Metaheuristic algorithms

Lab 5: Multiobjective optimization.
      Classification problems and classifiers. Weka – data mining platform
_____

    1. **Multiobjective optimization**

Multiobjective optimization means to simultaneously optimize several objective functions (criteria). The function to be optimized is vectorial $F:R^n->R^r$, and its components can be denoted as follows $F=(f_1,f_2,\dots,f_r)$.

The optimization criteria are usually conflicting, therefore the problem does not have a unique solution. In such a case we are looking for some compromise solutions (called Pareto optimal) characterized by the fact that they cannot be improved with respect to all their components (any improvement with respect to one criterion leads to a decrease of quality with respect to other criteria).

There are different approaches of this problem. The main approaches are:

- *Aggregation methods:* the multiobjective problem is transformed in a one-objective optimization problem by combining all optimization criteria in a single one. Thus the new objective function becomes: $f(x)=w_1f_1(x)+w_2f_2(x)+\dots+w_rf_r(x)$ where $w_1,w_2, \dots,w_r$ are weights associated to objective functions. For each set of weights one can obtain a different solution.

- *Direct approximation of the Pareto optimal set:* it uses a population of elements which will approximate the Pareto optimal set. The approximation process can be a evolutionary one. The main difference between multiobjective EAs and single objective EAs is related with the selection process. In the MOEAs the selection process is based on the dominance relationship between the elements (see Lecture 10).

Examples of test functions used to evaluate the performance of multiobjective algorithms are available at: http://en.wikipedia.org/wiki/Test_functions_for_optimization

Application 1. Let us consider the function $F:[0,4]->RxR$, $F(x)=((x-1)^2,(x-2)^2)$. Estimate the optimal Pareto set and the corresponding Pareto front.

Variant 1. By using the aggregation technique

    a) Construct the aggregated objective function:

```
function y=fw(x)
    w=0.1;
    y1=(x-1)*(x-1);
    y2=(x-2)*(x-2);
    y=w*y1+(1-w)*y2;
endfunction
```

    b) Apply an evolution strategy (for instance, that described in SE.sci) to optimize the aggregated objective for the following values of w: (0.1,0.2,0.3,…,0.9). The corresponding results are collected in a list.

c) Plot the points having as coordinates the values of the objective functions computed at the previous step (the plotted set of points will be illustrate an approximation of the Pareto front):

```
function pareto(x)
    f1=(x-1).^2;
    f2=(x-2).^2;
    plot(f1,f2,'*');
endfunction
```

Variant 2. Use the algorithm NSGA-II and MOGA algorithms implemented in SciLab (functions optim_nsga2 and optim_moga).

Hint: exNSGA.sci

2. Classification problems

The main aim of a classification task is to decide to which class a given object belongs. The classification task can be solved by using a classifier which is usually designed starting from some examples of labelled data.

2.1 Examples of classification tasks:
- Medical diagnosis:
    - the input data are symptoms, results of investigations, characteristics of the patient etc
    - the classes are related with the illness (a frequent case is when there are two classes: healthy and ill)
- Fault diagnosis:
    - the input data are characteristics of a given system
    - the classes corresponds to the system status (functional or with flaws)
- Character recognition:
    - the input data are vectors containing features of the characters
    - the classes corresponds to the characters (letter, digits, other symbols) to be recognized.
- Face recognition:
    - The input data are features extracted from digital images
    - The classes corresponds to the presence or absence of human faces in the image

2.2. Design of classifiers. Categories of data used when designing a classifier:

- *Training data*: data for which the corresponding class is known and are used to design the classifier through a learning process
- *Validation data*: data used to evaluate the behavior of the classifier during the learning process; the classification performance on the validation data is used to decide if the learning process should be stopped or not; early stopping of learning could be useful in order to avoid overtraining and enhance the generalization ability of the classifier
- *Testing data*: data used to evaluate the performance of a trained classifier

The classifiers can be of different types:
- Decision trees
- Decision (classification) rules

- Probabilistic approaches (Bayesian networks)
- Instance based classifiers (e.g. based on the nearest distance to some instances from the training set)
- Neural networks, Support Vector Machines

Main steps in designing a classifier:
- Preprocess the available data (e.g. select/ extract the relevant data features)
- Choose the classifier type (e.g. decision trees, classification rules, distance based, network-based etc.)
- Train the classifier, i.e. adapt the structure and/or estimate the adaptive parameters, such that it correctly classifies the data in the training set.
- Validate the classifier, i.e. evaluate the behavior of the classifier for data not included in the training set.

Measures of classifiers quality. The simplest quality measure is the classification accuracy: the ratio between the number of correctly classified instances and the total number of instances.

In the case to binary classification (there are two classes: a class of positive cases and a class of negative cases) the classifier performance can be illustrated using the so-called *confusion matrix*. The confusion matrix contains information about the percent of data which were correctly/ incorrectly classified. Let us consider the case when the first class corresponds to a positive diagnostic (e.g. malignant) and the second class corresponds to a negative diagnostic (e.g. benign). The confusion matrix contains the percent or absolute frequence of:

- *True positive cases* (TP): positive cases which were correctly classified in the first class
- *True negative cases* (TN): negative cases which were correctly classified in the second class
- *False positive cases* (FP): negative cases which were incorectly classified in the first class
- *False negative cases* (FN): positive cases which were incorrectly classified in the second class

Based on these values several other quality criteria can be computed:

sensitivity=TP / (TP + FN) (rmk: also called *recall*)   *specificity=TN/(TN+FP)*

precision=TP/(TP+FP)   *F=2*precision*recall/(precision+recall)*

Another possibility to illustrate the performance of a binary classifier depending on a discrimination threshold is the ROC (Receiver Operating Characteristic) curve which consists of points having the coordinates (*1-specificity*, *sensitivity*) and corresponding to different discrimination thresholds (in the case of binary classification the decision that an input belongs to a given class is usually taken based on a value in [0,1] and a threshold: if the value is larger than the threshold then the input belongs to the class otherwise it does not belong).

TP FP
TN
FN
1 1

100%

P(TP)

0%   P(FP)   100%
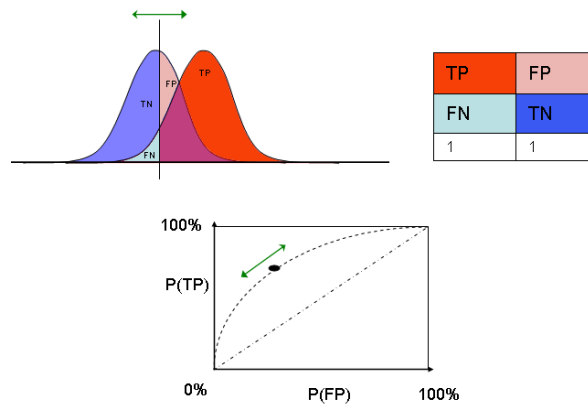
2.3. **Neural Networks for classification problems**

The most common neural network architecture used in solving classification problems is the *feedforward* one characterized by:
- An *input layer* having as many units as attributes are in data
- One or several *hidden layers* (as the number of hidden units is larger, the model extracted by the neural network is more complex – but this is not necessarily beneficial for the problem to be solved as it can lead to model overfitting)
- An *output layer* having as many units as classes

The output signal corresponding to an input vector (that containing the features of the object to be classified) specifies the class to which belongs the object. There are at least two ways of interpreting the outputs of a neural network used for classification:
- In the most common case, the class is specified by the index of the aoutput unit producing the largest value
- In the case when the output values belong to (0,1) and their sum equals 1 then they can be interpreted as probabilities (the value corresponding to unit i corresponds to the probability that the input vector belongs to class i).

If the values produced by the output neurons are in (0,1), the correct answers (to be used in the learning process) are vectors containing one value equal to 1 on the position corresponding to the right class and equal to 0 on all the other positions.

The *learning* process is controlled by a *cross-validation technique* which consists in dividing the initial set of data into three slices:
- *Training data*: the data used in the training algorithm to compute the adjustments for connections weights
- *Validation data*: when the classification error on these data starts to increase, the training process is stopped (these data are not used to compute the weight adjustments but only to decide if the network has generalization ability)
- *Test data*: these data are not used in the training process but only to evaluate the quality of the classifier

2.4. Weka data mining platform (http://www.cs.waikato.ac.nz/ml/weka/)

- Weka is a collection of machine learning algorithms (including several types of neural networks: multilayer perceptron, radial basis functions networks) designed to solve various data mining tasks: data visualization, attribute selection, classification, clustering, association rules extraction.

- It is open source and the methods can be used from the Explorer GUI or directly from a Java code. It allows to compare different methods (Experimenter GUI) or to define workflows of data mining tasks (Knowledge Flow GUI).

Exercise 1. Open the file *breast_cancer_nominal.csv* in Weka and compare the accuracy of the following classifiers:
- Classification rules: Rules-> OneR, Rules-> ConjunctiveRules, Rules->NNge
- Decision trees: Trees->J48, Trees->RandomForest
- Nearest Neighbor: Lazy->IBk
- Bayesian Networks: bayes->NaiveBayes
- Neural Networks: functions->MultilayerPerceptron, functions->RBFnetwork,
- Support Vector Machines: functions->SMO

Hint: All classifiers are in the Classify panel of the Weka Explorer. For neural networks classifiers in order to see a graphical representation of the network architecture set GUI option in the panel opened when click on the classifier bar.

2.5. Machine Learning Repository   (http://archive.ics.uci.edu/ml/)

This archive contains a lot of data from various domains (life sciences, physical sciences, computer science and engineering) which can be used to train, validate and test classifiers and other machine learning techniques.

Exercise 2. Visit the repository, download one of the dataset for classification (at your choice) and use it to test classifiers from Weka (see Ex. 1)