

SKETCH of the SOLUTIONS

STK2100 - spring 2019

Exercise 1

- a There is no protecting factor, as all regression coefficients are larger than 0.

The intercept includes the baseline effects, in this case the log-odds for a female, non-smoker who does not drink and has age = 0.

The latter is the reason why the intercept does not have a “physical” meaning in this case.

A solution is to center age, so the baseline is a female with average age. Since the regression coefficient is not significant, it may also be excluded from the model (although, strictly speaking, in this way the problem is removed, more than solved...).

- b The correct answer is model 2, as it has the best performance on the test set in terms of the ROC-curve (valid explanations: largest area under the curve, closest curve to the top-left corner, farthest curve from the intercept).

The ROC curve is a graphical tool to visualize the performance of a classification model, and it displays the sensitivity and (1 minus) specificity when moving the threshold used to discriminate between the two response classes.

Sensitivity = true positive / (true positive + false negative), where true positive are the observations correctly identified as positive by the model, and false negative the observations incorrectly classified as negative by the model.

Specificity = true negative / (false positive + true negative), where true negative are the observations correctly identified as negative by the model, and false positive the observations incorrectly classified as positive by the model.

- c The plot shows how many times a classification based on the selected model is better than a random choice when classifying a specific amount of observations.

In particular, for the highlighted point, this means that if we want to classify 35% of the observations, if we select those to which the model gives the highest probability to be, let us say, positive, we identify 1.6 times more positive than when repeating the procedure at random.

d The values in the table can be computed as:

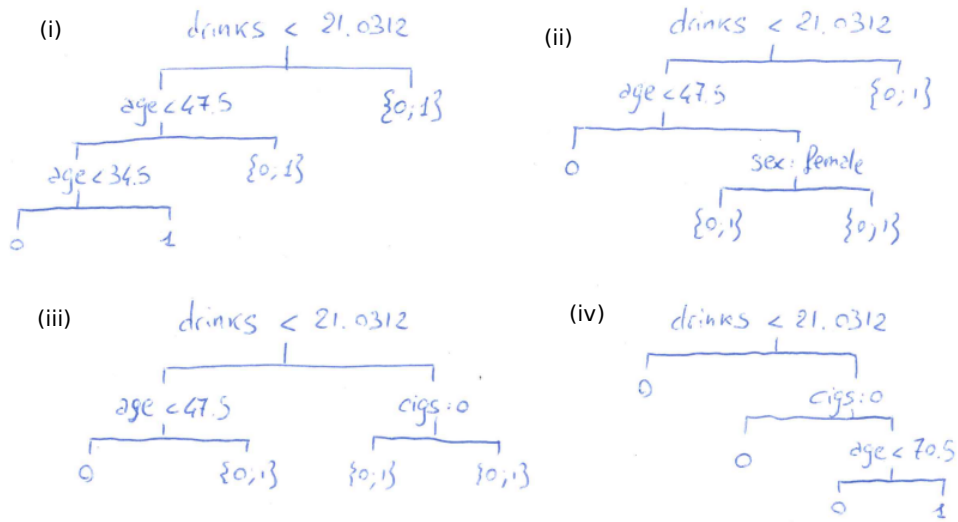
- $a_{21} = 47 \times \exp\{-0.96141\}/(1 + \exp\{-0.96141\}) \approx 13$;
- $a_{22} = 152 \times \exp\{-0.96141+1.11964\}/(1+\exp\{-0.96141+1.11964\})$
 ≈ 82 ;
- $a_{11} = 47 - 13 = 34$;
- $a_{12} = 152 - 82 = 70$.

e The predicted response for a 40-year old man who smokes and drinks, on average, 16 ounces of alcoholic drinks per week is 0 (no oral cancer). It could be found following the path: left (drink less than 21.0312 ounces per week), left (younger than 47.5 years), right (older than 34.5 years).

Cross-validation is a procedure to estimate the prediction error, often used to select the tuning parameter(s) of a model (here the number of leaves). It consists in splitting the observations in K approximately equal-sized folds and using in turn one fold as an independent test set, in which the performance of a model (or, more generally, a method) trained on the remaining $K - 1$ folds is evaluated. The K estimates of the performance measure (e.g., deviance) are then summarized (e.g., by averaging) into a single measure. If the goal is finding the best parameter, the procedure is performed for each parameter and that corresponding to the best performance is selected.

Here cross-validation is used to select the tuning parameter. In order to do this, it is important to evaluate the results obtained with a specific value of the tuning parameter in an independent set, to avoid overfitting. Having a moderately low number of observations (less than 200, see point d), it is not reasonable to split them in separated training and test sets. Cross-validation is used because it allows to take advantage of all observations to find the tuning parameter while keeping the independence between training and test set.

There are at least 4 possible ways to prune the tree while keeping the right response for the described subject. All the following (see next page) choices are correct:



Exercise 2

- a The two terms are the dissimilarity within groups and the dissimilarity between groups, respectively. The former is the quantity of dissimilarity present within the clusters, i.e. the sum of the dissimilarities between all observations belonging to the same group. The latter is the sum of all dissimilarities between observations belonging to different groups.

Since they sum to a fixed quantity, the total dissimilarity, that does not depend on the grouping structure, minimizing one corresponds to maximize the other.

A normalization procedure is usually necessary to avoid issues related to different scales (in the case of continuous variables) or different number of levels (in the case of categorical variables) in the dimensions (variables) used to separate the observations.

- b The algorithm is the “K-means”.

In step 2.(a) the observations are allocated to the best group based on their dissimilarities with the cluster centroids, guaranteeing that the total dissimilarity within groups is minimal once the centroids have been chosen.

In step 2.(b) the centroids are updated based on the current grouping, guaranteeing that the total dissimilarity within groups is minimal once the observations are allocated to the clusters.

The main limitation of the algorithm is that the number of clusters must be pre-specified. Other limitations are: it only works for quantitative variables; there is no guarantee to reach a global minimum.

- c The single link and the complete link measure the dissimilarity between two clusters in two different ways. The former by computing the smallest dissimilarity between two observations belonging to separate clusters, the latter measures the cluster dissimilarity as the largest dissimilarity between two observations belonging to separate clusters.

The different ways of measuring dissimilarities impact the process of agglomeration or division of the groups performed by the clustering algorithm, and may therefore lead to different clustering. In particular, the single link tends to be more efficient in recovering filiform group structures, while the complete link tends to work better with spheroidal structures.

Any choice of K between 2 and 10 is reasonable. Using the dendrogram, the idea is to select a cut-off where the vertical lines are longer, as they represent the improvements in terms of reduced dissimilarity when increasing the number of clusters.

Even dividing the observations in independent training and test sets, a larger number of clusters will always result in a smaller dissimilarity within clusters. Since there is no response, indeed, any additional division will represent an improvement in the test set as well. Basically, it does not make sense to implement a cross-validation procedure because clustering is an unsupervised task.

Exercise 3

- a In local regression methods h is a tuning parameter that controls the width of the window used to locally estimate the function. Smaller values of h correspond to smaller windows, and consequently to a reduced number of points used to estimate the value of interest. Larger values of h , instead, correspond to larger windows, and consequently to estimates that follow less the local behaviour of the curve, as they also take into account farthest points.

Therefore, smaller values of h mean smaller bias and larger variance (the estimated curve follows the local behaviour of the data, with the risk, if h is too small, of modelling random fluctuations). In contrast, larger values mean increased bias but reduced variance.

The concept of bias-variance trade-off is crucial to define the best prediction model. Basically, one desires a model which explains well the data (low bias) without following their random fluctuations (low variance). Unfortunately the two aspects are in contrast: decreasing the bias results in an increase of the variance, and the other way around. Both extremes, low bias but too much variance (overfitting) and low variance but too much bias (underfitting), should be avoided, and therefore the best balance (trade-off) should be found.

- b The “curse of dimensionality” defines the phenomenon for which the number of observations close to a point decreases exponentially with the increase of the number of dimensions. As a consequence, non-parametric estimations have to rely on fewer points, with an obvious decrease in precision, or farther points, which are less likely to be similar to the point under investigation.

A possible solution is to concentrate the information in fewer dimensions, for example by applying techniques such as principal component analysis.