

Introduction to Decision Trees

Alice Gao

Lecture 8

Readings: R & N 18.3

Based on work by K. Leyton-Brown, K. Larson, and P. van Beek

Outline

Learning Goals

Introduction to Decision Trees

The Order of Testing Features

Revisiting the Learning goals

Learning Goals

By the end of the lecture, you should be able to

- ▶ Describe the components of a decision tree.
- ▶ Construct a decision tree given an order of testing the features.
- ▶ Determine the prediction accuracy of a decision tree on a test set.
- ▶ Trace the execution of and implement the ID3 algorithm.

What is a decision tree

- ▶ Book a flight
- ▶ Handle late night cravings
- ▶ Decide how to spend Friday evening

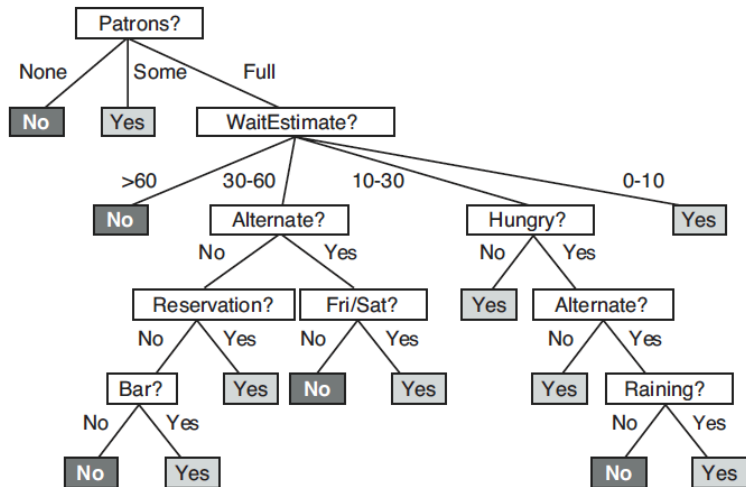
Jeeves the valet - training set

Day	Outlook	Temp	Humidity	Wind	Tennis?
1	Sunny	Hot	High	Weak	No
2	Sunny	Hot	High	Strong	No
3	Overcast	Hot	High	Weak	Yes
4	Rain	Mild	High	Weak	Yes
5	Rain	Cool	Normal	Weak	Yes
6	Rain	Cool	Normal	Strong	No
7	Overcast	Cool	Normal	Strong	Yes
8	Sunny	Mild	High	Weak	No
9	Sunny	Cool	Normal	Weak	Yes
10	Rain	Mild	Normal	Weak	Yes
11	Sunny	Mild	Normal	Strong	Yes
12	Overcast	Mild	High	Strong	Yes
13	Overcast	Hot	Normal	Weak	Yes
14	Rain	Mild	High	Strong	No

Jeeves the valet - test set

Day	Outlook	Temp	Humidity	Wind	Tennis?
1	Sunny	Mild	High	Strong	No
2	Rain	Hot	Normal	Strong	No
3	Rain	Cool	High	Strong	No
4	Overcast	Hot	High	Strong	Yes
5	Overcast	Cool	Normal	Weak	Yes
6	Rain	Hot	High	Weak	Yes
7	Overcast	Mild	Normal	Weak	Yes
8	Overcast	Cool	High	Weak	Yes
9	Rain	Cool	High	Weak	Yes
10	Rain	Mild	Normal	Strong	No
11	Overcast	Mild	High	Weak	Yes
12	Sunny	Mild	Normal	Weak	Yes
13	Sunny	Cool	High	Strong	No
14	Sunny	Cool	High	Weak	No

What is a decision tree?



How to construct a decision tree

- ▶ At each node, choose a remaining feature.
- ▶ Create arcs, one for each value of the feature.
- ▶ Each child is a subtree with all the remaining examples such that the feature takes the value on the arc.

Example: Decision tree

Construct a decision tree using the following order of testing features:

- ▶ Test Outlook first.
- ▶ For Outlook=Sunny, test Temp.
- ▶ For Outlook=Rain, test Wind.
- ▶ For any further branches, always test Humidity before testing Wind.

CQ: Testing accuracy

CQ: What is the testing accuracy of the decision tree you just generated?

Test Set

Day	Outlook	Temp	Humidity	Wind	Tennis?
1	Sunny	Mild	High	Strong	No
2	Rain	Hot	Normal	Strong	No
3	Rain	Cool	High	Strong	No
4	Overcast	Hot	High	Strong	Yes
5	Overcast	Cool	Normal	Weak	Yes
6	Rain	Hot	High	Weak	Yes
7	Overcast	Mild	Normal	Weak	Yes
8	Overcast	Cool	High	Weak	Yes
9	Rain	Cool	High	Weak	Yes
10	Rain	Mild	Normal	Strong	No
11	Overcast	Mild	High	Weak	Yes
12	Sunny	Mild	Normal	Weak	Yes
13	Sunny	Cool	High	Strong	No
14	Sunny	Cool	High	Weak	No

(A) 11/14

(B) 12/14

(C) 13/14

(D) 14/14

A different order of testing the features

What if we use a different order of testing the features?

What would the decision tree look like?

What would be the testing accuracy of the resulting decision tree?

Observations

- ▶ The order of testing features affects the size of the decision tree.
- ▶ Many decision trees are consistent with the training examples.

Which decision tree would you prefer?

Revisiting the Learning Goals

By the end of the lecture, you should be able to

- ▶ Describe the components of a decision tree.
- ▶ Construct a decision tree given an order of testing the features.
- ▶ Determine the prediction accuracy of a decision tree on a test set.
- ▶ Trace the execution of and implement the ID3 algorithm.