CPSC 425: Computer Vision

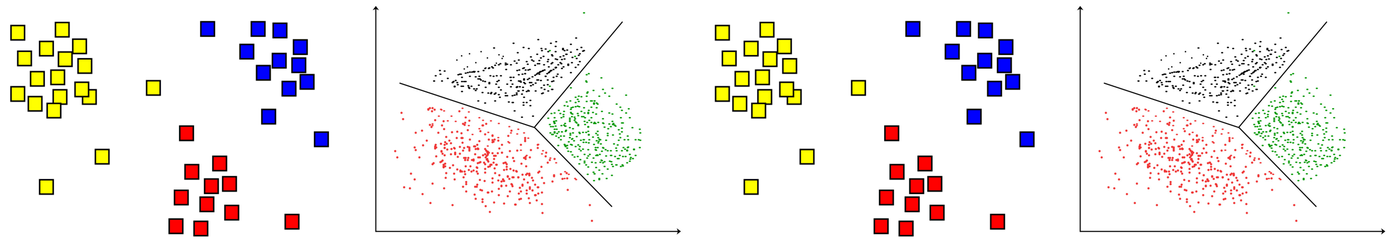

Lecture 23: Clustering

Grouping in Human Vision

Humans routinely group features that belong together when looking at a scene.
What are some cues that we use for grouping?

Grouping in Human Vision

Humans routinely group features that belong together when looking at a scene.
What are some cues that we use for grouping?

— Similarity
— Symmetry
— Common Fate
— Proximity
— ...

Grouping in Human Vision

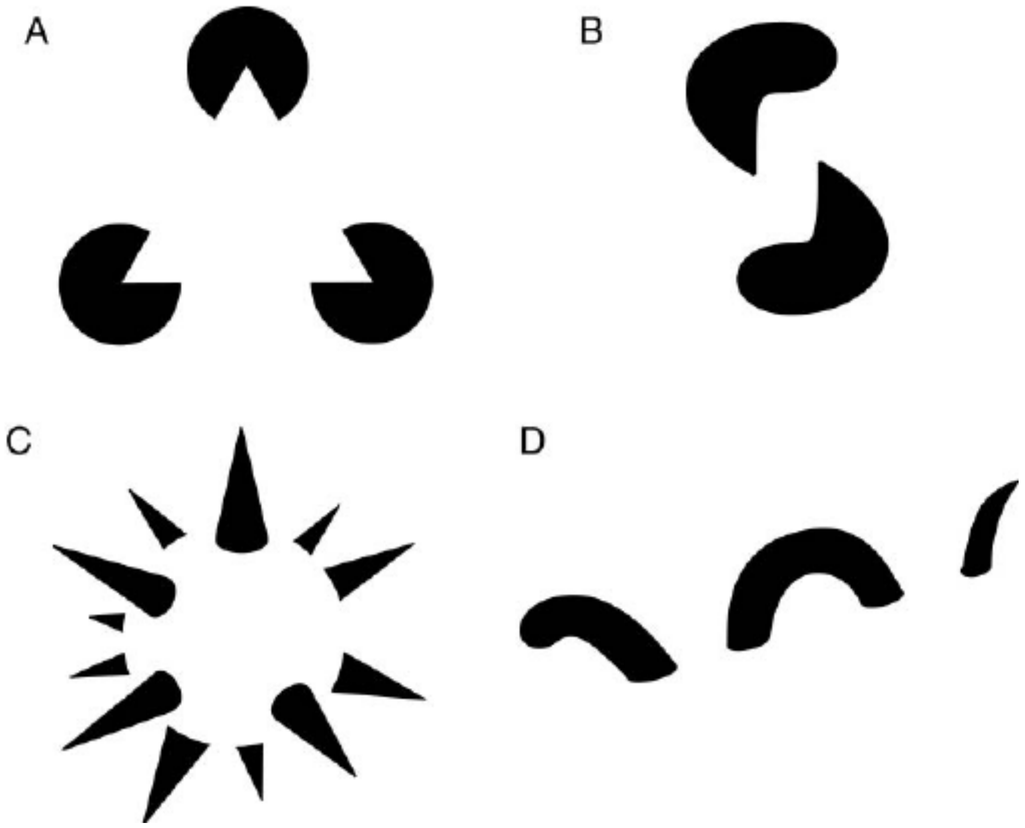

A. Kanizsa triangle
B. Tse's volumetric worm
C. Idesawa's spiky sphere
D. Tse's "sea monster"

Figure credit: Steve Lehar

Grouping in Human Vision

Slide credit: Kristen Grauman

Grouping in Human Vision

Slide credit: Kristen Grauman

Clustering

It is often useful to be able to **group** together **image regions** with similar appearance (e.g. roughly coherent colour or texture)

— image compression

— approximate nearest neighbour search

— base unit for higher-level recognition tasks

— moving object detection in video sequences

— video summarization

Clustering

Clustering is a set of techniques to try to find components that belong together (i.e., components that form clusters).

— Unsupervised learning (access to data, but no labels)

Two basic clustering approaches are

— **agglomerative clustering**

— **divisive clustering**

Agglomerative Clustering

Each data point starts as a separate cluster. Clusters are recursively merged.

Algorithm:

Make each point a separate cluster

Until the clustering is satisfactory

    Merge the two clusters with the smallest inter-cluster distance

end

Agglomerative Clustering

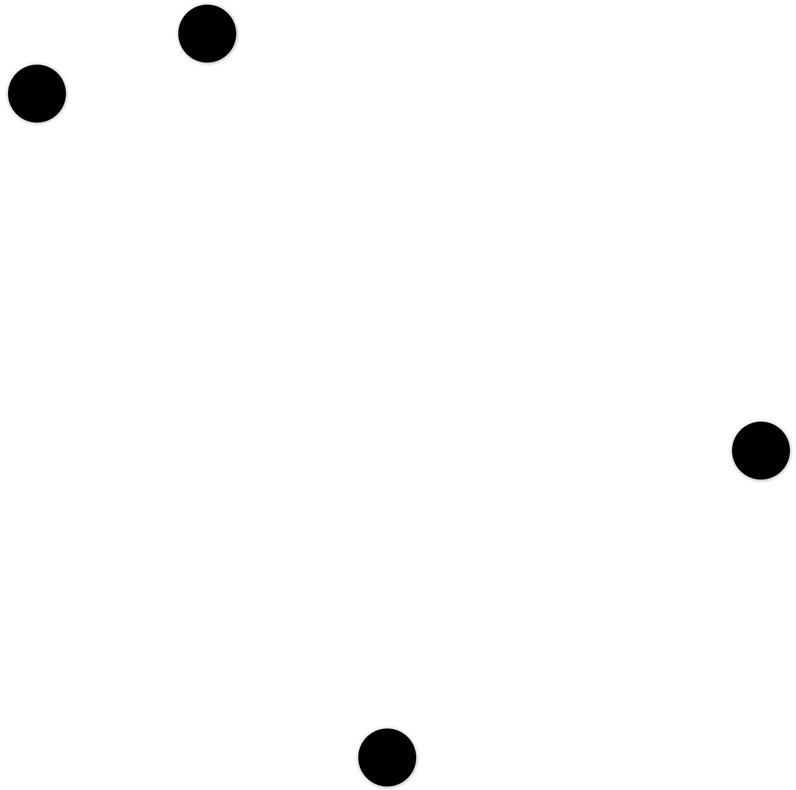

Agglomerative Clustering

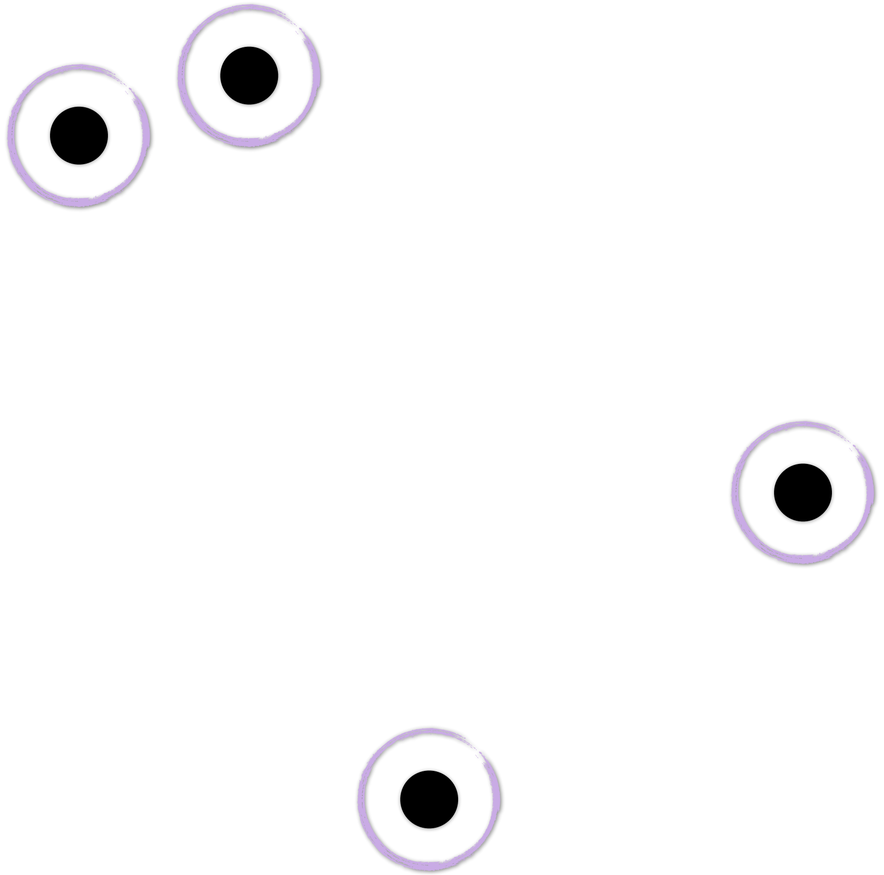

Agglomerative Clustering

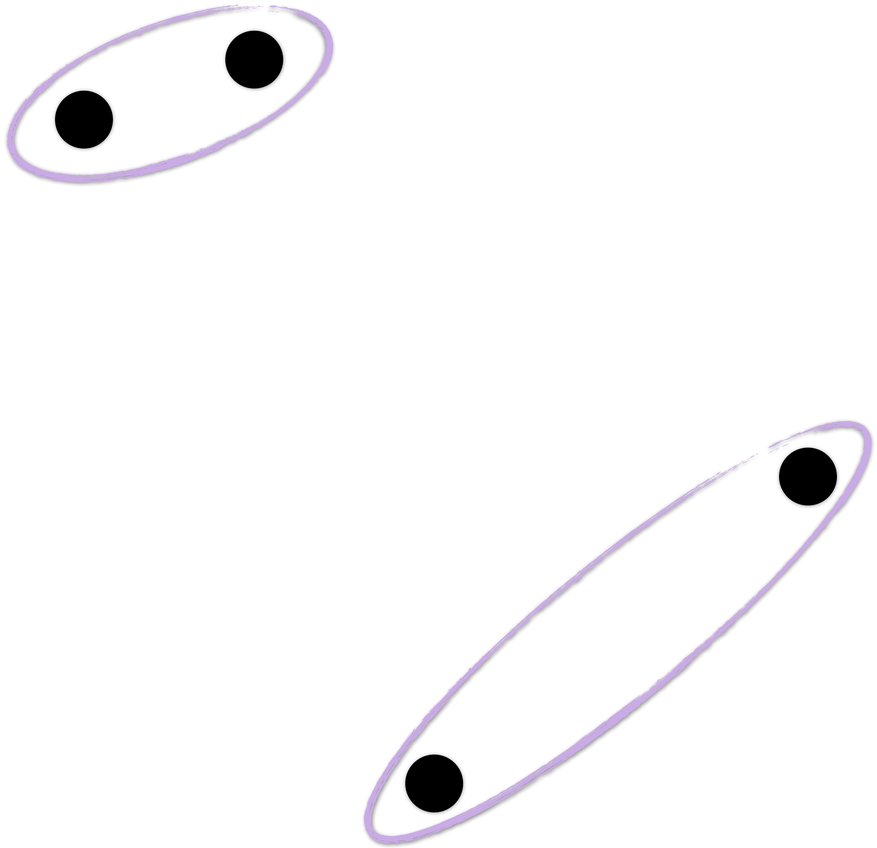

Divisive Clustering

The entire data set starts as a single cluster. Clusters are recursively split.

Algorithm:

Construct a single cluster containing all points

Until the clustering is satisfactory

    Split the cluster that yields the two components

    with the largest inter-cluster distance

end

Divisive Clustering

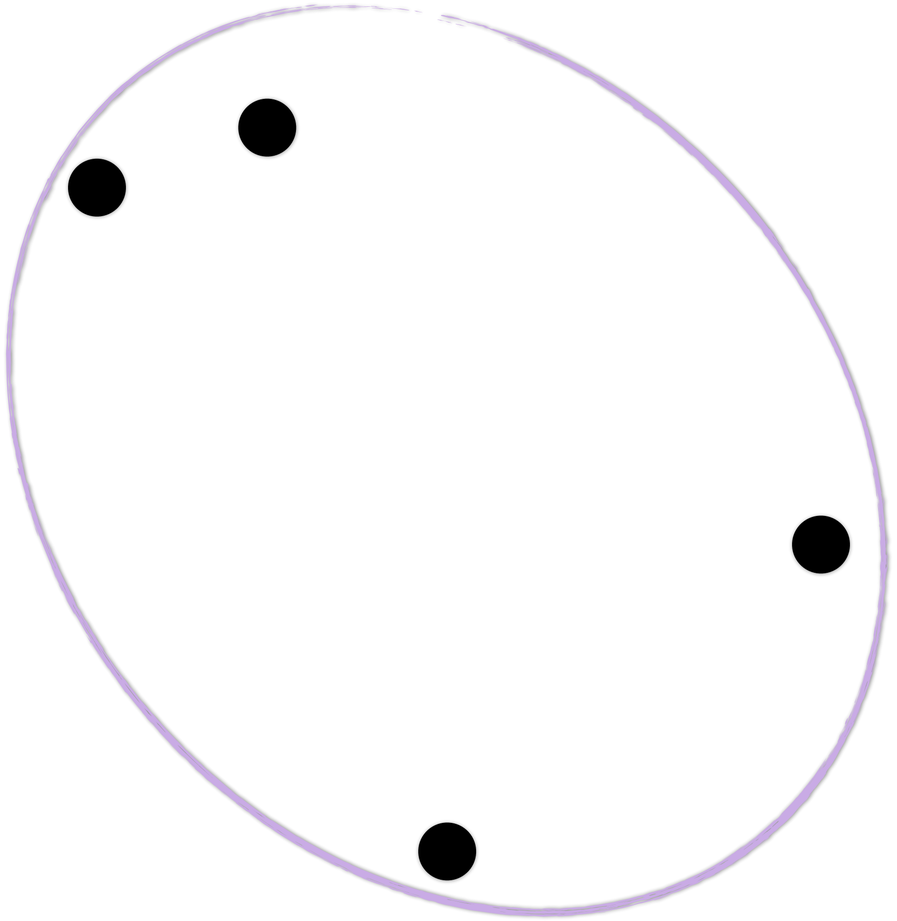

Divisive Clustering

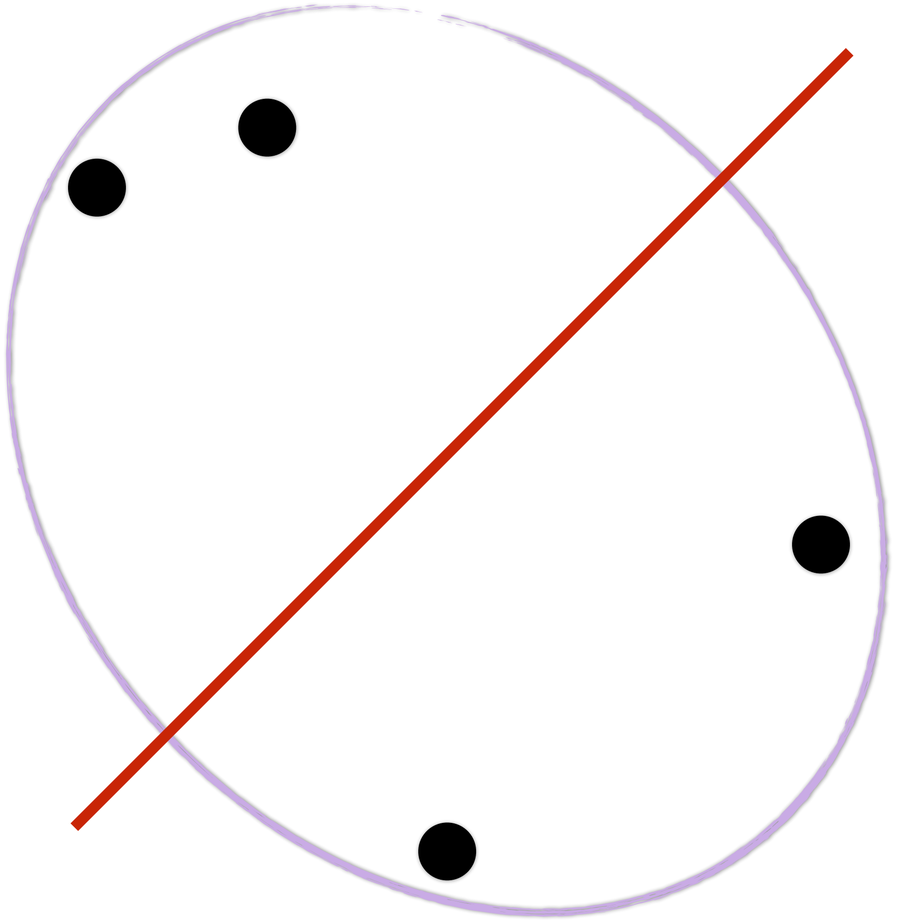

Divisive Clustering

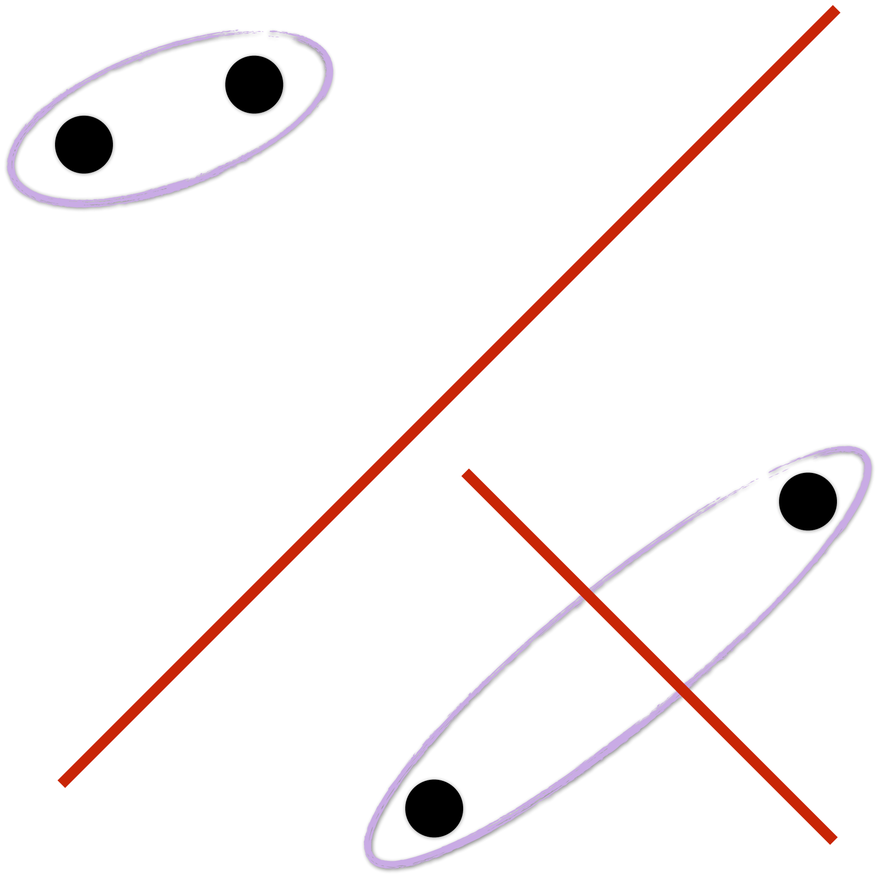

Divisive Clustering

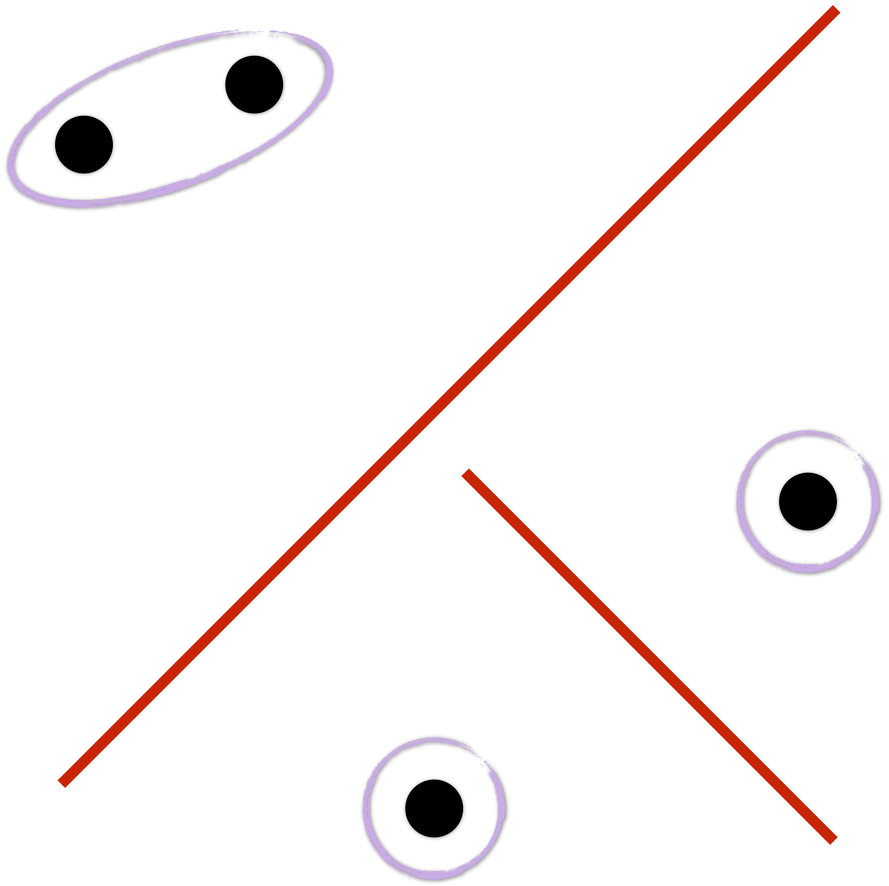

Inter-Cluster Distance

How can we define the cluster distance between two clusters C_1 and C_2 in agglomerative and divisive clustering? Some common options:

the distance between the closest members of C_1 and C_2

$$\min d(a, b), a \in C_1, b \in C_2$$

– single-link clustering

the distance between the farthest members of C_1 and a member of C_2

$$\max d(a, b), a \in C_1, b \in C_2$$

– complete-link clustering

Inter-Cluster Distance

How can we define the cluster distance between two clusters C_1 and C_2 in agglomerative and divisive clustering? Some common options:

an average of distances between members of C_1 and C_2

$$\frac{1}{|C_1||C_2|} \sum_{a \in C_1} \sum_{b \in C_2} d(a, b)$$

– group average clustering

Dendrogram

The algorithms described generate a hierarchy of clusters

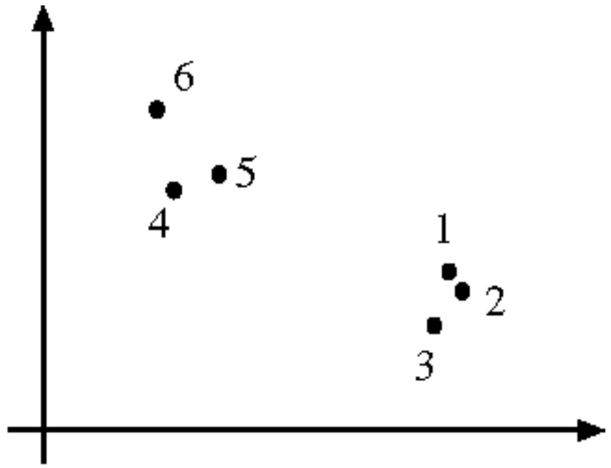

Forsyth & Ponce (2nd ed.) Figure 9.15

Dendrogram

The algorithms described generate a hierarchy of clusters, which can be visualized with a **dendrogram**.

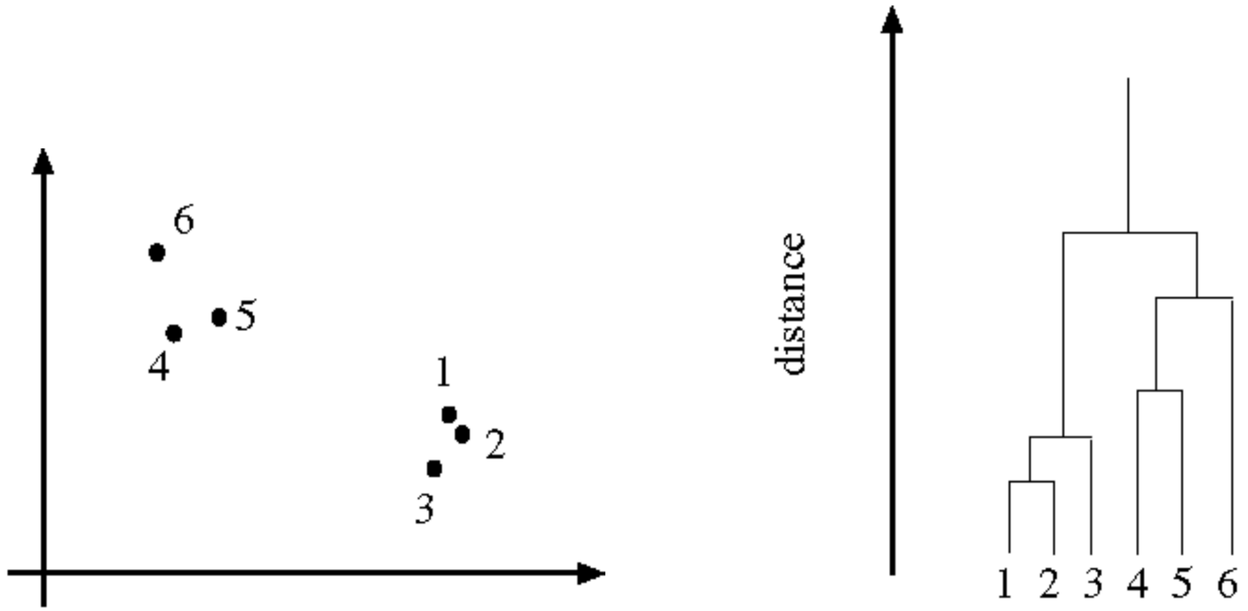

Forsyth & Ponce (2nd ed.) Figure 9.15