

Data Mining

Lab 3: Data classification

Outline:

- The classification problem
- Instance based classifiers (k- NearestNeighbor)
- Decision trees (J48)
- Rule-based classifiers (OneR, JRip)
- Multilayer Perceptron

1. Classification

Quick reminder of main concepts:

- Aim of a classifier: establish to which class the input data belongs
- Constructing a classifier: extract a classification model using the data from a training set; the construction process is called supervised training/learning
- Performance evaluation: confusion matrix, accuracy, sensitivity, specificity, recall, precision, F-measure

Accuracy = $(TP+TN)/(TP+TN+FP+FN)$ = nr of correctly classified data/total nr of data

Sensitivity = $TP/(TP+FN)$ (TP rate or recall)

Specificity = $TN/(TN+FP)$ (TN rate), 1-specificity = $FP/(TN+FP)$ = FP rate

Precision = $TP/(TP+FP)$ (nr of true positive cases/ nr of cases classified as positives)

F-measure = $2 * precision * recall / (precision + recall)$

ROC (Receiver Operating Characteristic) Curve: TP rate (specificity) vs. FP rate (1-sensitivity). It is constructed by computing pairs like (FP rate, TP rate) for each fold of the data set (in the context of cross-validation) and the points (0,0), $P_1, \dots, P_n, (1,1)$ (where P_i is the point corresponding to the i -th value in increasing order of the FP rate). Ideally would that all points have the coordinates (0,1), and the area under the curve equals the area of the square defined by (0,0) and (1,1), i.e. 1.

For the unbalanced datasets (the number of examples in one class is significantly smaller than the number corresponding to other classes) one use the Matthew correlation coefficient (MCC):

$MCC = (TP * TN - FP * FN) / \sqrt{((TP + FP) * (TP + FN) * (TN + FP) * (TN + FN))}$

Exercise 1: Visualization of decision surfaces

- a) Launch BoundaryVisualizer (Weka GUI Chooser -> Visualization -> Boundary Visualizer)

- b) Open the file “iris2D.arff” (it contains only two attributes: petalLength and petalWidth from iris.arff). Apply the following classifiers: **IBk** (k=1, k=2, k=5), **J48**, **OneR**, **JRip**, **NaiveBayes** and compare the results.

2. Decision trees

A decision tree is a classification tree which contains:

- *Internal nodes*: each internal node has an attribute associated and some splitting conditions (each splitting condition corresponds to a branch in the tree)
- *Leaf nodes*: each leaf node has a class associated

Top-down construction:

- For the root node one identifies the splitting attribute and the splitting conditions which leads to the highest informational gain based on the available dataset
- For each branch recursively apply the same procedure until a stopping condition is identified, e.g. the dataset corresponding to the current node is homogeneous (all data belong to the same class) or it is too small.

Remark: one of the most important aspects when constructing decision trees is the complexity level of the inferred tree. Simpler trees (i.e. fewer nodes) are always preferred. The simplification of the tree (pruning) can be done at different stages:

- During the tree construction (e.g. the branching process is stopped when the number of instances in the current dataset is smaller than a given threshold)
- After the tree construction some subtrees can be replaced with leaf nodes (as long as the classifier performance is not significantly decreased)

Decision tree usage: for a given instance one identify the tree branch which matches the corresponding attribute values; the label of the leaf node in that branch will be the predicted class.

Weka implementation: Id3, J48

Exercise 2: Comparing algorithms for decision trees construction

- a) Open in Weka the file “weather_nominal.arff”. Using Weka Explorer compare the performance of the following classifiers: **Id3**, **J48** (in this case both the “pruned” and “unpruned” variants are analyzed – this option is set in the list of parameters of J48)
- b) Open in Weka the file “weather_numeric.arff” and try to apply the same processing steps as at (a).
- c) Open in Weka the file “glass.arff” and analyze the impact on the J48 performance of the minimal number of objects which correspond to a leaf node (the value of the parameter **minNumObj**)

Exercise : Construction of a binary decision tree by interactive selection of data

[Remark: it can be tested only for versions of Weka older than 3.8]

Open in Weka the file “iris2D.arff” and apply the following processing steps:

- a) Select “**User Classifier**” (group **Tree** from **Classify**)
- b) In the window containing the root node go to **Data visualizer** and select a data subset with a high degree of purity (if it is possible all data should belong to the same class = the points have the same color):
 - from **Select Instance** choose **Rectangle**,
 - extend the rectangular frame over the data to be selected

- click on [Submit](#)

The effect of these steps is that the root node is split in two children nodes: one corresponding to the selected dataset and another one to the rest of data (the tree can be visualized in the panel [Tree visualizer](#)).

- c) Select another subset (out of the remaining data) and continue the process until the entire dataset is covered.

3. Rule based classifiers

The classification rules have the following structure: **IF** <antecedent> **THEN** <class label>, where the antecedent part contains conditions concerning the attribute values (the individual conditions are combined using conjunction and/or disjunction operators)

Construction of the classification rules set. There are two main variants:

- Starting from a *decision tree*: from each branch one constructs a classification rule (the conditions in the antecedent are obtained by using the conjunction operator)
- Directly from the data: by using the so-called *covering algorithms*

Weka implementation:

- Starting from decision trees: M5rules, PART
- Covering algorithms: OneR, JRip

Exercise 3: Analysis of algorithms for extracting classification rules

- a) Open the file “[weather.nominal.arff](#)” and apply successively (using Weka Explorer) the following algorithms: [OneR](#), [JRip](#), [M5Rules](#), [PART](#)
- b) Identify the best rule based classifier for the dataset “[weather.nominal.arff](#)”

4. Instance Based Learning Classifiers

The simplest instance based classifier is the kNN (k nearest neighbours) .

Construction:

- The classification model consists of the training set (there is no specific training step – this is why these classifiers are considered lazy classifiers)

Usage:

- For an input data the class is predicted based on the following steps:
 - Step 1: Find the most similar k instances from the training set
 - Step 2: Identify the dominant class among the classes corresponding to the k instances selected at the previous step

Weka implementation: [IB1](#) (only 1 neighbour is used), [IBk](#) (k neighbours are used)

Exercise 4: Choose the “right” number of neighbors for kNN

- a) Open in Weka the file “[breast-w.arff](#)”. For IBk identify the value(s) of k (out of {1,2,3,4,5,6,7,8,9,10}) which lead(s) to the best classifier performance (Hint: modify the value of k from IBk).
- b) Open in Weka the file “[glass.arff](#)” and apply the same processing steps.

5. Probabilistic models

They are based on the estimation of the probability that a data belongs to a class ($P(C_k|a_1, a_2, \dots, a_n)$). During the training process are computed, based on the training set, some probabilities which are later used to estimate the probability for a given data to belong to each of the classes. The class having the maximal probability will be considered the class to which the data belongs.

Examples:

- **Naïve Bayes:** it is based on the simplifying assumption that all attributes are independent (consequently the probability to observe a data instance is the product of the probabilities to observe each of the attributes, i.e. $P(a_1, a_2, \dots, a_n) = P(a_1)P(a_2) \dots P(a_n)$).
- **Bayesian Network (Belief network):** it allows to describe the relationships between attributes by using a directed acyclic graph.

Exercise 5:

- a) Open the file “[weather.nominal.arff](#)”. Apply the classifiers [Naïve Bayes](#) and [Bayes Net](#)
- b) Use the [Bayes Network Editor](#) (from Weka GUI Choser -> Tools) to construct a Bayesian network. Main steps:
 - Select the dataset: [Tools->Set Data](#)
 - Place the nodes (each attribute, including that corresponding to the class will have an associated node) by using [Edit->Add node](#)
 - Specify the relationships by using [Edit->Add arc](#) (for instance one can specify an arc from the class node to the nodes corresponding to the attributes)
 - Network training ([Tools->Learn network](#) (this corresponds to the learning of the structure), [Tools->Learn CPT](#) (this corresponds to the computation of the tables with probabilities))
 - Usage of the Bayes network classifier: select the values which correspond to the instance to be classified (by right clicking on each node associated to an attribute and by using [Set evidence](#)); the decision is based on the values of the probabilities associated to the class.

Remark: details on using the Weka implementation of Bayesian networks can be find at http://www.cs.waikato.ac.nz/~remco/weka_bn/

6. Neural networks

The main steps in designing a neural network for classification are:

- *The choice of the architecture.* In the case of multi-layer perceptrons one have to choose:
 - The number of input units = the number of attributes
 - The number of hidden layers and units – it depends on the complexity of the problem (there is no unique solution)

- Number of output units:
 - Binary classification: one unit (the result is interpreted by using a threshold value, e.g. if the output value is larger than 0.5 then class 1 otherwise class 2) or two units (the unit which produces the larger value will indicate the class)
 - Multiple classification: The number of output units is equal to the number of classes.
- *The choice of the activation functions.* For the neural networks to be trained using the backpropagation algorithm the activation functions should be differentiable, thus the typical choices are:
 - The hidden units have sigmoidal activation functions (logistic or tanh function).
 - The output functions have sigmoidal or linear functions.
- *The choice of the training algorithm and of its parameters.* In the case of the backpropagation algorithm the typical parameters (also used in Weka) are:
 - Number of training epochs
 - Learning rate
 - Coefficient of the momentum term

Weka implementation: MultilayerPerceptron = feedforward neural network trained using Backpropagation (the variant with momentum term); the network architecture has by default one hidden layer but the number of layers (and of units per each layer) can be changed.

Exercise 6: Analysis of the network architecture on the classification performance

- a) Open the file “[breast-w.arff](#)” \
- b) Train a neural network using the default settings (one hidden layer with $K=(nr\ attributes+nr\ classes)/2$ units)
- c) Compare the performance of the classification (accuracy values) for the following values of the [Hidden Layers \(H\)](#) parameter:
 - a. ‘a’ ($K=(nr\ attributes+nr\ classes)/2$)
 - b. ‘i’ ($K=nr\ attributes$)
 - c. ‘o’ ($K=nr\ classes$)
 - d. ‘t’ ($K = nr\ attributes+nr\ classes$)
 - e. 4,2 (two hidden layers with 4 and 2 units, respectively)
- d) Analyze the influence of the learning rate and of the coefficient of the “momentum” term.
- e) Compare the results obtained by using the [MultilayerPerceptron](#) with that obtained by using a RBF network – radial basis function network ([functions->RBF Network](#))

Remark. By activating the option GUI from Multilayer Perceptron one can visualize the network architecture.

7. Classification using Support Vector Machines (SVM)

The Weka implementation of SVM uses a fast algorithm for solving the optimization problem involved in the estimation of the weights corresponding to support vectors (Sequential Minimal Optimization). The standard SVM is designed for binary classification. In the case of the classification in M classes the classification problem is transformed in several binary classification problems.

Exercise 7:

- a) Open in Weka the file “[breast-w.arff](#)” and use SVM (functions->SMO); analyze the influence of the kernel function on the classification accuracy (variants: [PolyKernel](#), [NormalizedPolyKernel](#), [RBFKernel](#))
- b) Open in Weka the file “[arrhythmia.arff](#)” and apply the SVM classifier; try to solve the same problem using a neural network (multilayer perceptron or RBF network).