Artificial Intelligence

Exercise sheet 4
Software Agents & Rule Learning

1 Software Agents

Note: This is a group task, you may solve it in groups up to the size of two people, everybody must however be able to explain all parts of the implementation

The following exercise is adopted from the book "Artificial Intelligence: A Modern Approach" [1].

Exercise 1.1 - Simulator (10 points)

Implement a performance-measuring environment simulator for the vacuum-cleaner world. This world can be described as follows:

- **Percepts**. Each vacuum-cleaner agent gets a three-element percept vector on each turn:
  1. A touch sensor, should be a 1 if the machine has bumped into something and a 0 otherwise
  2. A photosensor under the machine, which emits a 1 if there is dirt there and a 0 otherwise
  3. An infrared sensor, which emits a 1 when the agent is in its home location and a 0 otherwise

[1] Chapter 2 (Intelligent Agent), Stuart Russell and Peter Norvig, 1995

- **Actions**. There are five actions available:
  1. Go forward
  2. Turn right by 90^o
  3. Turn left by 90^o
  4. Suck up dirt
  5. Turn off
- **Goals**. The goal for each agent is to **clean up** and **go home**. To be precise, the performance measure will be:
  - Earn 100 points for each piece of dirt vacuumed up
  - Minus 1 point for each action taken
  - Minus 1000 points if it is not in the home location when it turns itself off
- **Environment**. The environment consists of a **grid of squares**:
  - Some squares contain **obstacles** (walls and furniture) and other squares are **open space**
  - Some of the open squares contain **dirt**
  - Each **go forward** action moves one square unless there is an obstacle in that square, in which case the agent stays where it is, but the touch sensor goes on
  - A **suck up dirt** action always cleans up the dirt
  - A **turn off** command ends the simulation

Create a vacuum-cleaner world consisting of a at least 5 x 5 grid of squares, at least 8 squares contain obstacles and at most 2 squares will contain dirt. Start the vacuum-agent from upper-left corner facing right, run the simulation and record its performance.

Repeat the simulation for at least 3 different worlds and compare the performances.

2 Rule Learning

> *Note: This is a group task, you may solve it in groups up to the size of two people, everybody must however be able to explain all parts of the implementation*

The following exercises are adopted from the book "Machine Learning" [2].

[2] Machine Learning, Tom Mitchell, McGraw Hill, 1997

Exercise 2.1 (3 points)

Give decision trees to represent the following boolean functions:

1. $A \wedge \neg B$
2. $A \vee [B \wedge C]$
3. $[A \vee B] \wedge [C \vee D]$

Exercise 2.2 (5 points)

Consider the following set of training examples:

Instance	A_1	A_2	Classification
1	T	T	+
2	T	T	+
3	T	F	-
4	F	F	+
5	F	T	-
6	F	T	-

1. What is the entropy of this collection of training examples with respect to the target function "Classification"?
2. What is the information gain of A_2 relative to these training examples?

Exercise 2.3 (12 points)

Experiment with the PlayTennis data set as shown at figure 2.1.

1. Implement the basic decision tree learning program ID3 algorithm
2. Run your algorithm on all the training examples from figure 2.1. It should produce the decision tree shown in figure 2.2)
3. Try running the algorithm on a randomly chosen subset containing half of the examples for training, and using half for test. What are the training and test accuracies?

Day	Outlook	Temperature	Humidity	Wind	PlayTennis
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 2.1: Training Data

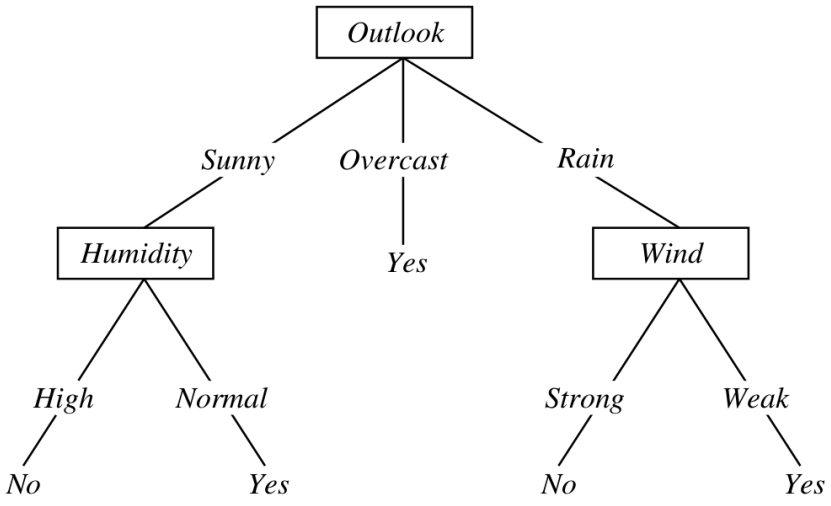

Figure 2.2: Decision Tree