

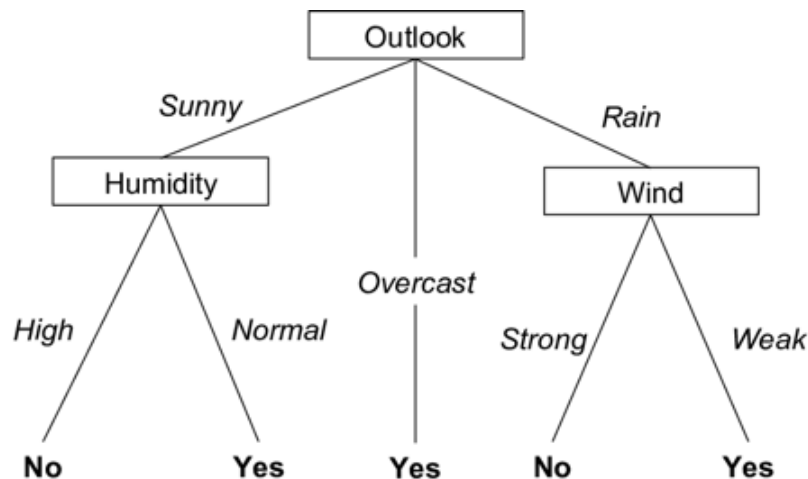
Intro to Decision Trees

(find and work with a partner)

1. First, what is n (number of data points)? What is p (number of features)? Given the training data and decision tree shown below, what is the classification error on this data?

Day	Outlook	Temperature	Humidity	Wind	PlayTennis (y)
x_1	Sunny	Hot	High	Weak	No
x_2	Sunny	Hot	High	Strong	No
x_3	Overcast	Hot	High	Weak	Yes
x_4	Rain	Mild	High	Weak	Yes
x_5	Rain	Cool	Normal	Weak	Yes
x_6	Rain	Cool	Normal	Strong	No
x_7	Overcast	Cool	Normal	Strong	Yes
x_8	Sunny	Mild	High	Weak	No
x_9	Sunny	Cool	Normal	Weak	Yes
x_{10}	Rain	Mild	Normal	Weak	Yes
x_{11}	Sunny	Mild	Normal	Strong	Yes
x_{12}	Overcast	Mild	High	Strong	Yes
x_{13}	Overcast	Hot	Normal	Weak	Yes
x_{14}	Rain	Mild	High	Strong	No


2. On the tree below, the children of each node divide the training data into partitions. Label each node (both internal nodes and leaves) with the counts of “No” and “Yes” labels based on the partition. For example, the counts for the node labeled *Outlook* would be [5, 9].



3. What if we had restricted the tree’s *depth* to be 1? What would the tree look like and what would be the classification error?

Data from Machine Learning by Tom Mitchell (Table 3.2)

4. For the dataset below, the label $y \in \{0, 1\}$. What is n ? What is p ? Devise a decision tree for this data that perfectly classifies the given examples. Internal node labels should be of the form " $x_j \leq a$ ", where a is some constant.



5. Repeat Question (2) for this decision tree (i.e. label each node with the "0" and "1" counts.)