Unsupervised Learning: k-Means, PAM, Hierarchical Clustering

Ahmed Eleish

Data Analytics ITWS/CSCI/MGMT - 4600/6600 September 30th , 2025

Tetherless World Constellation
Rensselaer Polytechnic Institute

Contents

Unsupervised Learning

- Machine learning paradigm where algorithms learn patterns from unlabeled data.

- **Clustering** is a type of unsupervised learning with the goal of finding structure in a dataset by identifying natural clusters of data points based on a similarity criterion (usually distance).

- The observations in these clusters are generally more similar (closer) to each other than they are to points in other clusters.

k-Means

• k-Means clustering is an unsupervised learning algorithm that, as the name hints, finds a fixed number (*k*) of clusters in a set of data.

• A *cluster* is a group of data points that are grouped together due to similarities in their features. When using a K-Means algorithm, a cluster is defined by a *centroid*, which is a calculated point at the center of a cluster.

• Every point in a data set is part of the cluster whose centroid is most closely located in feature space. To put it simply, K-Means finds *k* number of centroids, and then assigns all data points to the closest cluster while minimizing the sum of squared distances from the points to their assigned centroid.

• K-Means assumes spherical clusters.

• K-Means is sensitive to outliers.

K-Means Algorithm

Algorithm 10.1 *K-Means Clustering*

1. Randomly assign a number, from 1 to K, to each of the observations. These serve as initial cluster assignments for the observations.

2. Iterate until the cluster assignments stop changing:

   (a) For each of the K clusters, compute the cluster *centroid*. The kth cluster centroid is the vector of the p feature means for the observations in the kth cluster.

   (b) Assign each observation to the cluster whose centroid is closest (where *closest* is defined using Euclidean distance).

Reference: Introduction to Statistical Learning with Applications in R, 7th Edition, Chapter 10 – KMeans

K-Means Algorithm

Algorithm 10.1 *K-Means Clustering*

1. Randomly assign a number, from 1 to K, to each of the observations. These serve as initial cluster assignments for the observations.

2. Iterate until the cluster assignments stop changing:

   (a) For each of the K clusters, compute the cluster *centroid*. The kth cluster centroid is the vector of the p feature means for the observations in the kth cluster.

   (b) Assign each observation to the cluster whose centroid is closest (where *closest* is defined using Euclidean distance).

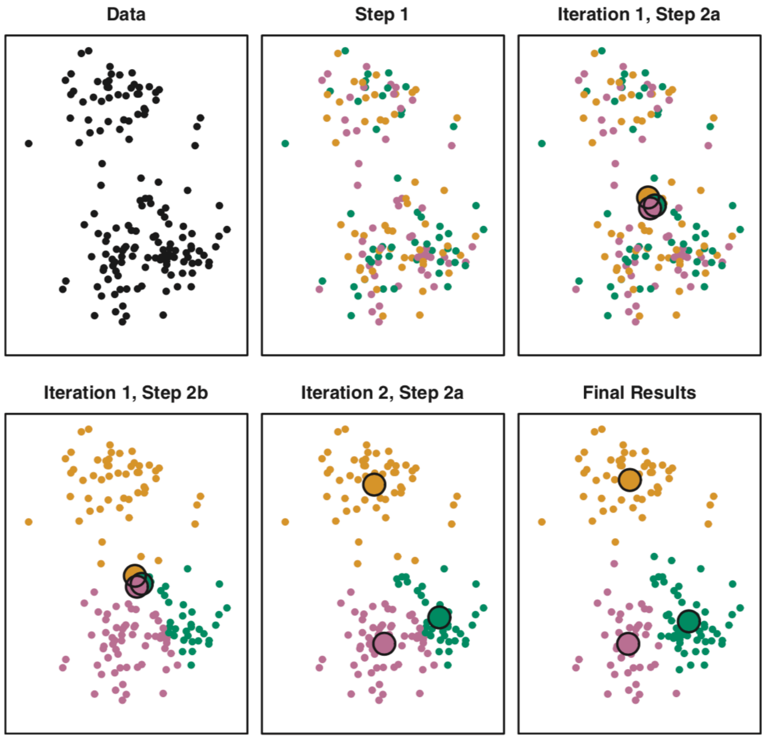

Reference: Introduction to Statistical Learning with Applications in R, 7th Edition, Chapter 10 – KMeans

K-Means Algorithm

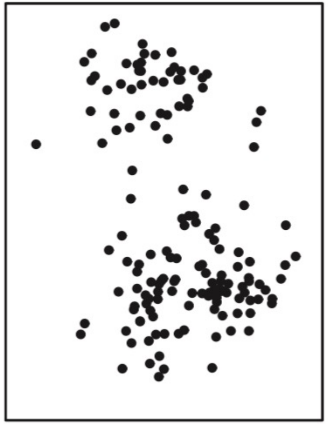

Observations (data) is shown

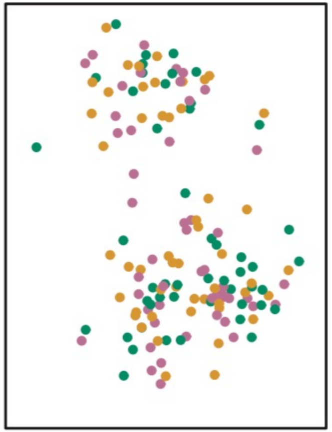

Step 1 of the algorithm: each observation is randomly assigned to a cluster

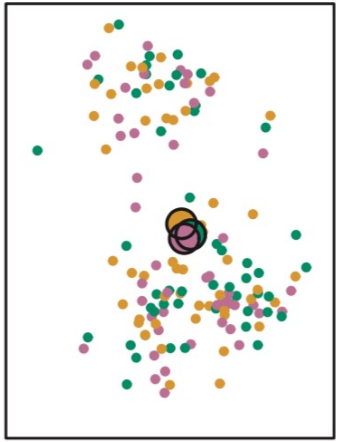

Iteration1 Step 2(a): The cluster centroids are computed; these are shown in large colored disks. Initially centroids are almost completely overlapping because the initial cluster assignment were chosen at random

Reference: Introduction to Statistical Learning with Applications in R, 7th Edition, Chapter 10 – KMeans

K-Means Algorithm

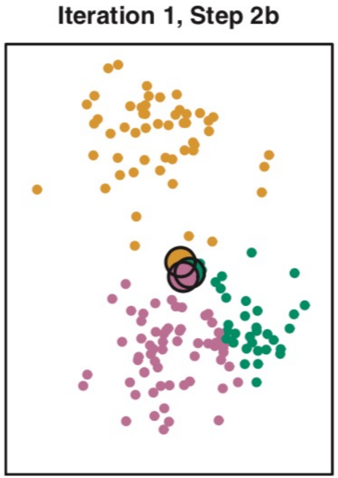

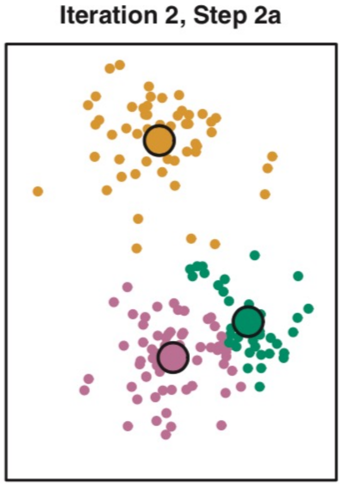

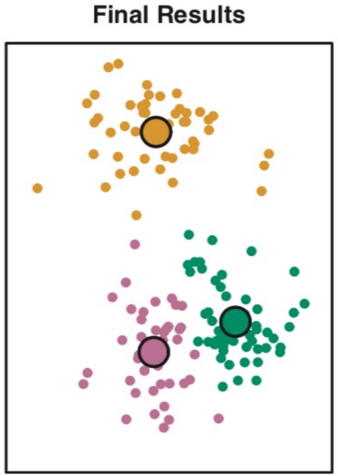

Iteration 1 Step 2(b) : each observation is assigned to the nearest centroid

Iteration 2, Step 2(a): the step 2(a) is once again performed, leading to new cluster centroids.

Final Results: the results obtained after ten iterations. You can see the distinct clusters with their centroids.

Image/Photo Credit: Introduction to Statistical Learning with Applications in R, 7th Edition, Chapter 10 – KMeans
Reference: Introduction to Statistical Learning with Applications in R, 7th Edition, Chapter 10 – KMeans

- K-Means clustering Animation

http://shabal.in/visuals/kmeans/6.html

Evaluating K-Means Models (Elbow Method)

Within-Cluster Sum of Squares: sum of squared Euclidean distances between all points in a cluster and the cluster's centroid.

$$WCSS = \sum_{i=1}^{k} \sum_{\mathbf{x} \in C_i} ||\mathbf{x} - \mathbf{c}_i||^2$$

- The elbow method is a heuristic that can be subjective and unreliable.

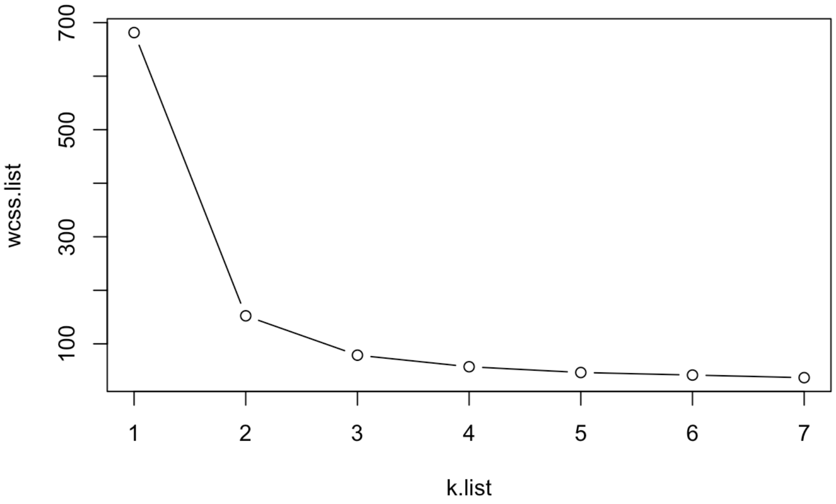

Plot of total within cluster sum of squares with values of *k*

Evaluating K-Means Models (Silhouette Method)

Silhouette value: a measure of similarity between a point and its cluster

$$a(i) = \frac{1}{|C_i| - 1} \sum_{j \in C_i, i \neq j} d(i, j)$$

Mean distance between *i* and other points in the cluster

$$b(i) = \min_{J \neq I} \frac{1}{|C_j|} \sum_{j \in C_j} d(i, j)$$

Minimum meqan distance between *i* and all points in any other clsuter

$$s(i) = \frac{b(i) - a(i)}{\max\{a(i), b(i)\}}, \text{ if } |C_i| > 1$$

$$s(i) = 0, \text{ if } |C_i| = 1$$

https://en.wikipedia.org/wiki/Silhouette_(clustering)

Evaluating K-Means Models (Silhouette Method)

Silhouette value: a measure of similarity between a point and its cluster

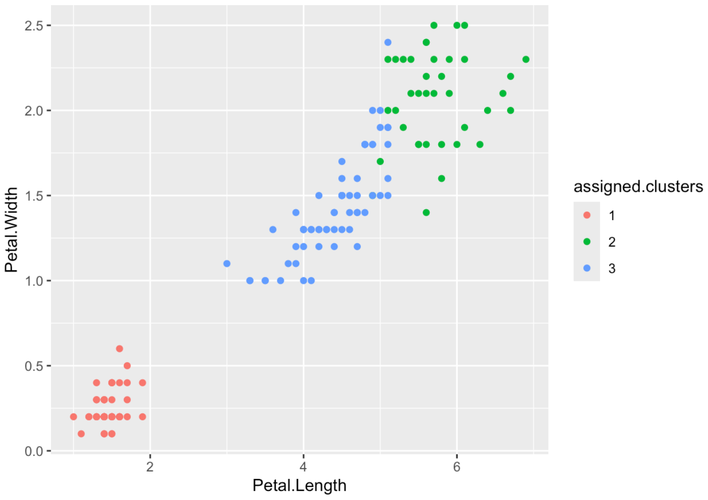

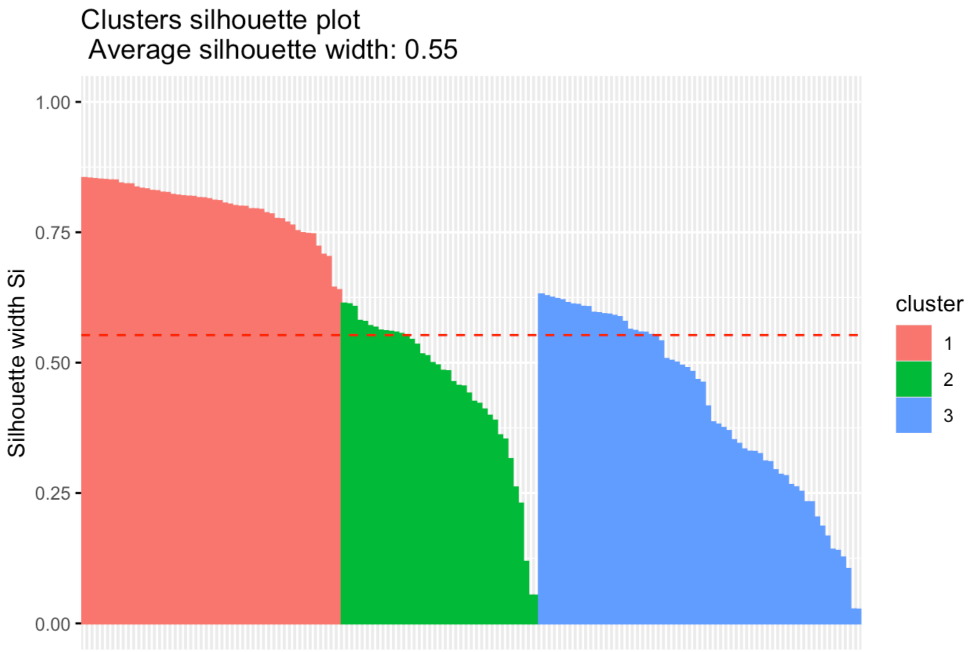

https://en.wikipedia.org/wiki/Silhouette_(clustering)

Partitioning Around Medoids (PAM)

- PAM, also called K-Medoids, is similar to K-Means but instead of calculating centroids, this algorithm selects actual data points (medoids) as cluster centers

- PAM is more robust to outliers and noise.

- The algorithm can be computationally expensive because it calculates pairwise distances between all data points.

Sadeghi, B. (2025). Clustering in geo-data science: Navigating uncertainty to select the most reliable method. *Ore Geology Reviews*, 106591. https://doi.org/10.1016/j.oregeorev.2025.106591

Partition Around Medoids (PAM)

Algorithm:

- **Build Phase** (find initial *k* medoids):

  - Select *k* points with the least cost, i.e. sum of distances to all other points
  - Assign all non-medoid points to the cluster whose medoid is closest

- **Swap Phase** (find best *k* medoids):

  - For each medoid, for each non-medoid:

    - Consider swapping the points, calculate the cost (summed distances)

  - Make the best swap

Partition Around Medoids (PAM)

Algorithm:

- **Build Phase** (find initial *k* medoids):

  - Select *k* points with the least cost, i.e. sum of distances to all other points
  - Assign all non-medoid points to the cluster whose medoid is closest

- **Swap Phase** (find best *k* medoids):

  - For each medoid, for each non-medoid:

    - Consider swapping the points, calculate the cost (summed distances)

  - Make the best swap

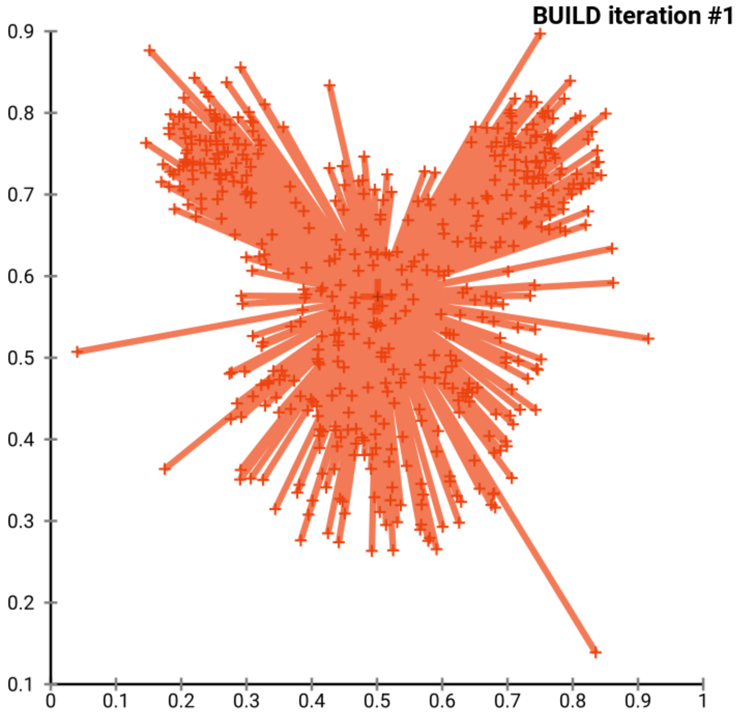

Evaluating PAM Models (Elbow Method)

Objective (Cost) Function: sum of distances between all points and their closest medoid.

$$\sum_{i=1}^{K} \sum_{x \in C_i} d(x, m_i)$$

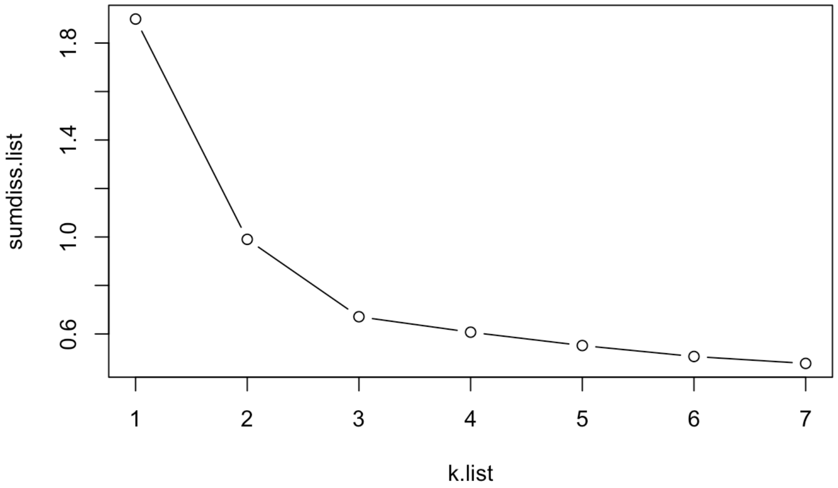

Plot of PAM cost function with values of *k*

Evaluating K-Means Models (Silhouette Method)

Silhouette value: a measure of similarity between a point and its cluster

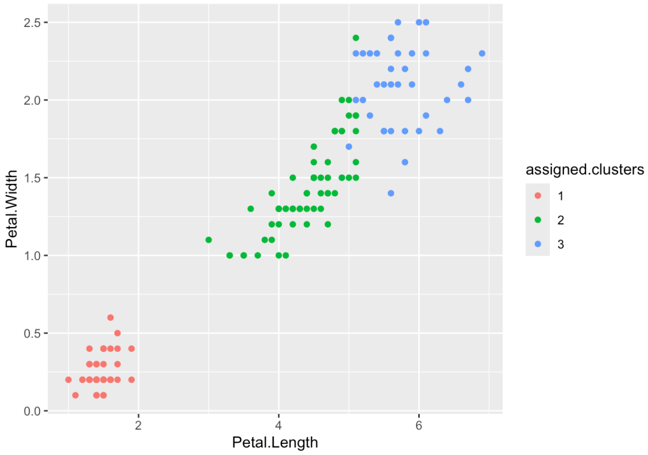

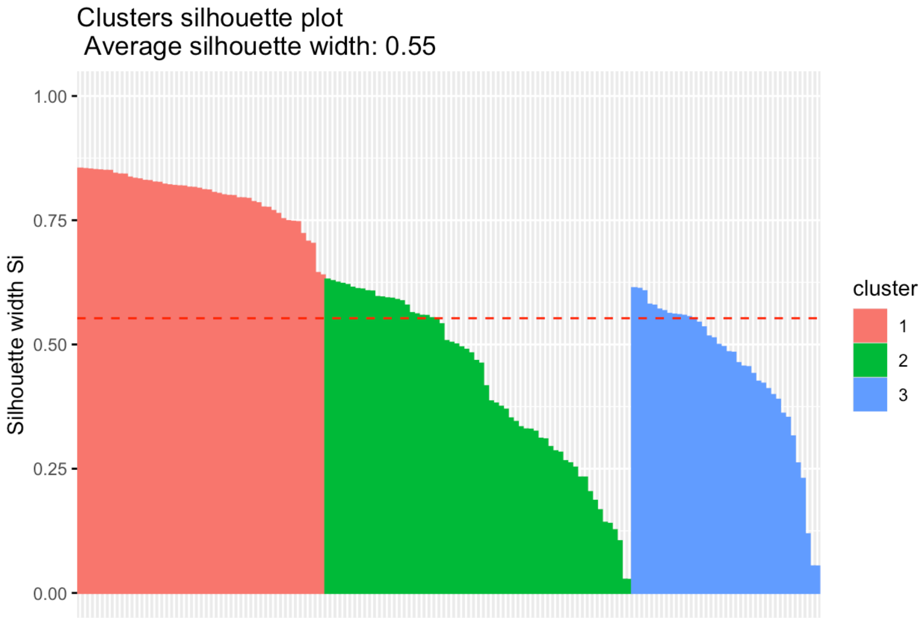

https://en.wikipedia.org/wiki/Silhouette_(clustering)

Hierarchical Clustering

- Contrary to partitional clustering, hierarchical clustering creates hierarchies of clusters through a bottom-up (agglomerative) or top-down (divisive) approach.

- Agglomerative clustering starts with each point in its own cluster, then iteratively merges clusters until it ends with a single cluster.

- Divisive clustering starts with all points in a single cluster, then recursively splits into multiple cluster.

- Hierarchical clustering offers flexibility and interpretability.

- Hierarchical clustering produces a dendrogram.

Sadeghi, B. (2025). Clustering in geo-data science: Navigating uncertainty to select the most reliable method. *Ore Geology Reviews*, 106591.
https://doi.org/10.1016/j.oregeorev.2025.106591

Hierarchical Clustering

Algorithm (agglomerative clustering)

1. Assign each point to a separate cluster (*n* cluster)
2. Merge the two closest* clusters
3. Repeat (2) Until all points are in a single cluster

* Distance between clusters:

- Complete linkage: maximum distance between points in two clusters
- Single linkage: minimum distance between points in two clusters
- Average linkage: average distance between points in two clusters

Evaluating Hierarchical Clustering (Dendrogram)

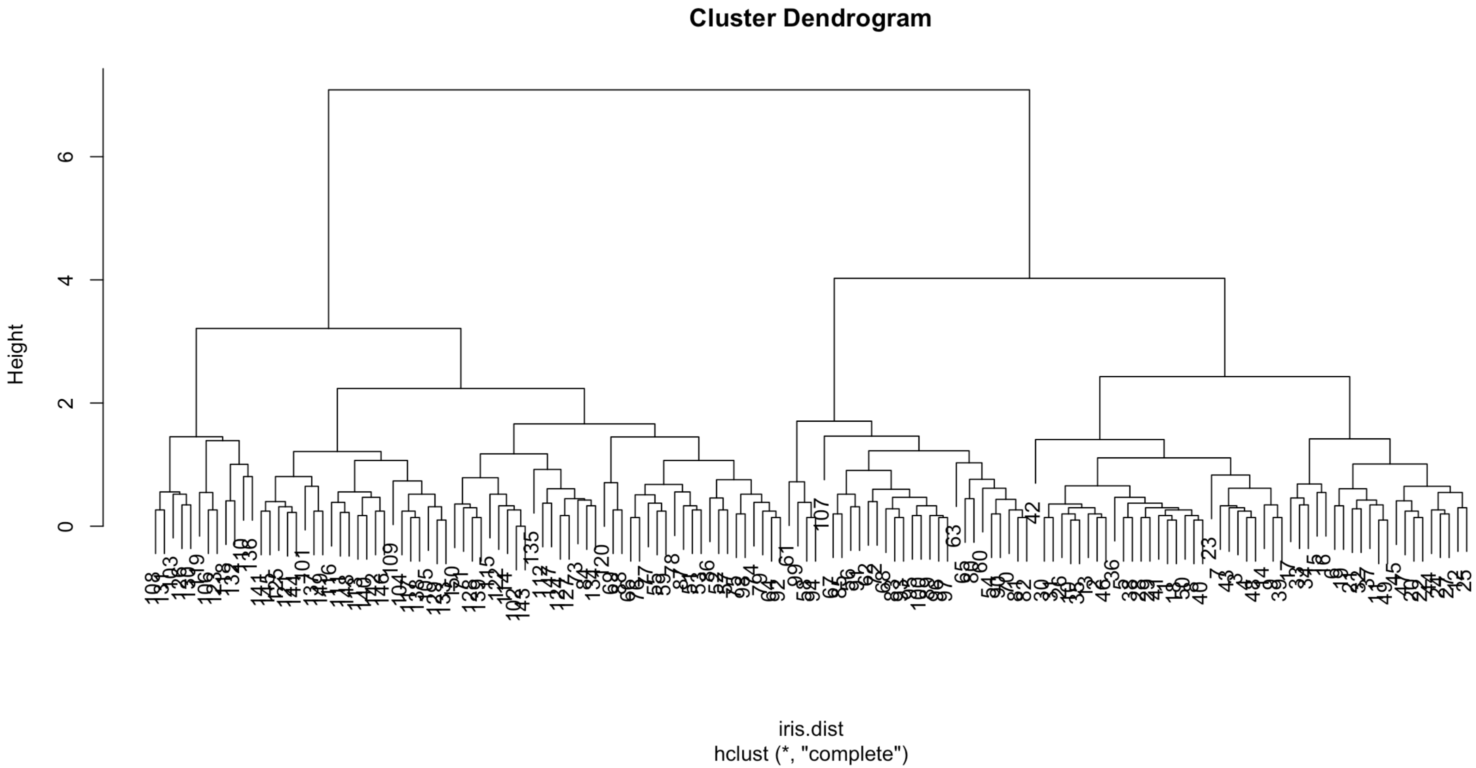

Evaluating Hierarchical Clustering (Dendrogram)

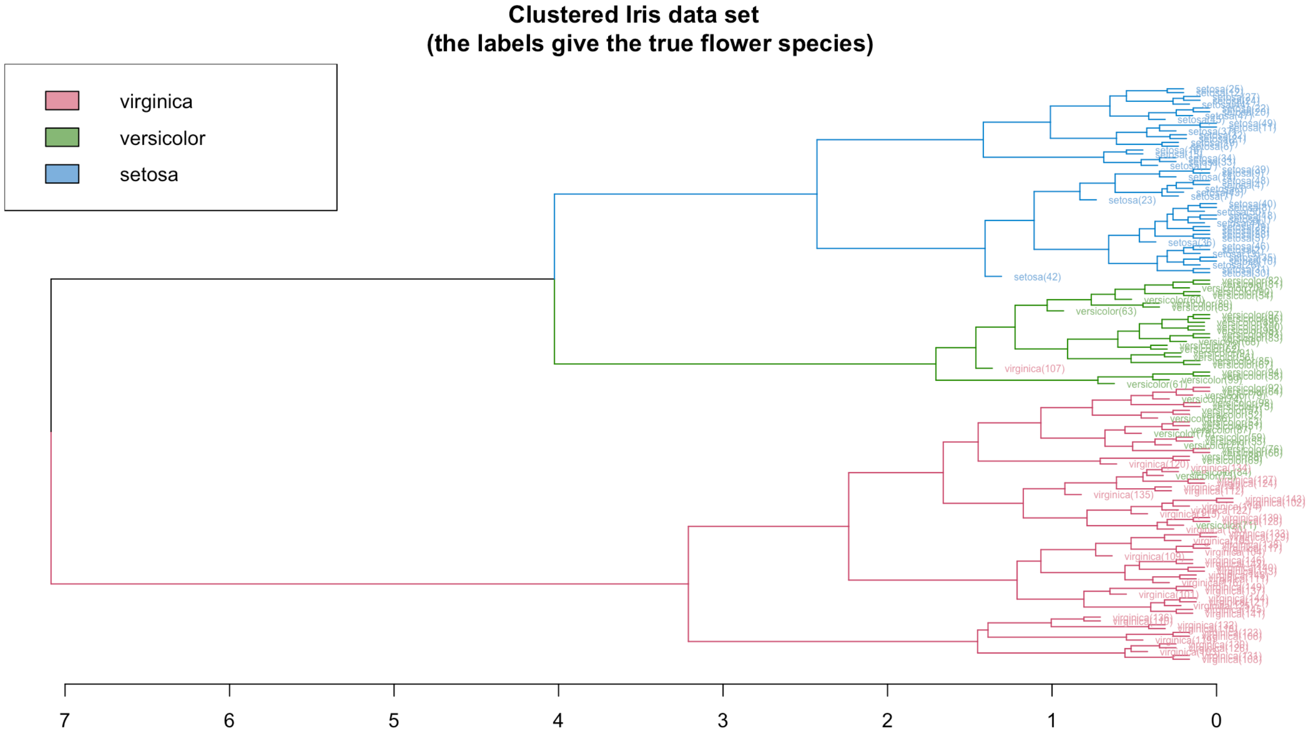

https://cran.r-project.org/web/packages/dendextend/vignettes/Cluster_Analysis.html

In class exercise code:
https://rpi.box.com/s/2wg4obl8ajrc1qm12rirdffylz96yn1d

Thanks!