

Decision Trees

1 Classification Trees With Numerical Features

In this problem you will be working with two datasets:

- **Iris:** has three classes and the task is to accurately predict one of the three sub-types of the Iris flower given four different physical features. These features include the length and width of the sepals and the petals. There are a total of 150 instances with each class having 50 instances.
- **Spambase:** is a binary classification task and the objective is to classify email messages as being spam or not. To this end the dataset uses fifty seven text based features to represent each email message. There are about 4600 instances

Since, both datasets have continuous features you will **implement** decision trees that have **binary splits**. For determining the optimal threshold for splitting you will need to search over all possible thresholds for a given feature (refer to class notes and discussion for an efficient search strategy). Use **information gain** to measure node impurity in your implementation.

1.1 Growing Decision Trees (*60 points*)

Instead of growing full trees, you will use an early stopping strategy. To this end, we will impose a limit on the minimum number of instances at a leaf node, let this threshold be denoted as η_{min} , where η_{min} is described as a percentage relative to the size of the training dataset. For example if the size of the training dataset is 150 and $\eta_{min} = 5$, then a node will only be split further if it has more than eight instances.

- (a) For the Iris dataset use $\eta_{min} \in \{5, 10, 15, 20\}$, and calculate the accuracy using ten fold cross-validation for each value of η_{min}
- (b) For the Spambase dataset use $\eta_{min} \in \{5, 10, 15, 20, 25\}$, and calculate the accuracy using ten fold cross-validation for each value of η_{min}

You can summarize your results in two separate tables, one for each dataset (report the average accuracy and standard deviation across the folds).

1.2 Interpreting the results (*Extra-credit*)

Select the best value of η_{min} for the Iris dataset, and create a class confusion matrix using ten-fold cross validation (*use only the test set for populating the confusion matrix*). How do you interpret the confusion matrix, and why?

2 Classification Trees With Categorical Features (*Extra-credit*)

In this problem you will be working with a new dataset:

- **Mushroom:** is binary classification dataset and the task is to accurately predict whether a mushroom is poisonous or edible given 21 different categorical (ordinal) features for each mushroom. These features describe various physical properties of the mushrooms such as length, diameter, etc. There are a total of 8124 instances.

As this dataset has only ordinal features we will **implement** both binary and multiway decision trees. Use **information gain** to measure node impurity in your implementation.

2.1 Multiway vs Binary Decision Trees (*50 points*)

In this problem we will grow both multiway and binary decision trees using categorical features. Similar to the last problem we will be using an early stopping strategy to keep our decision tree from overfitting the data.

- (a) Grow a multiway decision tree using $\eta_{min} \in \{5, 10, 15\}$, and calculate the accuracy using ten fold cross-validation for each value of η_{min} .
- (b) Replace each categorical feature $F \in \{f_1, f_2, \dots, f_v\}$ with v binary features, corresponding to each distinct value of the feature. For example given the second feature *cap-surface*, that takes values $\{\text{fibrous}=\text{f}, \text{grooves}=\text{g}, \text{scaly}=\text{y}, \text{smooth}=\text{s}\}$, we will replace the original feature with four new boolean features: $\{\text{is cap-surface}=\text{f}\}$, $\{\text{is cap-surface}=\text{g}\}$, $\{\text{is cap-surface}=\text{y}\}$ and $\{\text{is cap-surface}=\text{s}\}$. Only one of these features will be 1 for a given instance and the rest will be 0. Grow a binary decision tree with this new extended feature set, using $\eta_{min} \in \{5, 10, 15\}$, and calculate the accuracy using ten fold cross-validation for each value of η_{min} .

You can summarize your results in two separate tables, one for each dataset (report the average accuracy and standard deviation across the folds). Do you see a difference between the two approaches? Explain your answer briefly based on the results.

2.2 Interpreting the results

Select the best value of η_{min} for the the above two cases i.e., multiway and binary splits, and create a class confusion matrix using ten-fold cross validation (*use only the test set for populating the confusion matrix*). How do you interpret the confusion matrix, and why?

3 Entropy (20 points)

Consider training a binary decision tree using entropy splits.

- (a) Prove that the decrease in entropy by a split on a binary yes/no feature can never be greater than 1 bit.
- (b) Generalize this result to the case of arbitrary multiway branching.

4 Gain and Impurity Measures (20 points)

Consider an arbitrary impurity measure $\iota(q)$, where q is a node in the decision tree and a feature V that takes $|V|$ distinct values. For a split resulting in $|V|$ children, the gain for the impurity measure is:

$$Gain(q, V) = \iota(q) - \sum_{i=1}^{|V|} \frac{N_i}{N_q} \iota(i)$$

where N_i is the number of instances in the i^{th} child node. Show that maximizing the $Gain(., .)$ is equivalent to minimizing the impurity measure $\iota(.)$ over the $|V|$ children.

5 Regression Trees (50 points)

In this problem you will **implement** regression trees using a new dataset:

- **Housing:** This is a regression dataset where the task is to predict the value of houses in the suburbs of Boston based on thirteen features that describe different aspects that are relevant to determining the value of a house, such as the number of rooms, levels of pollution in the area, etc.
- (a) As this dataset has only numerical features we will be growing decision trees using only binary splits. Use **mean squared error (MSE)** to define the splits. Use an early stopping strategy similar to the previous decision tree problems and use $\eta_{min} \in \{5, 10, 15, 20\}$. Calculate the MSE using ten fold cross-validation for each value of η_{min} and report the average and standard deviation across the folds (summarize your results in a table) (40 points).
 - (b) Does η_{min} impact the results significantly? Briefly explain your answer (10 points).