Clustering Algorithms

CS345a: Data Mining
Jure Leskovec and Anand Rajaraman
Stanford University

Problem Statement

- Given a set of data points, group them into a clusters so that:
  - points within each cluster are similar to each other
  - points from different clusters are dissimilar
- Usually, points are in a high-dimensional space, and similarity is defined using a distance measure
  - Euclidean, Cosine, Jaccard, edit distance, …

Example: Doggie Data

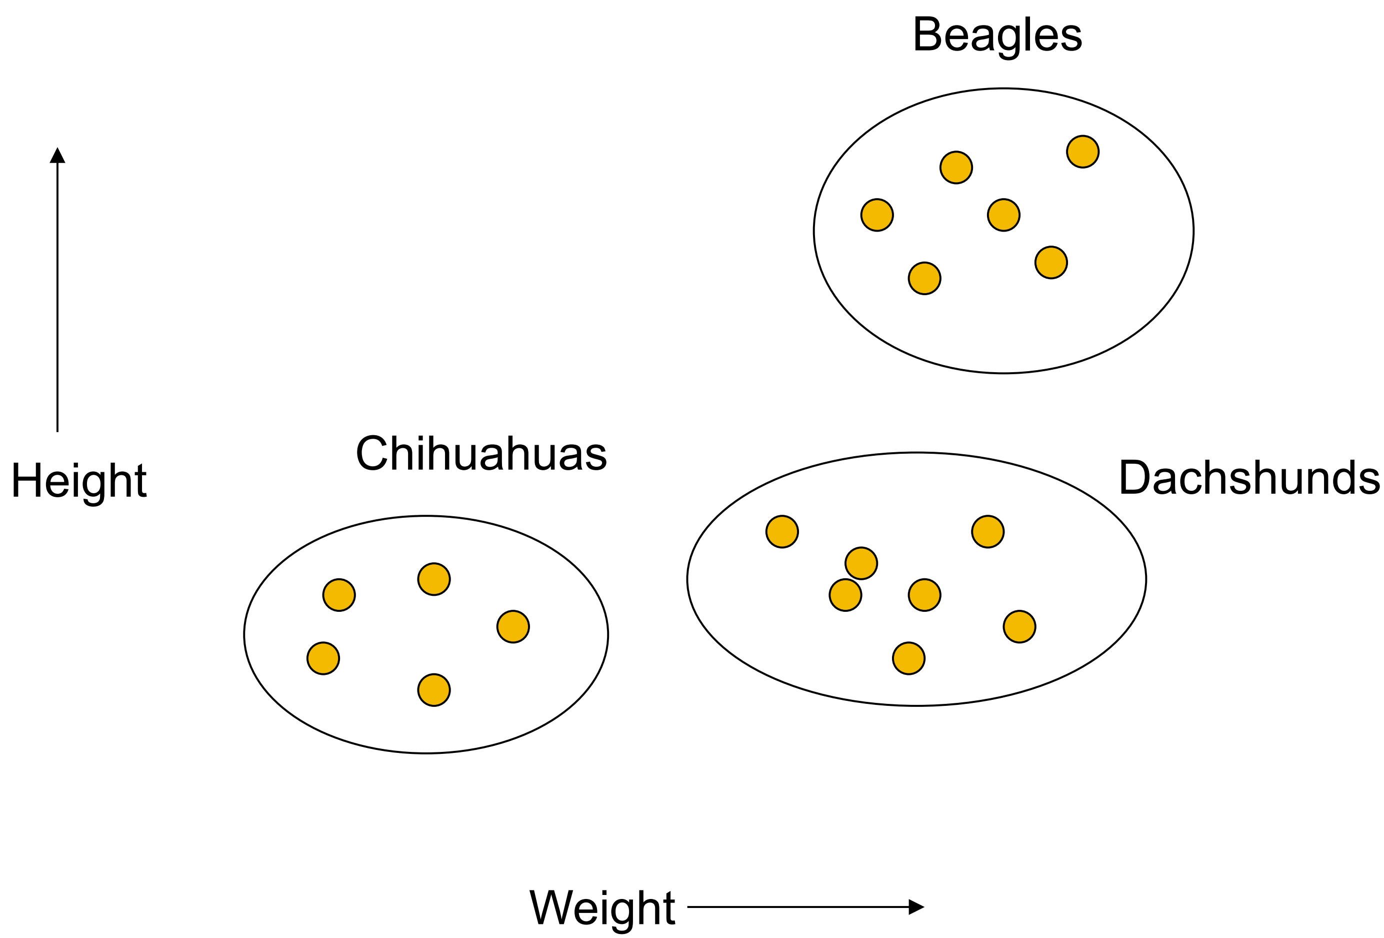

Application Example: SkyCat

- A catalog of 2 billion "sky objects" represents objects by their radiation in 7 dimensions (frequency bands).
- Problem: cluster into similar objects, e.g., galaxies, nearby stars, quasars, etc.
- Sloan Sky Survey is a newer, better version.

More Examples

- Cluster customers based on their purchase histories
- Cluster products based on the sets of customers who purchased them
- Cluster documents based on similar words or shingles
- Cluster DNA sequences based on edit distance

Methods of Clustering

- Hierarchical (Agglomerative):
  - Initially, each point in cluster by itself.
  - Repeatedly combine the two "nearest" clusters into one.
- Point Assignment:
  - Maintain a set of clusters.
  - Place points into their "nearest" cluster.

Hierarchical Clustering

- Key Operation: repeatedly combine two nearest clusters
- Three important questions:
  - How do you represent a cluster of more than one point?
  - How do you determine the "nearness" of clusters?
  - When to stop combining clusters?

Euclidean Case

- Each cluster has a well-defined centroid
  - i.e., average across all the points in the cluster
- Represent each cluster by its centroid
- Distance between clusters = distance between centroids

Example

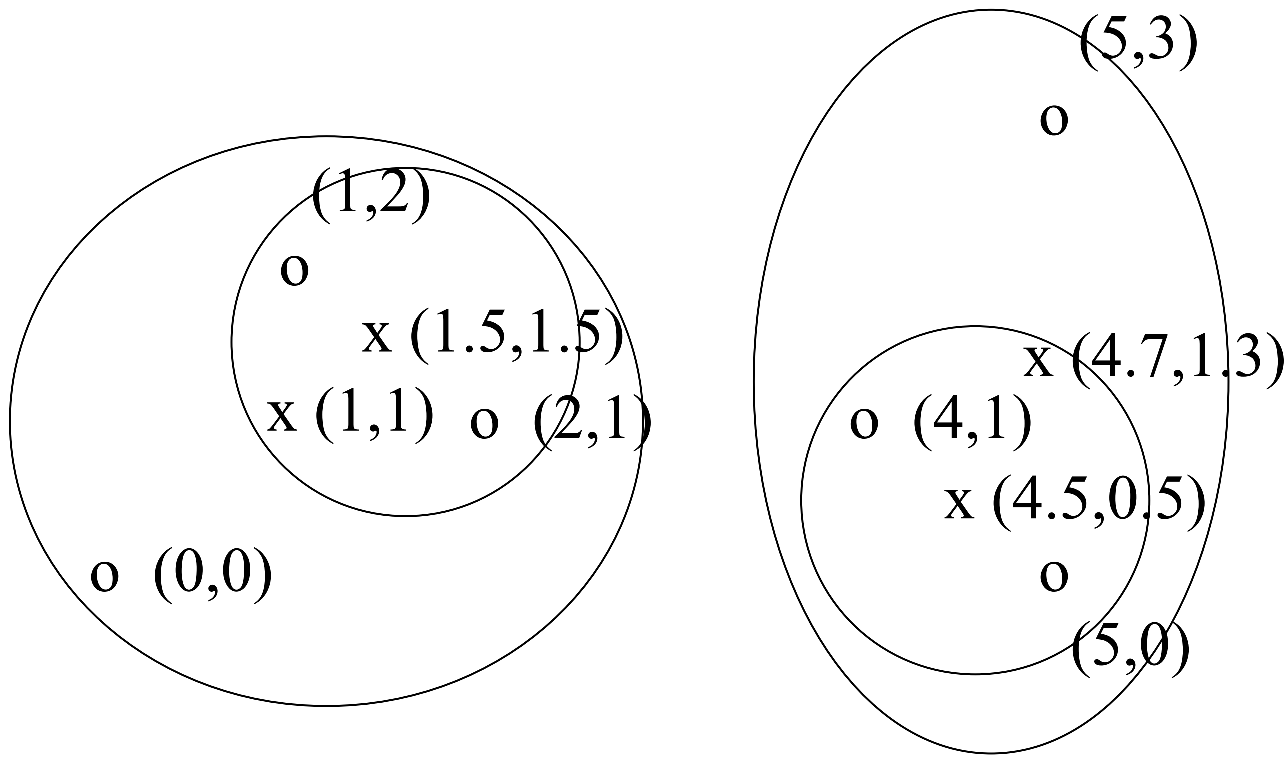

Non-Euclidean Distances

- The only "locations" we can talk about are the points themselves.
  - I.e., there is no "average" of two points.
- Approach 1: *clustroid* = point "closest" to other points.
  - Treat clustroid as if it were centroid, when computing intercluster distances.

“Closest” Point?

Possible meanings:

1. Smallest maximum distance to the other points.
2. Smallest average distance to other points.
3. Smallest sum of squares of distances to other points.
4. Etc., etc.

Example

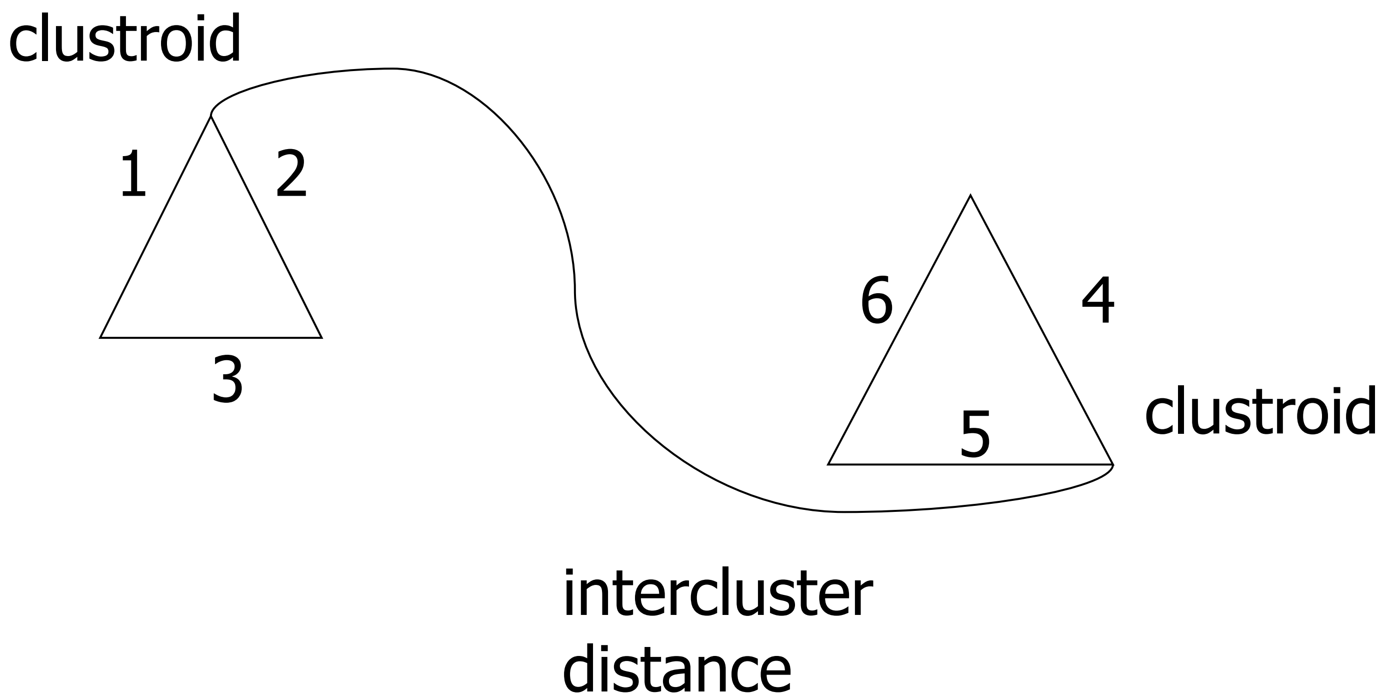

Other Approaches

- Approach 2: intercluster distance = minimum of the distances between any two points, one from each cluster.
- Approach 3: Pick a notion of "cohesion" of clusters, e.g., maximum distance from the clustroid.
  - Merge clusters whose *union* is most cohesive.

Cohesion

- Approach 1: Use the *diameter* of the merged cluster = maximum distance between points in the cluster.
- Approach 2: Use the average distance between points in the cluster.

Cohesion – (2)

- Approach 3: Use a density-based approach: take the diameter or average distance, e.g., and divide by the number of points in the cluster.
  - Perhaps raise the number of points to a power first, e.g., square-root.

Stopping Criteria

- Stop when we have k clusters
- Stop when the cohesion of the cluster resulting from the best merger falls below a threshold
- Stop when there is a sudden jump in the cohesion value

Implementing Hierarchical Clustering

- Naïve implementation:
  - At each step, compute pairwise distances between each pair of clusters
  - $O(N^3)$
- Careful implementation using a priority queue can reduce time to $O(N^2 \log N)$
- Too expensive for really big data sets that don't fit in memory

k – Means Algorithm(s)

- Assumes Euclidean space.
- Start by picking *k*, the number of clusters.
- Initialize clusters by picking one point per cluster.
  - Example: pick one point at random, then *k* -1 other points, each as far away as possible from the previous points.

Populating Clusters

1. For each point, place it in the cluster whose current centroid it is nearest, and update the centroid of the cluster.
2. After all points are assigned, fix the centroids of the *k* clusters.
3. Optional: reassign all points to their closest centroid.
   - Sometimes moves points between clusters.

Example: Assigning Clusters

Reassigned
points

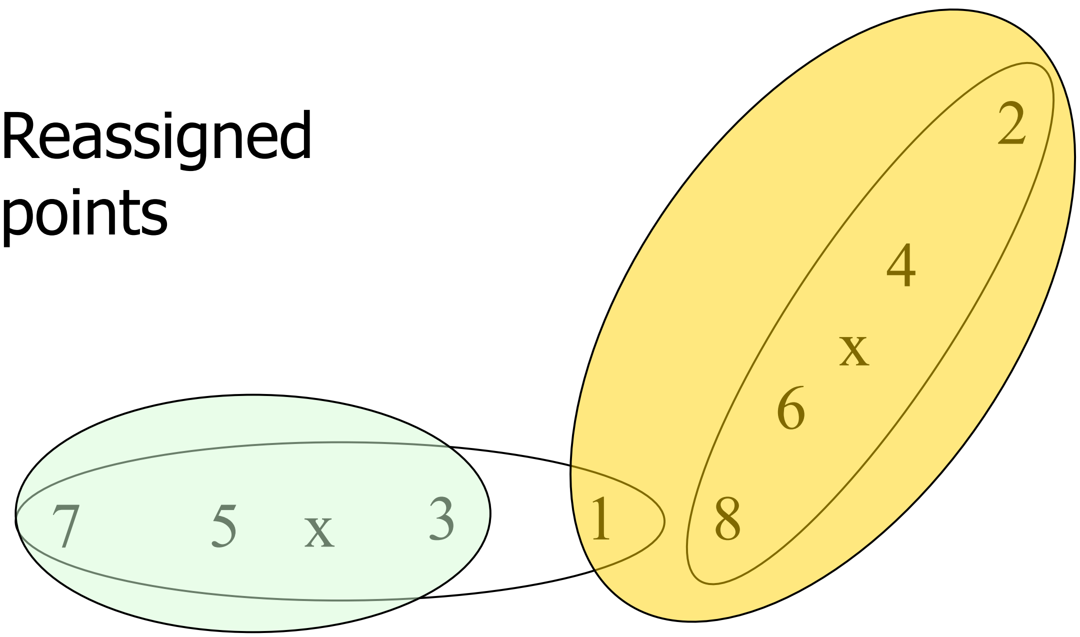

Clusters after first round

Getting *k* Right

- Try different *k*, looking at the change in the average distance to centroid, as *k* increases.
- Average falls rapidly until right *k*, then changes little.

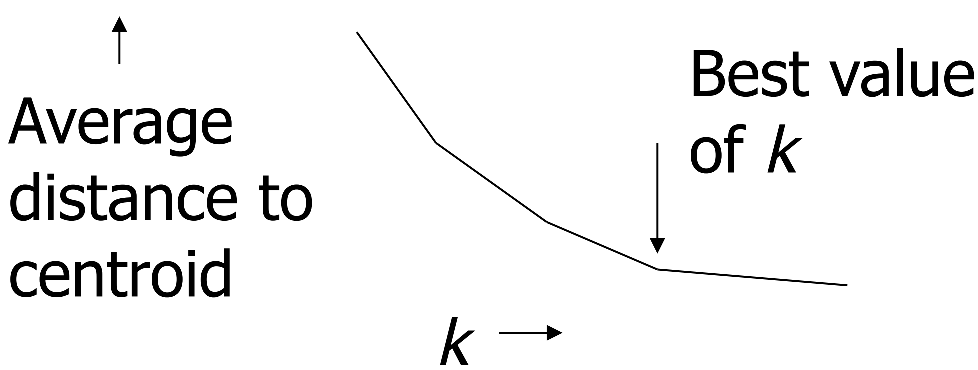

Example: Picking *k*

Too few;
many long
distances
to centroid.

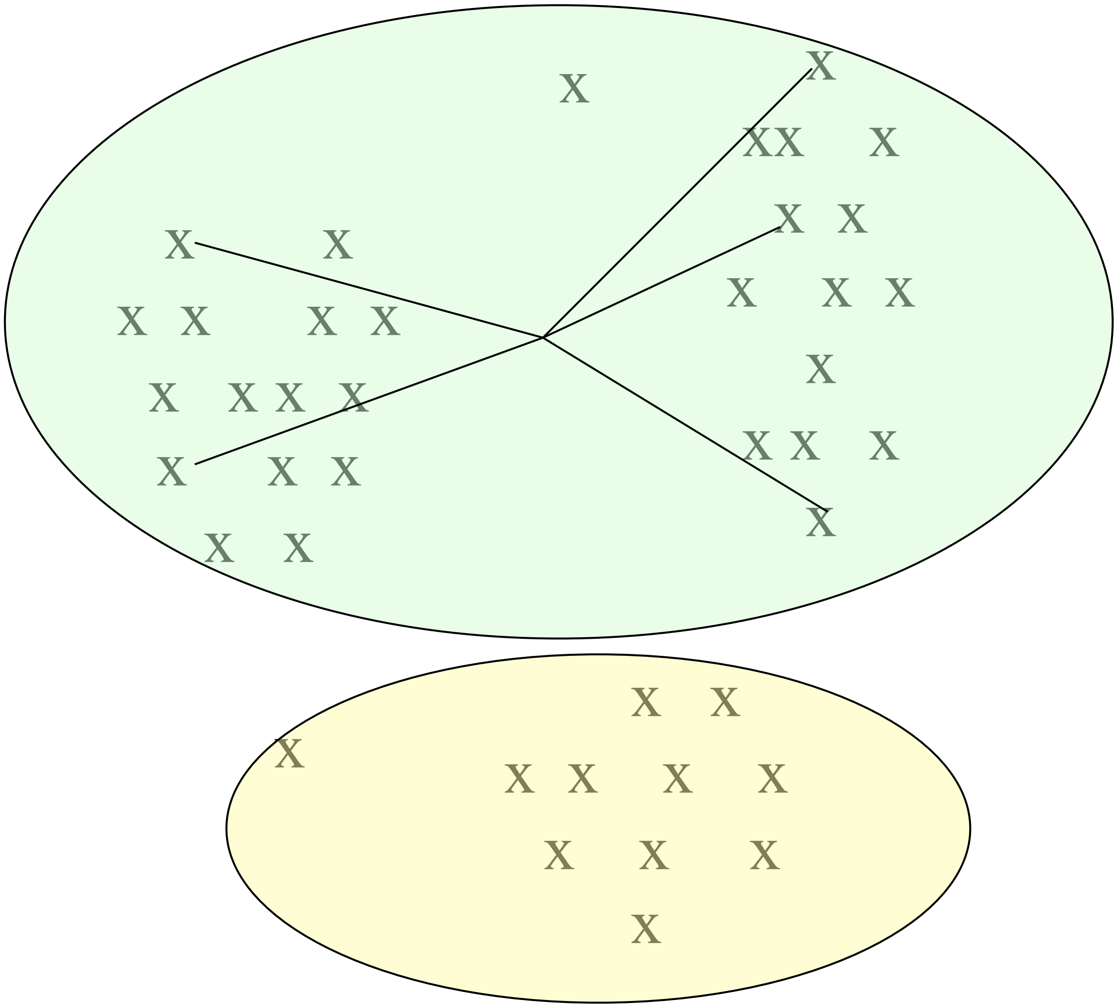

Example: Picking *k*

Just right;
distances
rather short.

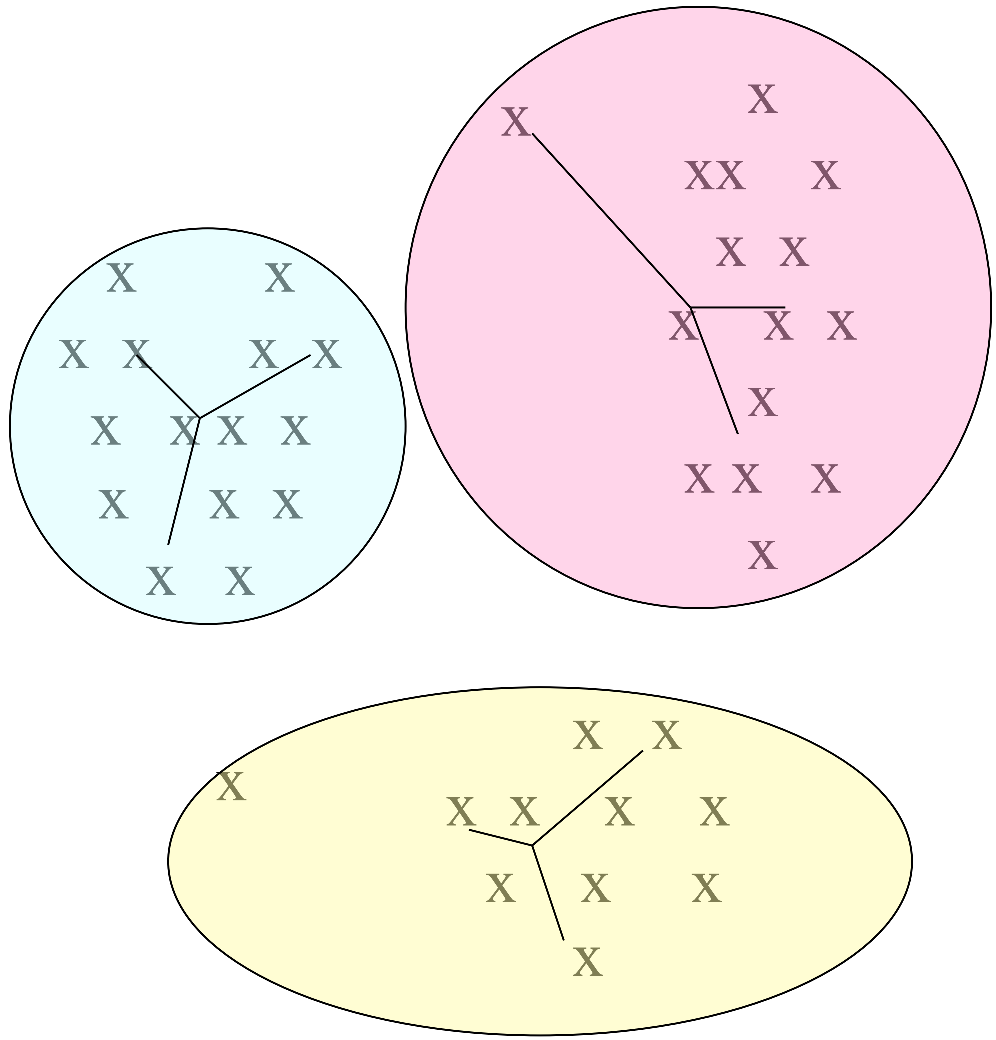

Example: Picking *k*

Too many;
little improvement
in average
distance.

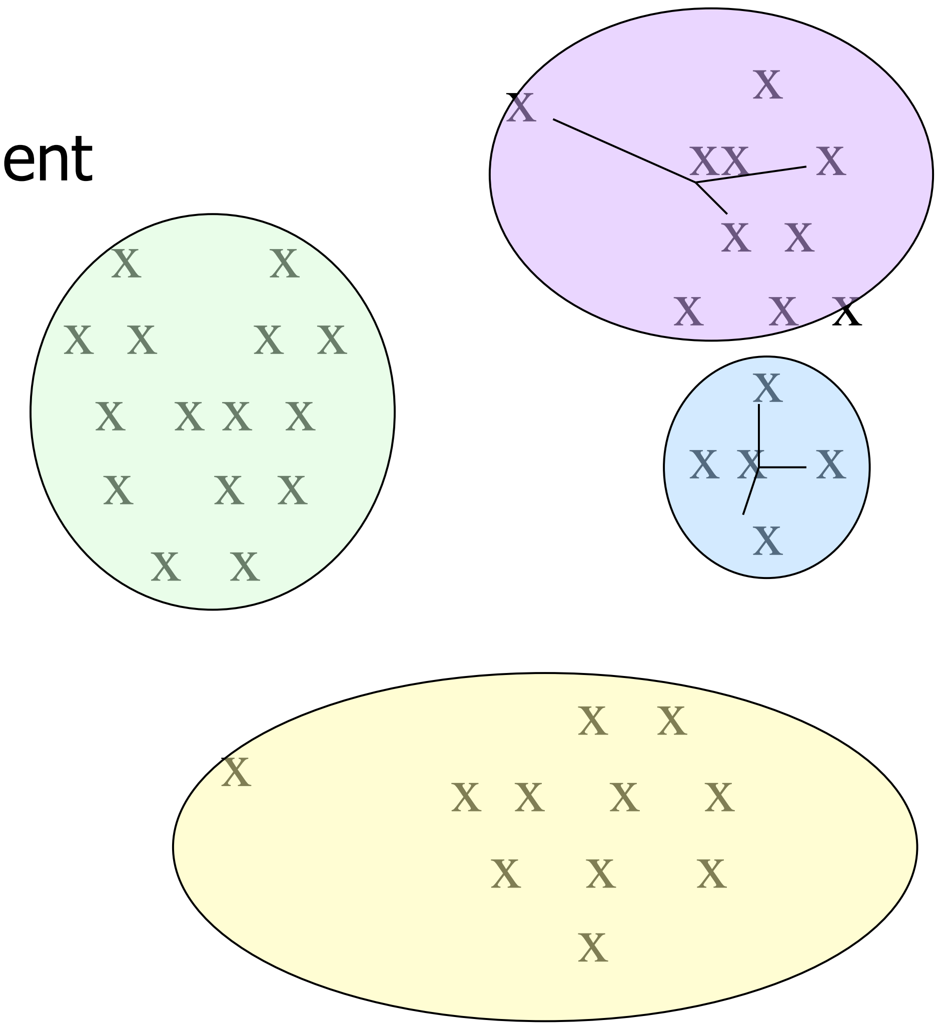

BFR Algorithm

- BFR (Bradley-Fayyad-Reina) is a variant of *k* -means designed to handle very large (disk-resident) data sets.
- It assumes that clusters are normally distributed around a centroid in a Euclidean space.
  - Standard deviations in different dimensions may vary.

BFR – (2)

- Points are read one main-memory-full at a time.
- Most points from previous memory loads are summarized by simple statistics.
- To begin, from the initial load we select the initial *k* centroids by some sensible approach.

Initialization: *k* -Means

- Possibilities include:
  1. Take a small random sample and cluster optimally.
  2. Take a sample; pick a random point, and then *k* – 1 more points, each as far from the previously selected points as possible.

Three Classes of Points

1. The *discard set*: points close enough to a centroid to be summarized.
2. The *compression set*: groups of points that are close together but not close to any centroid. They are summarized, but not assigned to a cluster.
3. The *retained set*: isolated points.

"Galaxies" Picture

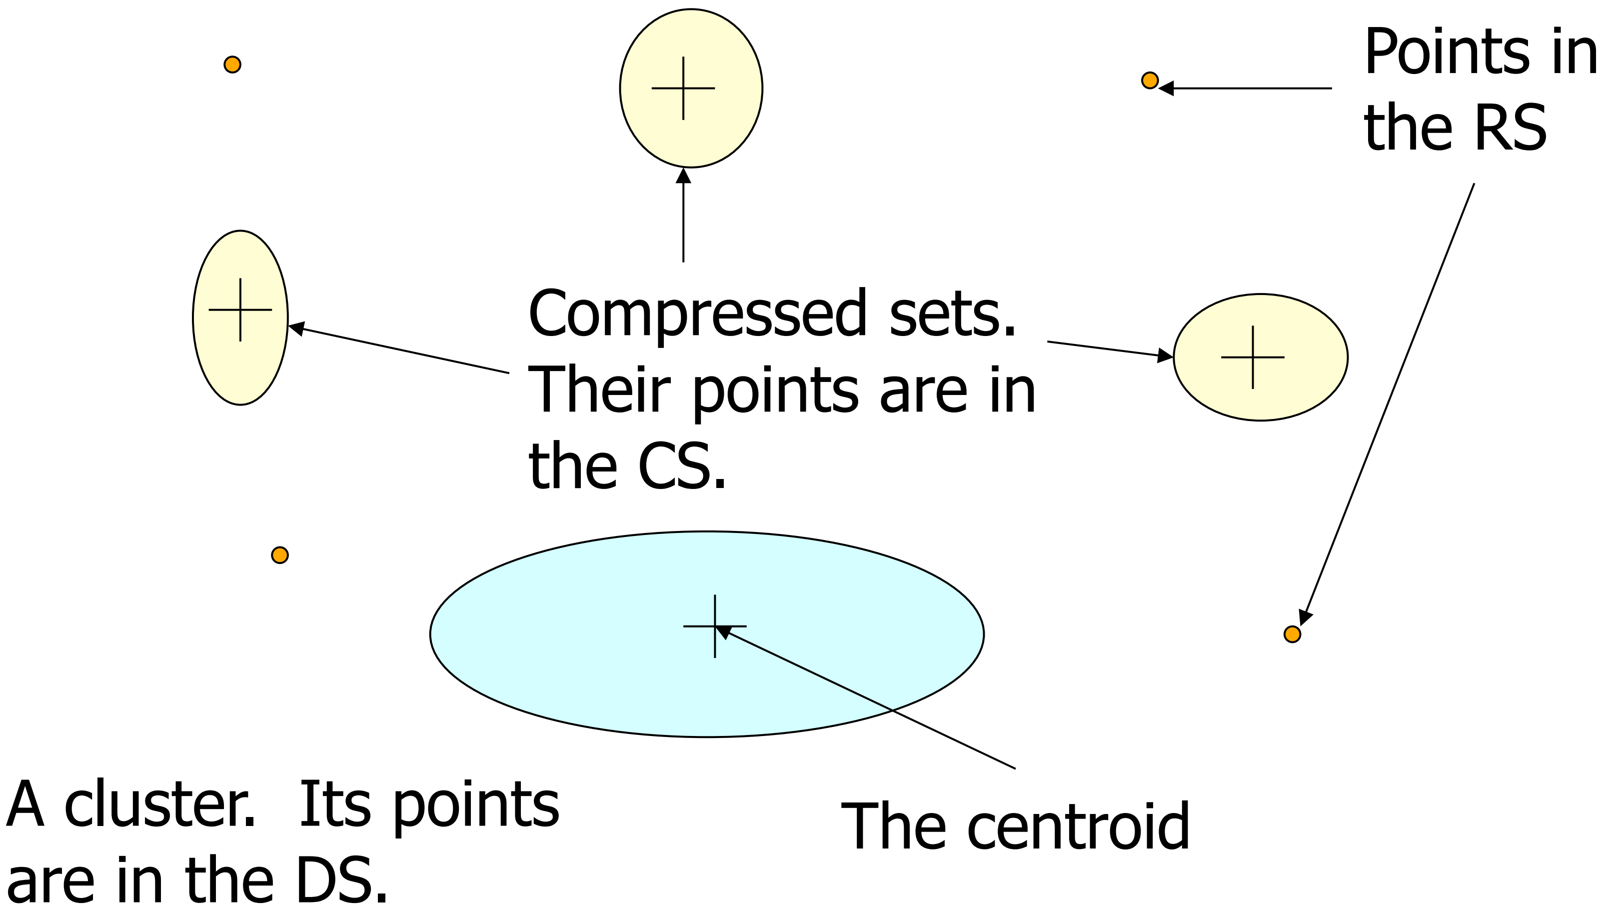

Summarizing Sets of Points

- For each cluster, the discard set is summarized by:
  1. The number of points, N.
  2. The vector SUM: i^{th} component = sum of the coordinates of the points in the i^{th} dimension.
  3. The vector SUMSQ: i^{th} component = sum of squares of coordinates in i^{th} dimension.

Comments

- $2d + 1$ values represent any number of points.
  - d = number of dimensions.
- Centroid (mean) in i^{th} dimension = SUM_i / N.
  - SUM_i = i^{th} component of SUM.
- Variance in dimension i can be computed by:

  $$(\text{SUMSQ}_i / N) - (\text{SUM}_i / N)^2$$
- Question: Why use this representation rather than directly store centroid and standard deviation?

Processing a "Memory-Load" of Points

1. Find those points that are "sufficiently close" to a cluster centroid; add those points to that cluster and the DS.
2. Use any main-memory clustering algorithm to cluster the remaining points and the old RS.
    - Clusters go to the CS; outlying points to the RS.

Processing – (2)

3. Adjust statistics of the clusters to account for the new points.
   - Add N's, SUM's, SUMSQ's.
4. Consider merging compressed sets in the CS.
5. If this is the last round, merge all compressed sets in the CS and all RS points into their nearest cluster.

A Few Details . . .

- How do we decide if a point is "close enough" to a cluster that we will add the point to that cluster?
- How do we decide whether two compressed sets deserve to be combined into one?

How Close is Close Enough?

- We need a way to decide whether to put a new point into a cluster.
- BFR suggest two ways:
  1. The *Mahalanobis distance* is less than a threshold.
  2. Low likelihood of the currently nearest centroid changing.

Mahalanobis Distance

- Normalized Euclidean distance from centroid.
- For point $(x_1,\ldots,x_k)$ and centroid $(c_1,\ldots,c_k)$:
  1. Normalize in each dimension: $y_i = (x_i - c_i)/\sigma_i$
  2. Take sum of the squares of the y_i's.
  3. Take the square root.

Mahalanobis Distance – (2)

- If clusters are normally distributed in d dimensions, then after transformation, one standard deviation = $\sqrt{d}$.
  - I.e., 70% of the points of the cluster will have a Mahalanobis distance < $\sqrt{d}$.
- Accept a point for a cluster if its M.D. is < some threshold, e.g. 4 standard deviations.

Picture: Equal M.D. Regions

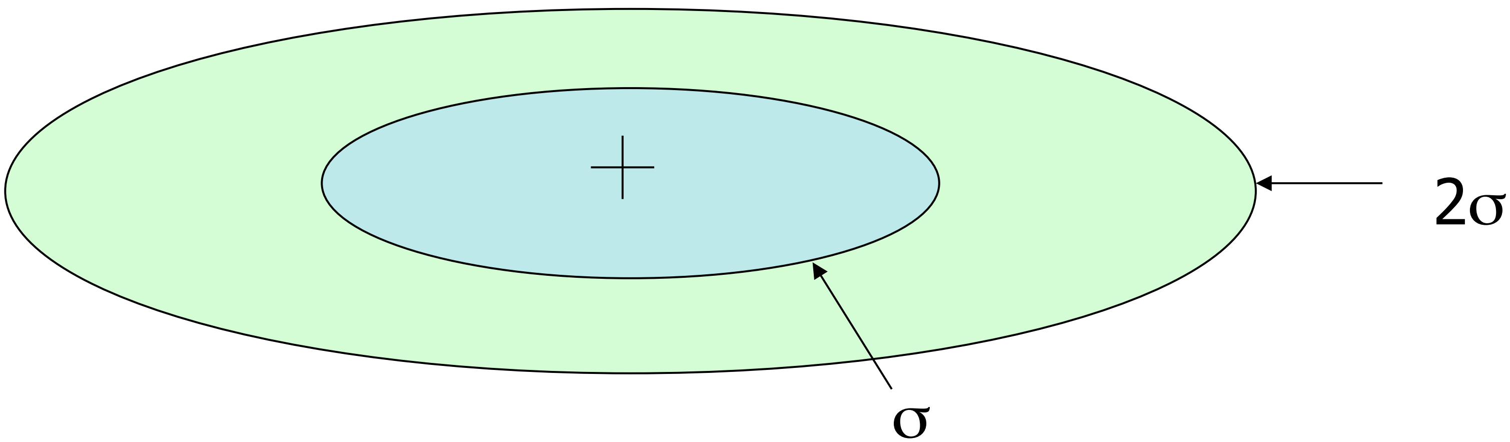

Should Two CS Subclusters Be Combined?

- Compute the variance of the combined subcluster.
  - *N*, SUM, and SUMSQ allow us to make that calculation quickly.
- Combine if the variance is below some threshold.
- Many alternatives: treat dimensions differently, consider density.

The CURE Algorithm

- Problem with BFR/*k* -means:
  - Assumes clusters are normally distributed in each dimension.
  - And axes are fixed – ellipses at an angle are *not* OK.
- CURE:
  - Assumes a Euclidean distance.
  - Allows clusters to assume any shape.

Example: Stanford Faculty Salaries

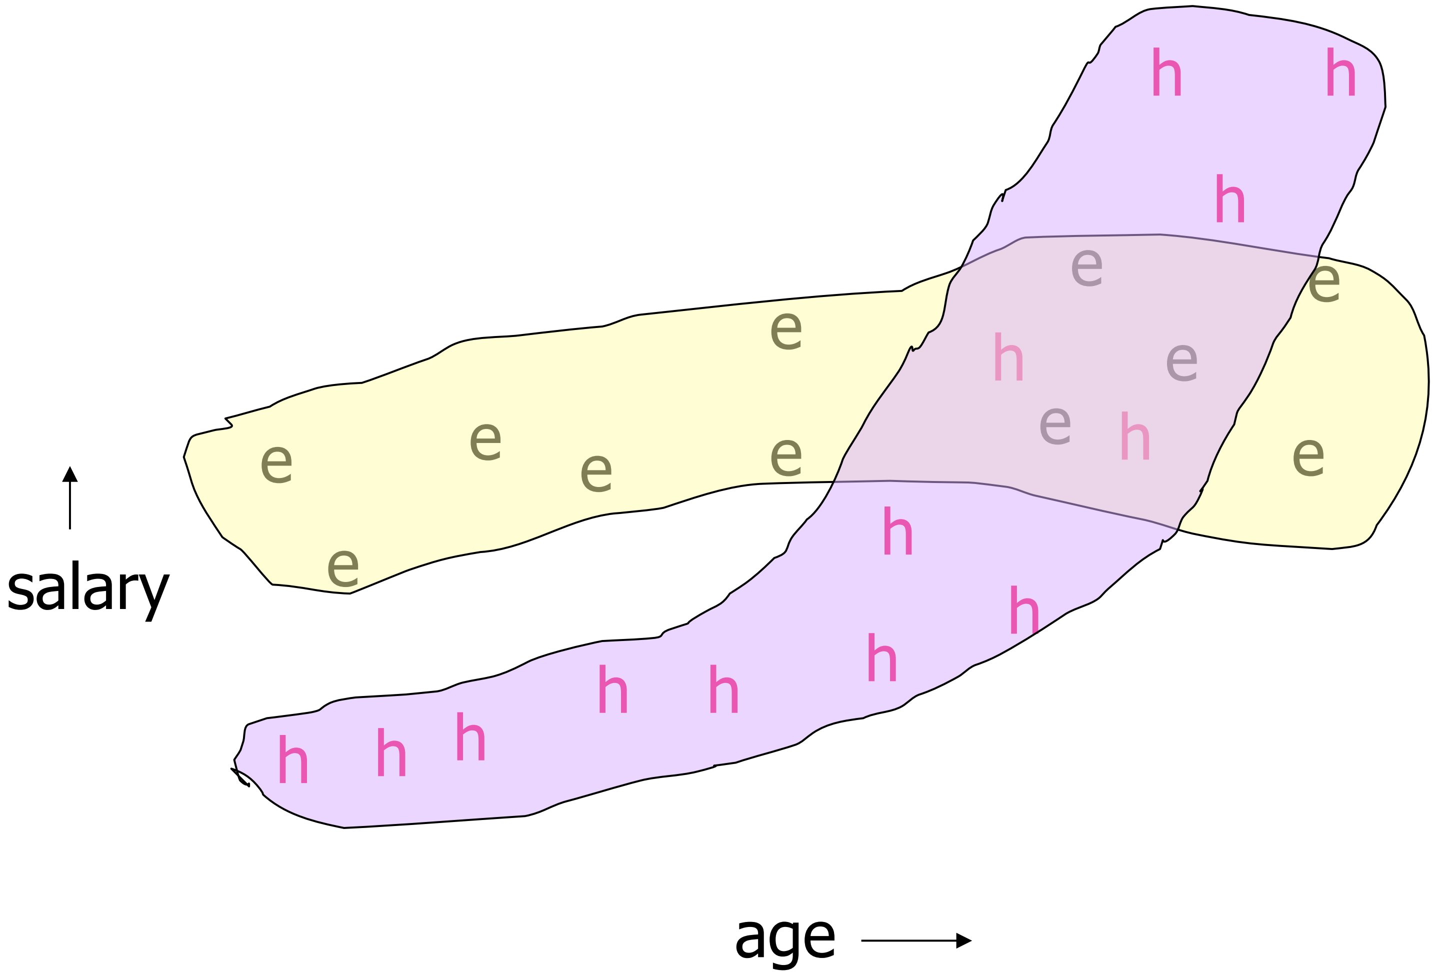

Starting CURE

1. Pick a random sample of points that fit in main memory.
2. Cluster these points hierarchically – group nearest points/clusters.
3. For each cluster, pick a sample of points, as dispersed as possible.
4. From the sample, pick representatives by moving them (say) 20% toward the centroid of the cluster.

Example: Initial Clusters

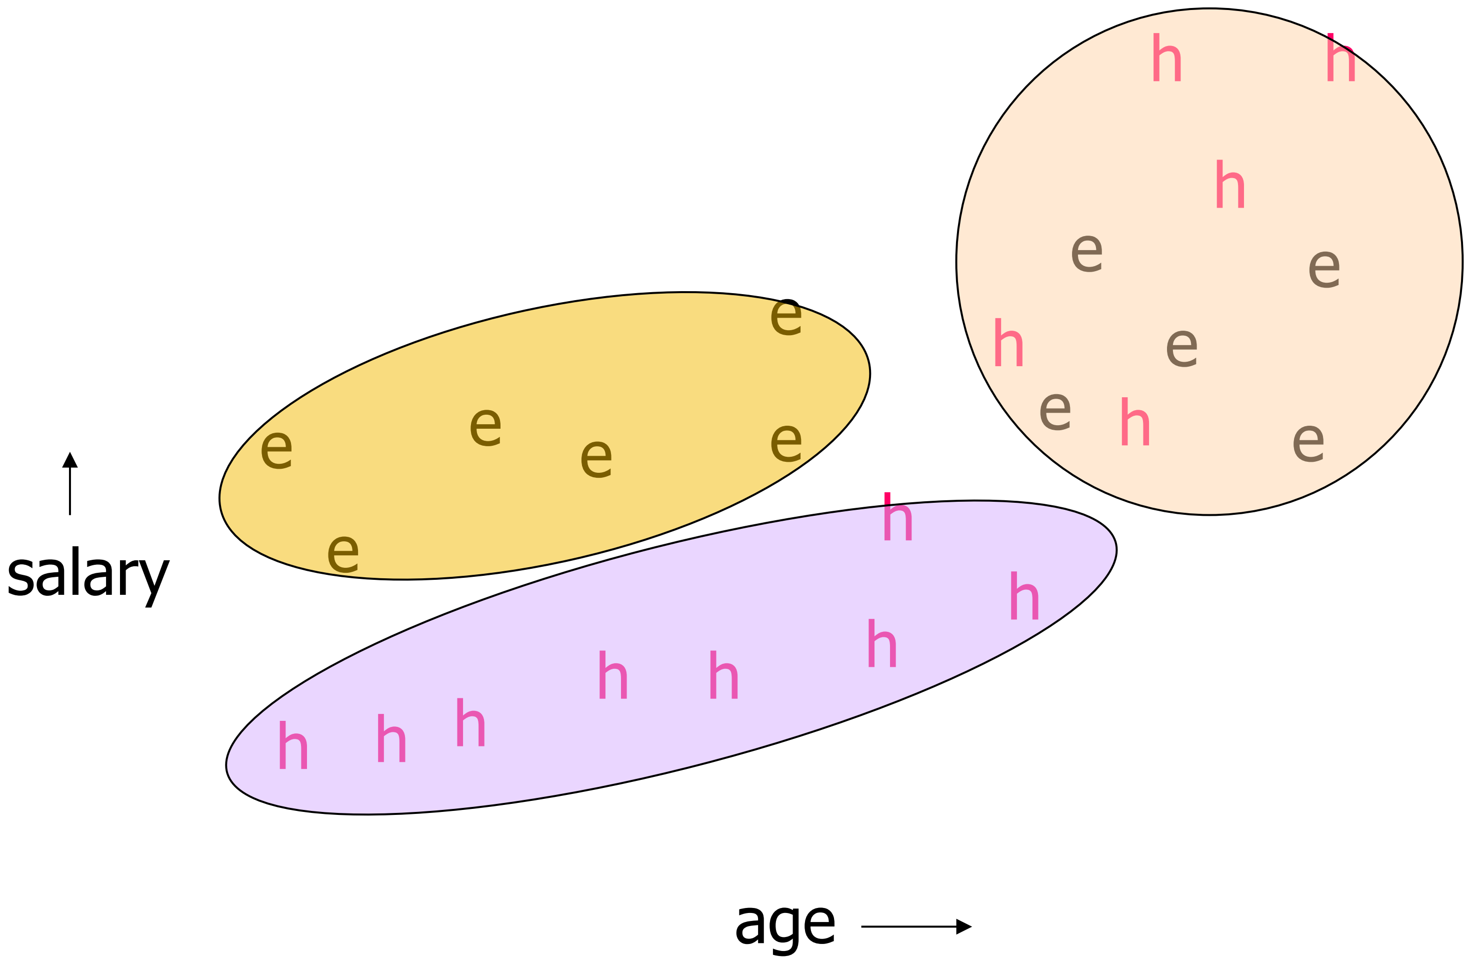

Example: Pick Dispersed Points

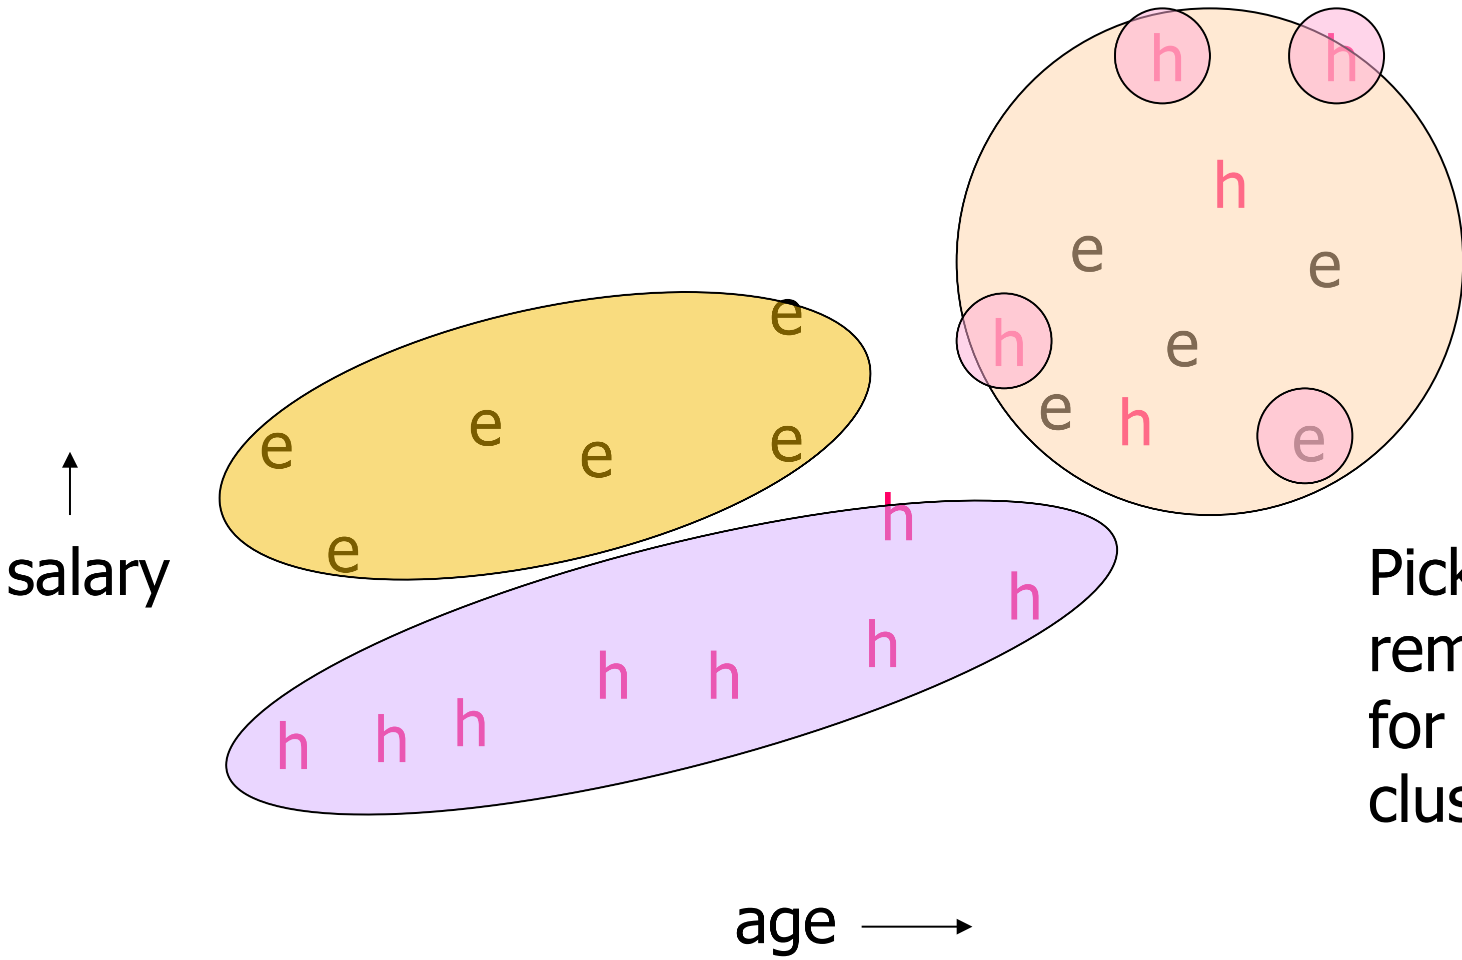

Pick (say) 4 remote points for each cluster.

Example: Pick Dispersed Points

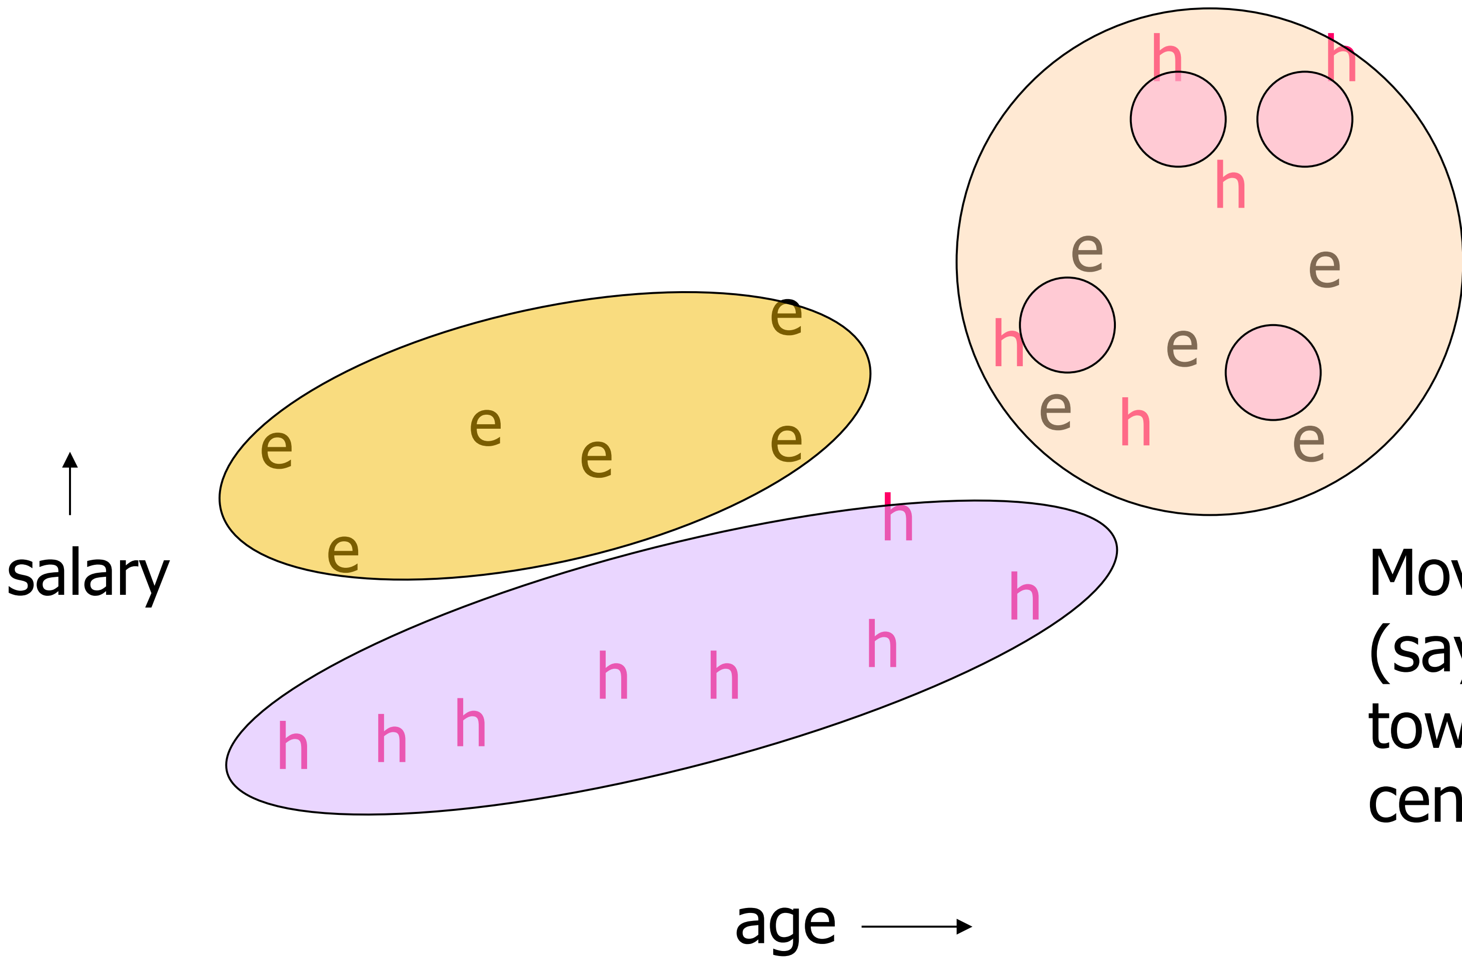

Move points (say) 20% toward the centroid.

Finishing CURE

- Now, visit each point *p* in the data set.
- Place it in the "closest cluster."
  - Normal definition of "closest": that cluster with the closest (to *p*) among all the sample points of all the clusters.