Mean Shift

Paper by Comaniciu and Meer
Presentation by Carlo Lopez-Tello

What is the Mean Shift Algorithm?

- A method of finding peaks (modes) in a probability distribution
- Works without assuming any underlying structure in the distribution
- Works on multimodal distributions
- Works without assuming the number of modes

Why do we care about modes?

- Given a data set we can assume that it was sampled from some pdf
- Samples are most likely to be drawn from a region near a mode
- We can use the modes to cluster the data
- Clustering has many applications: filtering, segmentation, tracking, classification, and compression.

Why do we care about modes?

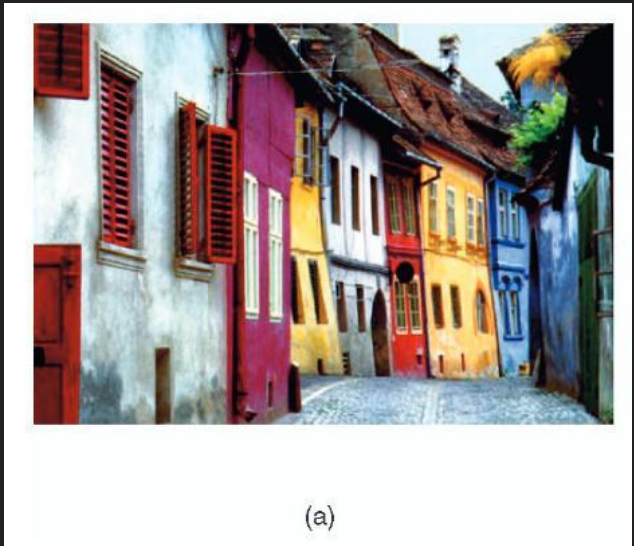

(a)

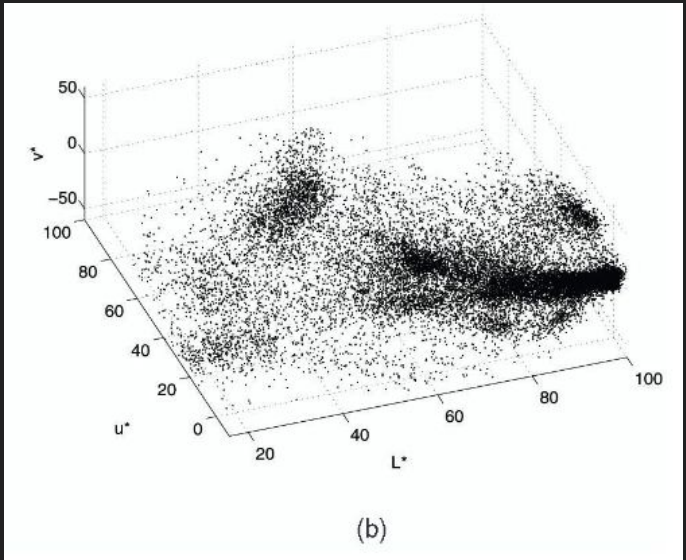

(b)

Fig. 1. Example of a feature space. (a) A 400×276 color image. (b) Corresponding L*u*v* color space with $110,400$ data points.

Why use mean shift for clustering?

- K-means needs to know how many clusters to use. Clusters data into voronoi cells.

- Histograms require bin size and number of bins

- Mixture models require information about pdf structure

Intuition

- We have a set of data that represents discrete samples of a distribution
- Locally we can estimate the density of the distribution with a function
- Compute the gradient of this estimation function
- Use gradient ascent to find the peak of the distribution

How does it work?

- We estimate the density using:

$$\hat{f}_{h,K}(\mathbf{x}) = \frac{c_{k,d}}{nh^d} \sum_{i=1}^{n} k\left(\left\| \frac{\mathbf{x} - \mathbf{x}_i}{h} \right\|^2 \right).$$

- Where h (bandwidth) is the region around x where we are trying to estimate the density and k is some kernel function

- Instead of using the gradient of f, we use the mean shift vector:

$$\mathbf{m}_{h,G}(\mathbf{x}) = \frac{1}{2} h^2 c \frac{\hat{\nabla} f_{h,K}(\mathbf{x})}{\hat{f}_{h,G}(\mathbf{x})}.$$

How to find a mode?

1. Start at any point
2. Compute mean shift
3. if mean shift is zero: possible mode found
4. else move to where mean shift is pointing

   go to 2

- To find multiple modes we need to try all points that are more than h distance apart
- Prune modes by perturbing them and checking for convergence
- Combine modes that are close together. Take the higher one.

How to cluster using mean shift?

- Every point in the data set will converge to some mode using mean shift
- We cluster points together if they converge to the same mode

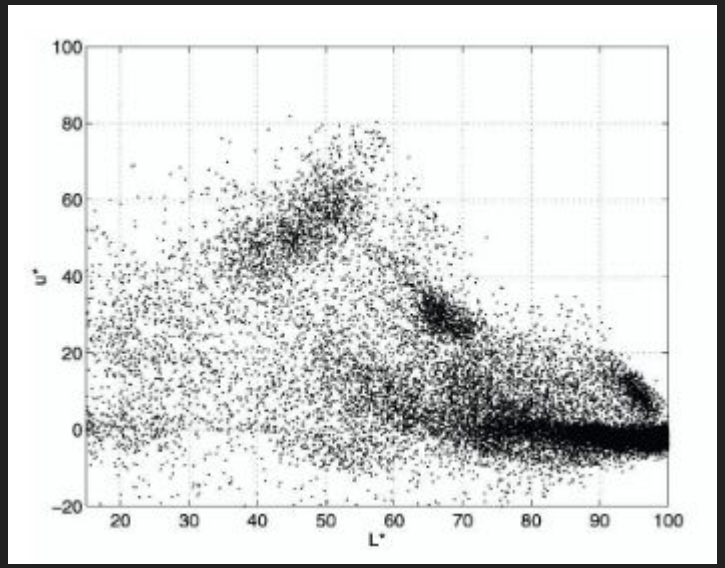

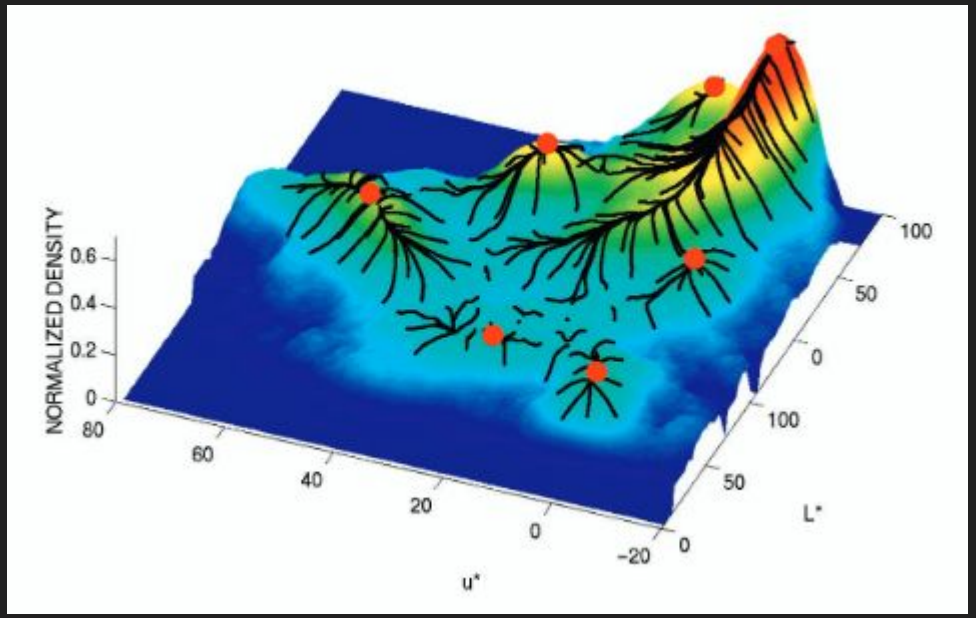

Mean Shift Filtering

- Cluster using intensity and position. Then change the intensity to match the cluster.

(a) (b)

Fig. 3. *Cameraman* image. (a) Original. (b) Mean shift filtered $(h_s, h_r) = (8, 4)$.

Mean Shift Segmentation

- Cluster using intensity and spatial information. Each cluster represents a segment of the image.

Discussion

- Need to select bandwidth and kernel function
- Gaussian kernel performs better, but takes longer to converge
- Kernel density estimation does not scale well with the dimension of the space.

Questions