

Intelligent Systems
Exercise sheet 10
Rule Learning

Exercise 1¹ (8 points)

Give decision trees to represent the following boolean functions:

- (a) $A \wedge \neg B$
- (b) $A \vee [B \wedge C]$

Exercise 2¹ (10 points)

Consider the following set of training examples:

Instance	Classification	a_1	a_2
1	+	T	T
2	+	T	T
3	-	T	F
4	+	F	F
5	-	F	T
6	-	F	T

Figure 1: Test set of training examples

- (a) What is the entropy of this collection of training examples with respect to the target function "Classification"?
- (b) What is the information gain of a_2 relative to these training examples?

¹ Exercise from T. Mitchell

Exercise 3¹ (13 points)

Experiment with the PlayTennis data

- (a) Implement the basic ID3 algorithm (for discrete attributes only).
- (b) Take a look at the training data in Figure 2, and then run your decision tree learning program on all the training examples to be sure it produces the correct decision tree for this data (it should produce the decision tree shown in Figure 3). Now try running it on a randomly chosen subset containing half of the examples for training, and using half for test. What are the training and test accuracies?

<i>Day</i>	<i>Outlook</i>	<i>Temperature</i>	<i>Humidity</i>	<i>Wind</i>	<i>PlayTennis</i>
D1	Sunny	Hot	High	Weak	No
D2	Sunny	Hot	High	Strong	No
D3	Overcast	Hot	High	Weak	Yes
D4	Rain	Mild	High	Weak	Yes
D5	Rain	Cool	Normal	Weak	Yes
D6	Rain	Cool	Normal	Strong	No
D7	Overcast	Cool	Normal	Strong	Yes
D8	Sunny	Mild	High	Weak	No
D9	Sunny	Cool	Normal	Weak	Yes
D10	Rain	Mild	Normal	Weak	Yes
D11	Sunny	Mild	Normal	Strong	Yes
D12	Overcast	Mild	High	Strong	Yes
D13	Overcast	Hot	Normal	Weak	Yes
D14	Rain	Mild	High	Strong	No

Figure 2 - Training data

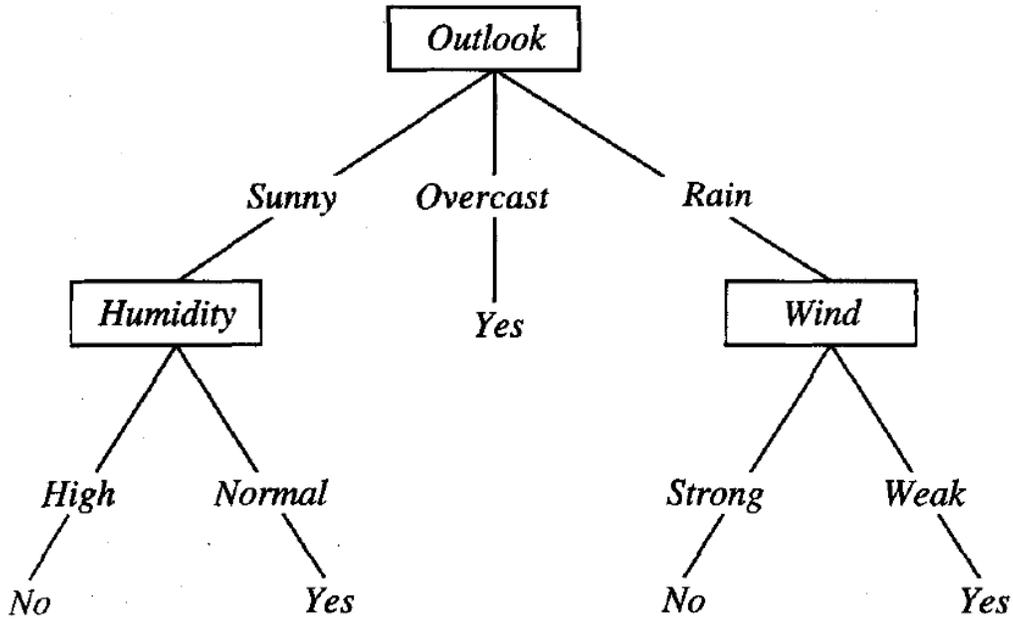


Figure 3 - PlayTennis decision tree

- (c) Answer whether or not the actions in the next three questions are possible or not. If possible, demonstrate it by running the decision tree learner, and turn in the parameter settings you used and your data set. If impossible, explain why. When answering the questions below, assume that (i) the target concept is the concept described by the decision tree of Figure 3; (ii) all training examples you add must be consistent with this target concept (meaning that the classification you include for the example must be the one with which the tree in Figure 3 would label it), and (iii) you are not allowed to include post pruning when learning decision trees.
- Is it possible to get ID3 to further elaborate the tree below the rightmost leaf in Figure 3 (and make no other changes to the tree), by adding a single new correct training example to the original fourteen examples?
 - Is it possible to get ID3 to learn an incorrect tree (i.e., a tree that is not equivalent to the target concept of Figure 3) by adding new correct training examples to the original fourteen?
 - Is it possible to produce some set of correct training examples that will get ID3 to include the attribute "Temperature" in the learned tree, even though the true target concept is independent of "Temperature"?