Ensembles of Classifiers and Clusterings

Lecture 1

Reasons for using Ensembles

- **Statistical reasons**:
  - Combining the output of several classifiers may reduce the risk of an unfortunate selection of a poorly performing classifier

Reasons for using Ensembles

- **Large Volumes of Data**:
  - Sometimes, the amount of data to be analyzed can be too large to be handled by a single classifier. Thus, we can:
    - Partition the data into smaller subsets;
    - Train different classifiers;
    - Combine their outputs using a combination rule

Reasons for using Ensembles

- **Too Little Data**:
  - A reasonable sized set of training data is crucial to learn the underlying data distribution. When available data is scarce, we can:
    - Draw overlapping random subsets of the available data using resampling techniques
    - Train different classifiers, creating the ensemble

Reasons for using Ensembles

- **Divide and Conquer**:
  - The given task may be too complex, or lie outside the space of functions that can be implemented by the chosen classifier method (e.g.: non-linear problem, and linear classifiers)
  - Appropriate combinations of simple (e.g., linear) classifiers can learn complex (e.g., non-linear) boundaries

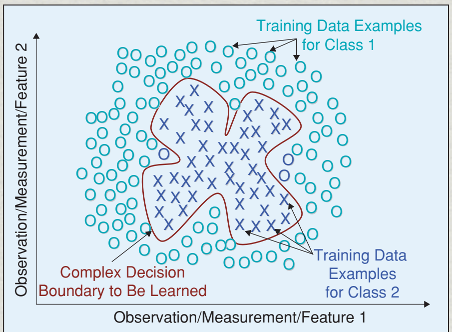

Figure 1. Complex decision boundary that cannot be learned by linear or circular classifiers.

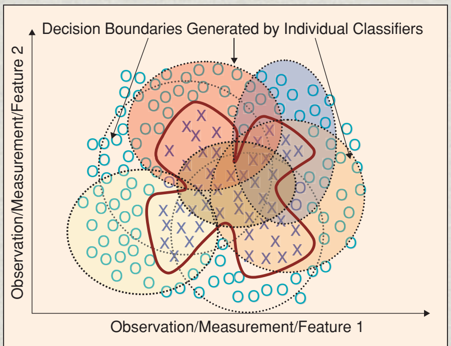

Figure 2. Ensemble of classifiers spanning the decision space.

Reasons for using Ensembles

- **Data Fusion**:
  - Several sets of data obtained from different sources, where the nature of features is different (e.g.: categorical and numerical features)
  - Data from each source can be used to train a different classifier, thus creating an ensemble

Components of an Ensemble

- Two key components:
  - A method to generate the individual classifiers of the ensemble
  - A method for combining the outputs of these classifiers

Diversity: The Key Feature

- The individual classifiers must be diverse, i.e., they make errors on different data
- Intuition: if they make the same errors, such mistakes will be carried into the final prediction
- Thus: the errors the classifiers make should be uncorrelated

Accuracy

- The component classifiers need to be “reasonably accurate” to avoid poor classifiers to obtain the majority of votes.
- Intuition: If the components of the ensemble are poor classifiers, they make a lot of errors, and those errors are carried out to the final prediction.

Accuracy and Diversity

- Requirements for accuracy and diversity have been quantified:
  - Under simple majority voting and *independent error conditions*, if all classifiers have the same probability of error of *less than 50%*, then the error of the ensemble decreases monotonically with an increasing number of classifiers.

How to achieve diversity

- Use different training data sets to train individual classifiers
- Such data sets are often obtained through resampling techniques (***bootstrapping*** or ***bagging***): training data subsets are drawn randomly, usually with replacement, from the entire training data

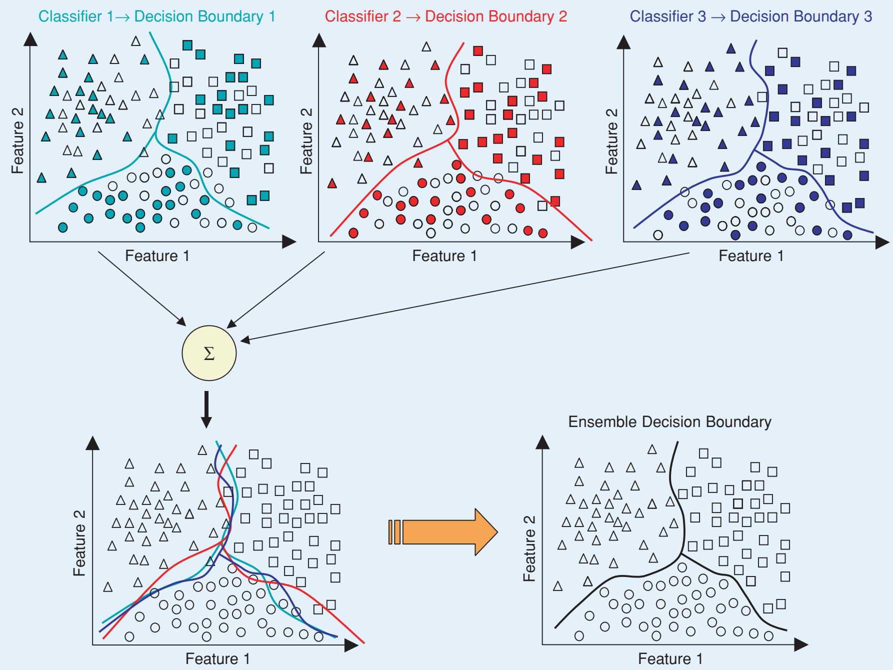

Figure 3. Combining classifiers that are trained on different subsets of the training data.

How to achieve diversity

- Use different training data sets to train individual classifiers
- If the training data subsets are drawn without replacement, the procedure is also called ***jackknife*** or ***k-fold*** data split: the entire data set is split into k blocks, and each classifier is trained only on k-1 of them. A different subset of k blocks is selected for each classifier

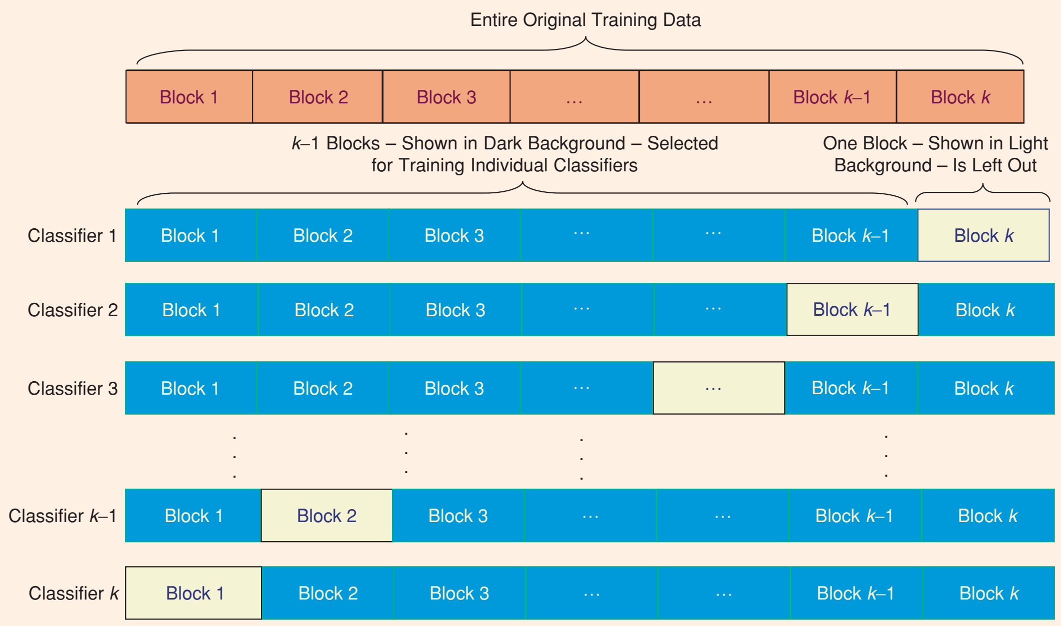

Figure 4. *k*-fold data splitting for generating different, but overlapping, training datasets.

How to achieve diversity

- When is bagging (bootstrapping) effective?
- To ensure diverse classifiers, the base classifier should be ***unstable***, that is, *small changes* in the training set should lead to *large changes* in the classifier output.

How to achieve diversity

- When is bagging (bootstrapping) effective?
- Large error reductions have been observed with *decision trees* and bagging. This is because decision trees are highly sensitive to small perturbations of the training data.

How to achieve diversity

- When is bagging (bootstrapping) effective?
- Bagging is not effective with *nearest neighbor classifiers*. Why? NN classifiers are highly stable with respect to variations of the training data
- It has been shown that the probability that any given training point is included in a data set bootstrapped by bagging is approximately 63.2%. It follows that the nearest neighbor will be the same in 63.2% of the classifiers
- Thus, the errors are highly correlated, and bagging becomes ineffective

How to achieve diversity

- Use different training parameters for different classifiers
- E.g., ensemble of neural networks trained with different weight initialization, or different number of layers/nodes
- If the base classifier is unstable with respect to the tuning parameters, diverse classifiers can be generated

How to achieve diversity

- Use different type of classifiers
- E.g., an ensemble of neural networks, decision trees, nearest neighbor classifiers, and support vector machines

How to achieve diversity

- Use different subsets of features to train the individual classifiers
- E.g., random feature subsets (random subspace method)
- This approach is effective with nearest neighbor (NN) methods, because NN techniques are highly sensitive to the chosen features

Clustering Ensembles

- Clustering ensembles leverage the diversity of the input clusterings to generate a consensus clustering that is superior to the component ones;
- Clustering ensembles offer a solution to challenges inherent to clustering arising from its ill-posed nature;
- The major challenge is to find a consensus clustering that achieves an improved clustering of the data

The Clustering Ensemble process

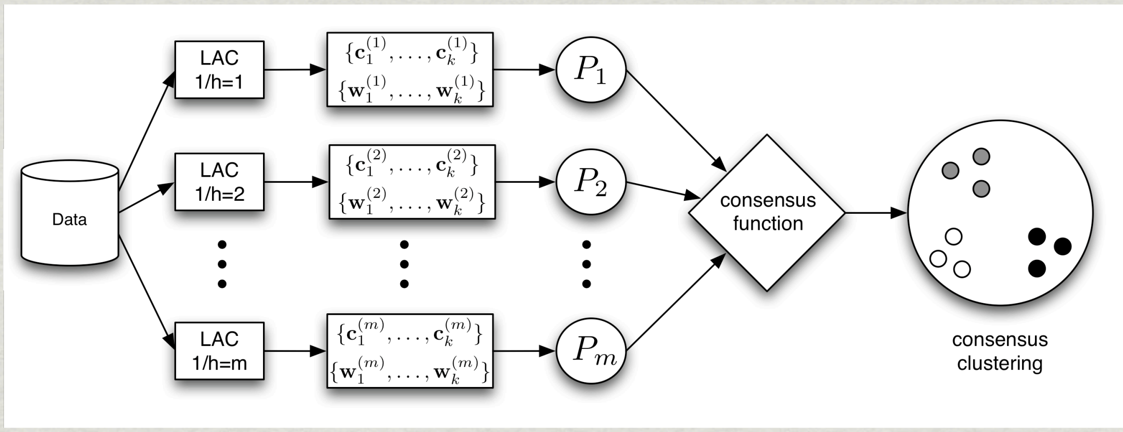

- Goal: Aggregate a collection of base clusterings to produce a partition of the data that is more accurate that the component ones

Clustering Ensembles

- A clustering ensemble technique is characterized by two components:
  - The mechanism to generate diverse clusterings
  - The consensus function to combine the input clusterings into a final clustering

Clustering Ensembles

- Diverse component clusterings can be generated by:
  - Varying the number and/or location of initial centroids
  - Using different clustering algorithms
  - Sub-sampling features or data

Clustering Ensembles

- A popular methodology to build a consensus function is to use the co-association matrix:
  - Two points have similarity 1 if they belong to the same cluster; similarity 0 otherwise
  - This defines a binary similarity matrix for each clustering
  - Lets consider an example...

Clustering Ensembles

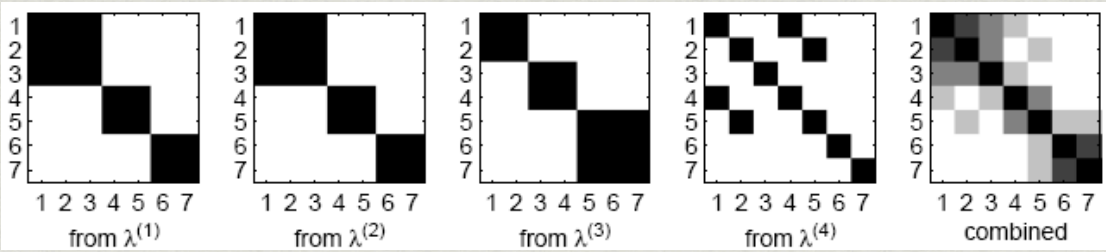

- Overall similarity matrix S: entry-wise average of the m individual matrices (m=4 above)
- An element of S represents the fraction of clusterings in which two data are in the same cluster
- S is used to re-cluster the data using a similarity-based clustering algorithm, e.g., hierarchical clustering

Clustering Ensembles

- A different popular mechanism for constructing a consensus maps the problem onto a graph-based partitioning setting:
  - From S, a similarity graph is induced: vertices correspond to data, and edge weights represent the similarity between the corresponding two vertices
  - A k-way partitioning of the vertices that minimizes the edge weight-cut is computed
  - The result gives the consensus clustering.