
Assignment 5


Introduction to Data Analytics

Prof. Nandan Sudarsanam & Prof. B. Ravindran

1. In a binary classification scenario where x is the independent variable and y is the dependent variable, logistic regression assumes that the conditional distribution $y|x$ follows a
 - (a) Bernoulli distribution
 - (b) binomial distribution
 - (c) normal distribution
 - (d) exponential distribution
2. To control the size of the tree, we need to control the number of regions. One approach to do this would be to split tree nodes only if the resultant decrease in the sum of squares error exceeds some threshold. For the described method, which among the following are true?
 - (a) it would, in general, help restrict the size of the trees
 - (b) it has the potential to affect the performance of the resultant regression/classification model
 - (c) it is computationally infeasible
3. For a K-NN classification model, as the value of K increases
 - (a) bias increases and variance decreases
 - (b) bias decreases and variance increases
4. For a supervised learning problem, we have seen that the data used in the testing phase, i.e., the test set, is not used for training or building the model. Does this mean that we do not require labels for data points in the test set?
 - (a) no
 - (b) yes
5. For a two-class classification problem, we use an SVM classifier and obtain the following separating hyperplane.

We have marked 4 instances of the training data. Identify the point which will have the most impact on the shape of the boundary on its removal.

 - (a) 1



- (b) 2
 - (c) 3
 - (d) 4
6. For a linearly non-separable data set, are the points which are misclassified by the SVM model support vectors?
- (a) no
 - (b) yes
7. In the linearly non-separable case, what effect does the C parameter have on the SVM model?
- (a) it determines the count of support vectors
 - (b) it is a count of the number of data points which do not lie on their respective side of the hyperplane
 - (c) it determines how many data points lie within the margin
 - (d) it allows us to trade-off the number of misclassified points in the training data and the size of the margin
8. Suppose that we use a radial basis function kernel with appropriate parameters to perform classification on a particular two class data set where the data is not linearly separable. In this scenario
- (a) the decision boundary in the original feature space is linear
 - (b) the decision boundary in the original feature space is non-linear
 - (c) the decision boundary in the transformed feature space is linear
 - (d) the decision boundary in the transformed feature space is non-linear
9. Consider the following data set:

price	maintenance	capacity	airbag	profitable
low	low	2	no	yes
low	med	4	yes	no
low	low	4	no	yes
low	high	4	no	no
med	med	4	no	no
med	med	4	yes	yes
med	high	2	yes	no
med	high	5	no	yes
high	med	4	yes	yes
high	high	2	yes	no
high	high	4	yes	yes
high	high	5	yes	yes

Considering 'profitable' as the binary values attribute we are trying to predict, which of the attributes would you select as the root in a decision tree with multi-way splits using the cross-entropy impurity measure?

- (a) price
 - (b) maintenance
 - (c) capacity
 - (d) airbag
10. For the same data set, suppose we decide to construct a decision tree using binary splits and the Gini index impurity measure. Which among the following feature and split point combinations would be the best to use as the root node assuming that we consider each of the input features to be unordered?
- (a) price - {low, med}|{high}
 - (b) maintenance - {high}|{med, low}
 - (c) maintenance - {high, med}|{low}
 - (d) capacity - {2}|{4, 5}

Weka-based programming assignment questions

For the questions on SVMs, you will need to download the LibSVM package and use it along with Weka. You can install the package using Weka's package manager (in the Tool's menu of Weka's GUI Chooser window), or follow the instructions given here.

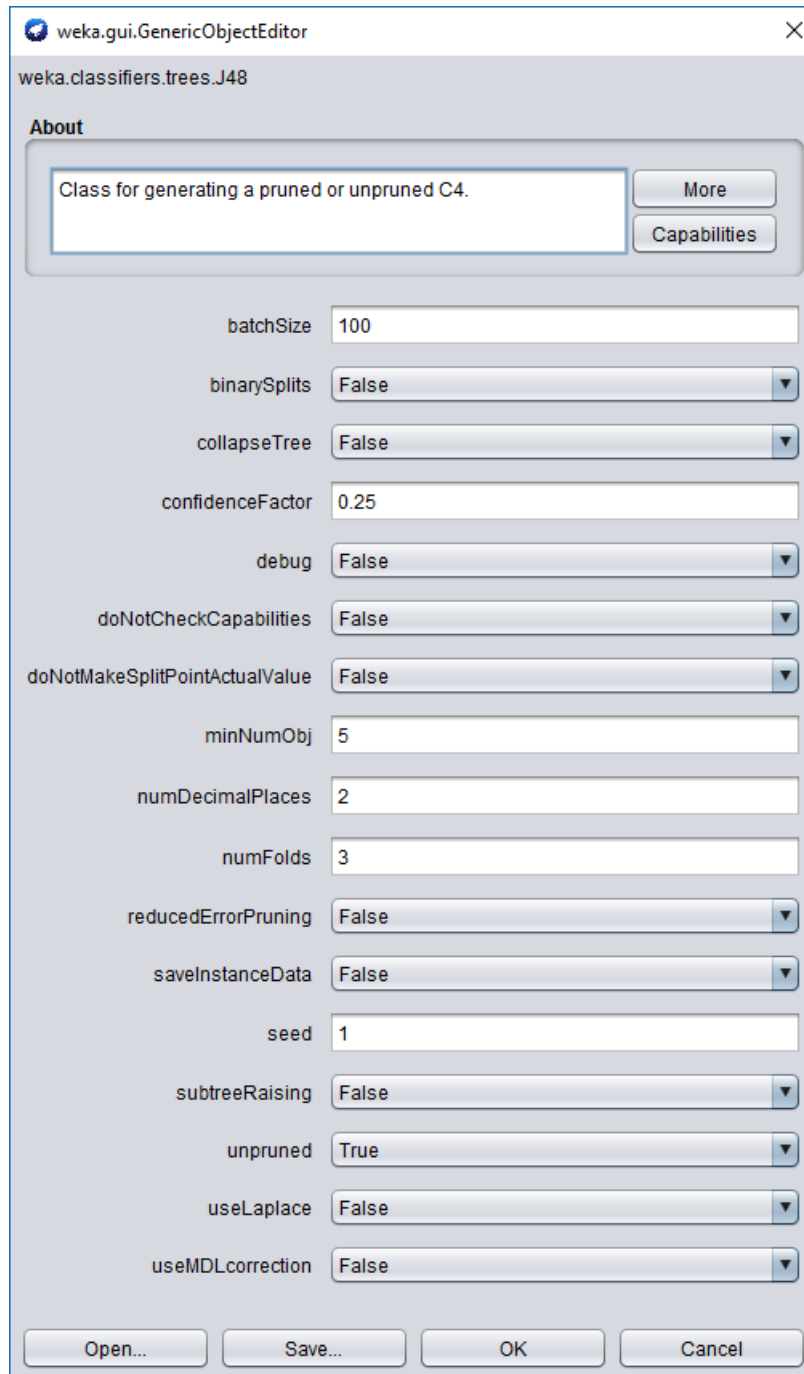
Once successfully installed, you will find the option LibSVM under 'functions' when choosing the classifier. The options that we will need to considered are debug (set to True), degree (for use with polynomial kernel), and kernelType (for choosing the kernel). Note that you will need to launch Weka from the command line (or using the 'Weka 3.8 (with console)' shortcut) to view the output from the optimiser where the parameter nSV will indicate the number of support vectors. Make use the visualise options to view the data. As in the last assignment, we have given both train and test sets for each data set.

11. For data set 1, train linear and radial basis function kernel SVMs (default parameter values). What are the number of support vectors in each case?

- (a) 100, 100
 - (b) 3, 104
 - (c) 3, 116
 - (d) 550, 100
12. For data set 2, train 5 degree polynomial (5 degree, $\text{coef0} = 0$), 10 degree polynomial (10 degree, $\text{coef0} = 0$) and radial basis kernel functions (default parameter values). What are the number of support vectors in each case?
- (a) 230, 20, 27
 - (b) 324, 20, 27
 - (c) 425, 21, 25
 - (d) 362, 20, 25
13. In the previous question, which classifier among the three would you prefer to use for unseen data?
- (a) 5 degree polynomial
 - (b) 10 degree polynomial
 - (c) radial basis function

For the questions on decision tree, we will be using the UCI Tic-Tac-Toe Endgame data set (Dataset5). Use the decision tree model J48 under the trees folder. Make use of 10-fold cross validation for all experiments with this data set. Also, make use of the visualise tree option (by right clicking in the results list) once a decision tree model is built.

We will consider the following to be the default parameter settings:



14. How many levels does the unpruned tree contain considering multi-way and binary splits respectively, with the other parameters remaining the same as above?

(a) 10, 6

- (b) 6, 8
 - (c) 8, 6
 - (d) 6, 10
15. How many levels does the pruned tree (unpruned = false, reducedErrorPruning = false) contain considering multi-way and binary splits respectively?
- (a) 10, 6
 - (b) 6, 8
 - (c) 8, 6
 - (d) 6, 10
16. Considering the weighted average F-measure as the performance indicator (higher values indicate better performance), do the pruned trees result in increased performance or decreased performance in the multi-way and binary split cases respectively?
- (a) decrease, increase
 - (b) increase, decrease
 - (c) increase, increase
 - (d) decrease, decrease
17. Which among the following attributes seems to be the most important for this particular classification task?
- (a) top-left-square
 - (b) top-right-square
 - (c) middle-middle-square
 - (d) bottom-middle-square