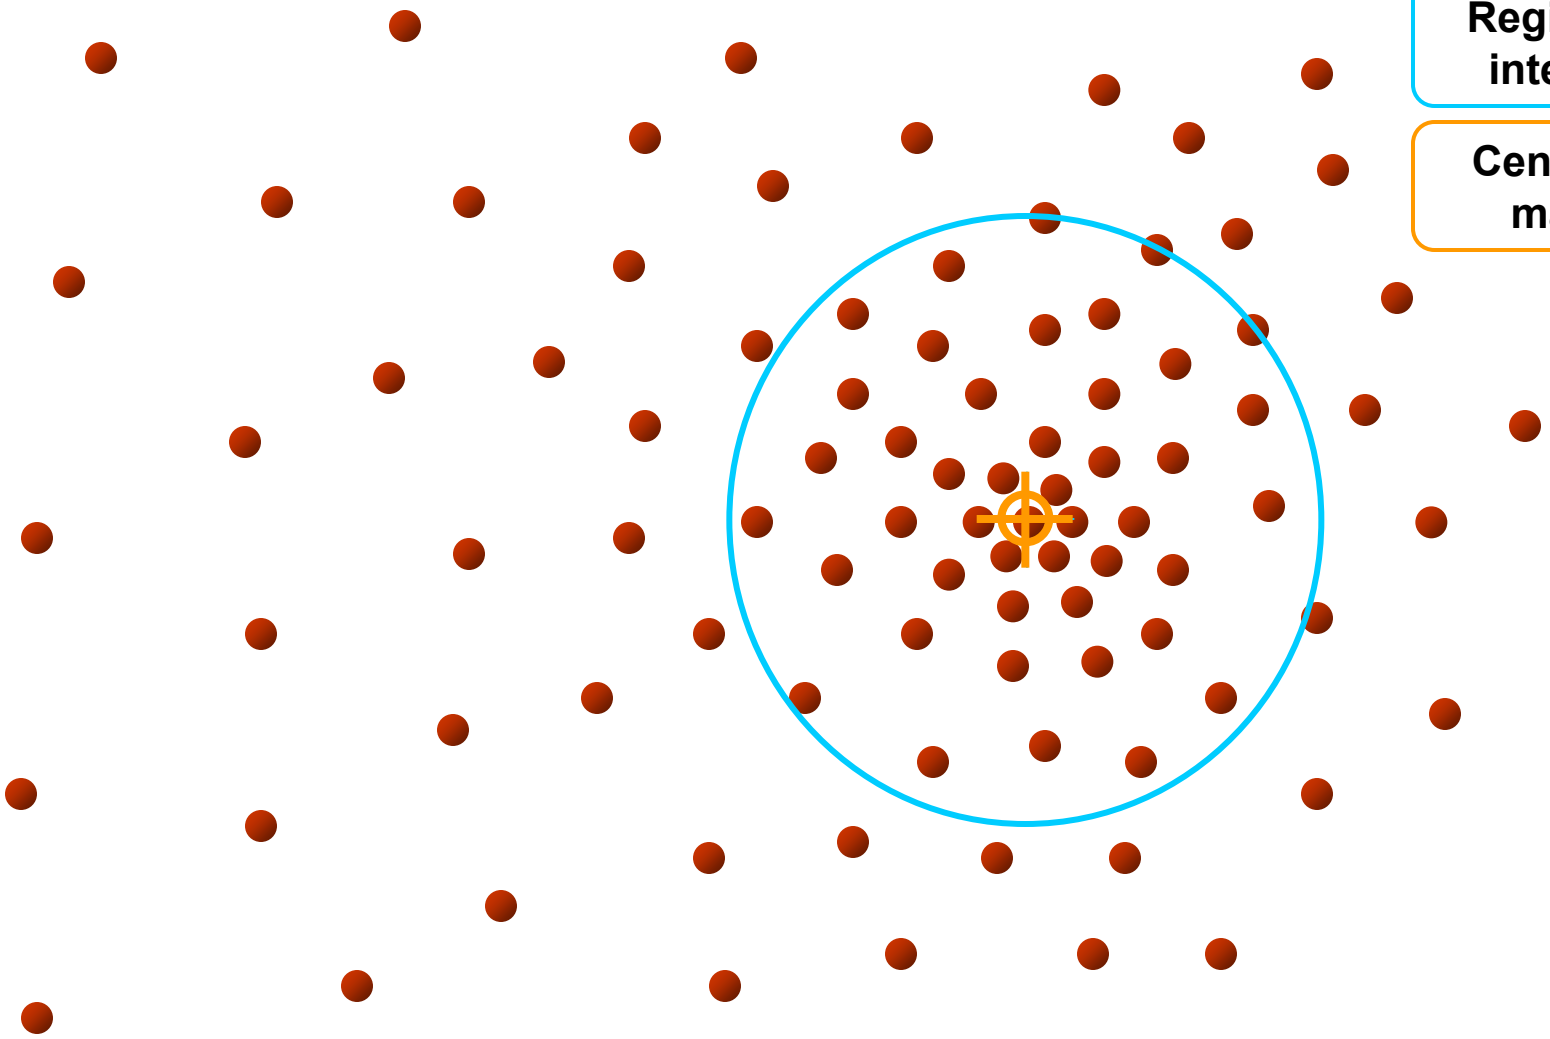

Slide by Y. Ukrainitz & B. Sarel

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
- Start mean-shift at every data point

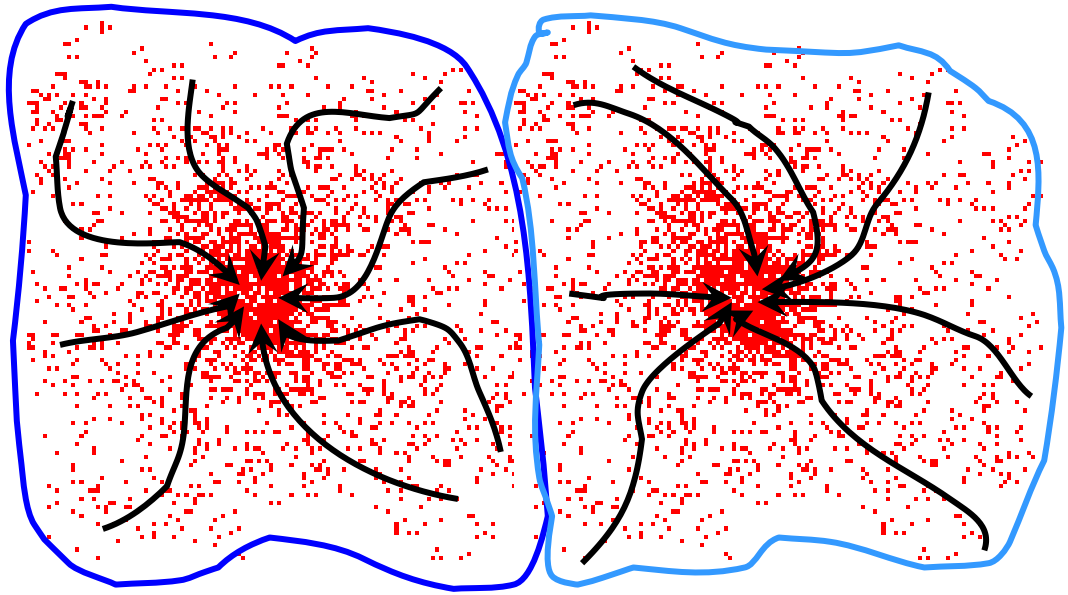

Slide by Y. Ukrainitz & B. Sarel

Mean-Shift Clustering/Segmentation

- Choose features (color, gradients, texture, etc)
- Initialize windows at individual pixel locations
- Start mean-shift from each window until convergence
- Merge windows that end up near the same “peak” or mode

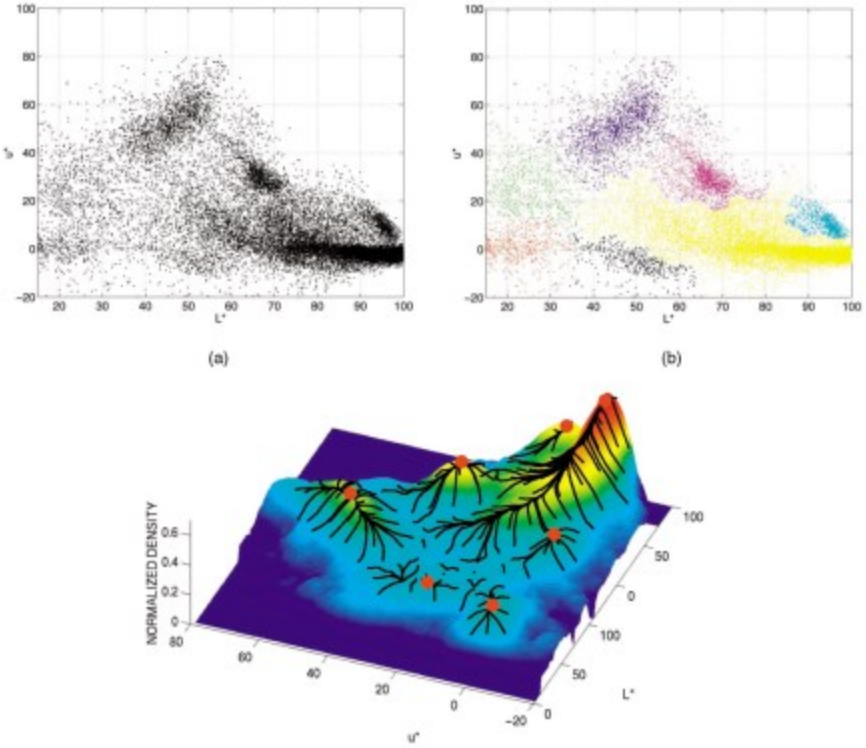

Slide adapted from Svetlana Lazebnik

Segmenting Pixels in Same Colour Clusters

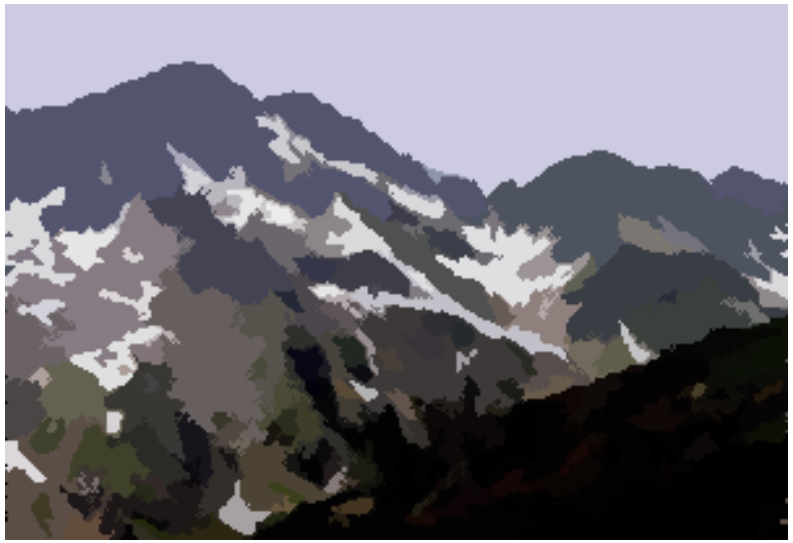

http://www.caip.rutgers.edu/~comanici/MSPAMI/msPamiResults.html

Slide credit: Svetlana Lazebnik

More Examples

Slide credit: Svetlana Lazebnik

Mean-Shift Lecture Overview

- **Pros**
  - General, application-independent tool
  - Model-free, does not assume any prior shape (spherical, elliptical, etc.) on data clusters
  - Just a single parameter (window size h)
    - h has a physical meaning (unlike k-means) == scale of clustering
  - Finds variable number of modes given the same h
  - Robust to outliers

- **Cons**
  - Output depends on window size h
  - Window size (bandwidth) selection is not trivial
  - Computationally rather expensive
  - Does not scale well with dimension of feature space

Slide adapted from Svetlana Lazebnik