


Artificial Intelligence 691-B

Basic Textbooks



Artificial Intelligence :
Structures and Strategies for Complex Problem Solving
Fifth Edition by George F. Luger
(2005, Addison Wesley)




Artificial Intelligence :
A Modern Approach
Second Edition by Stuart J. Russell & Peter Norvig
(2003, Prentice Hall)

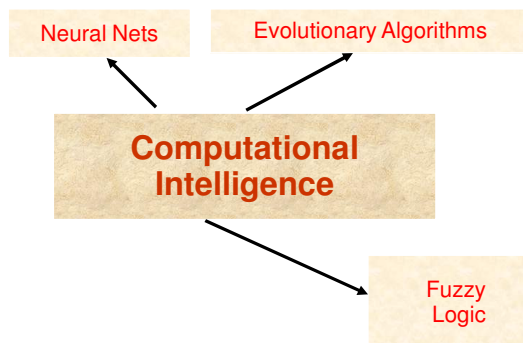
Lecture Notices will be available on the cu scholar website
<http://scholar.cu.edu.eg/?q=areegsaid/classes>

1

Recall



"Soft Computing"



Machine Learning (ML)

- ❑ Machine Learning is the study of how to build computer systems that adapt and improve with experience. It is a subfield of Artificial Intelligence and intersects with cognitive science, information theory, and probability theory, ...etc
- ❑ Intelligent agents must be able to manage through the course of their interactions with the world, as well as through the experience of their own internal states and processes.
- ❑ Simon's definition (1983) describes learning as allowing the system to "perform better the second time." : "Learning is any process by which a system improves performance from experience."

Task of ML

- ❑ **Prediction:** To predict the desired output for a given input based on previous input/output pairs. E.g., to predict the value of a stock given other inputs like interest rates etc.
- ❑ **Categorization:** To classify an object into one of several categories based on features of the object. E.g., a spam email based on subject
- ❑ **Clustering:** To organize a group of objects into homogeneous segments. E.g., a satellite image analysis system which groups land areas into forest, urban...etc, for better utilization of natural resources.
- ❑ **Planning:** To generate an optimal sequence of actions to solve a particular problem. E.g., an Automated Vehicle which plans its path to avoid obstacles

Classification

Assign object/event to one of a given finite set of categories:

Medical diagnosis
Credit card applications or transactions
Fraud detection in e-commerce
Worm detection in network packets
Spam filtering in email
Recommended articles in a newspaper
Recommended books, movies, music, or jokes
Financial investments
Spoken words
Handwritten letters

6

Planning

Performing actions in an environment in order to achieve a goal:

- Playing checkers, chess, or backgammon
- Driving a car or a jeep
- Flying a plane, helicopter, or rocket
- Controlling an elevator
- Controlling a character in a video game
- Controlling a mobile robot

7

Models of Learning

- Classical AI deals mainly with deductive reasoning, learning represents inductive reasoning.
 - **Deductive reasoning** arrives at answers to queries relating to a particular situation starting from a set of general axioms,
 - **inductive reasoning** arrives at general axioms from a set of particular instances.
- Classical AI often suffers from the knowledge acquisition problem in real life applications where obtaining and updating the knowledge base is costly and prone to errors.
- Machine learning serves to solve the knowledge acquisition bottleneck by obtaining the result from data by induction.

Why Machine Learning

- Machine learning is important in several real life problem because of the following reasons:
 - Some tasks cannot be defined well except by example
 - Working environment may not be known at design time
 - Explicit knowledge encoding may be difficult and not available
 - Environments change over time
- Learning is widely used in a number of application areas such as:
 - Data mining and knowledge discovery
 - Speech/image/video (pattern) recognition
 - Adaptive control
 - Autonomous vehicles/robots
 - Decision support systems
 - Bioinformatics and WWW

Models of Learning

- Induction**, which is learning a generalization from a set of examples, is one of the most fundamental learning tasks.
- Different Approaches:
- Symbolic approach
 - Neural Nets
 - genetic and evolutionary learning.
- Strongest models of learning we have, may be seen in the human and animal systems that have evolved towards equilibration with the world. This approach to learning through adaptation is reflected in **genetic algorithms, genetic programming**
 - In the real world this information is often not immediately available AI needs to be able to **learn from experience**

Different kinds of learning

Supervised learning:

Someone gives us examples and the right answer for those examples

We have to predict the right answer for unseen examples

Unsupervised learning:

We see examples but get no feedback

We need to find patterns in the data

Reinforcement learning:

We take actions and get rewards

Have to learn how to get high rewards

Classification

To lend money to people. We have to predict whether they will pay you back or not. People have various (say, binary) features:

do we know their Address? do they have a Criminal record? high Income? Educated? Old? Unemployed?

We see examples: (Y = paid back, N = not)

+a, -c, +i, +e, +o, +u: Y

-a, +c, -i, +e, -o, -u: N

+a, -c, +i, -e, -o, -u: Y

-a, -c, +i, +e, -o, -u: Y

-a, +c, +i, -e, -o, -u: N

-a, -c, +i, -e, -o, +u: Y

+a, -c, -i, -e, +o, -u: N

+a, +c, +i, -e, +o, -u: N

Next person is +a, -c, +i, -e, +o, -u. Will we get paid back?

Classification...

We want some hypothesis h that predicts whether we will be paid back

+a, -c, +i, +e, +o, +u: Y
 -a, +c, -i, +e, -o, -u: N
 +a, -c, +i, -e, -o, -u: Y
 -a, -c, +i, +e, -o, -u: Y
 -a, +c, +i, -e, -o, -u: N
 -a, -c, +i, -e, -o, +u: Y
 +a, -c, -i, -e, +o, -u: N
 +a, +c, +i, -e, +o, -u: N

Lots of possible hypotheses: will be paid back if...

Income is high (*wrong on 2 occasions in training data*)

Income is high and no Criminal record (*always right in training data*)

(Address is known AND ((NOT Old) OR Unemployed)) OR ((NOT Address is known) AND (NOT Criminal Record)) (*always right in training data*)

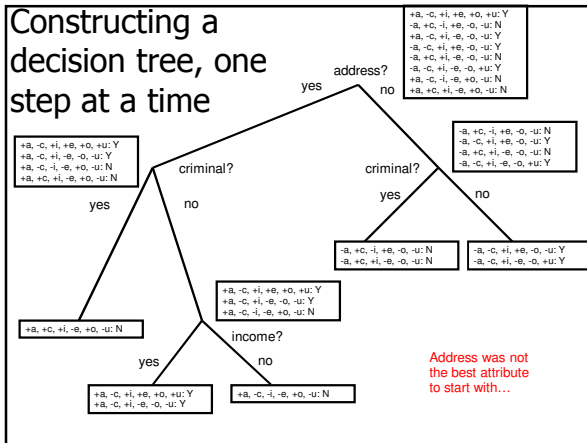
Which one seems best? Anything better?

Occam's Razor


(William of Occam in 1324)

- preferring simplicity and avoiding unnecessary assumptions.
- we should always accept the simplest answer that correctly fits our data.
- *simpler hypotheses tend to generalize to better data*
- Intuition: given limited training data, it is likely that there is some **complicated** hypothesis that is not actually good but that happens to perform well on the training data
- it is less likely that there is a **simple** hypothesis that is not actually good but that happens to perform well on the training data

Constructing a decision tree, one step at a time



Starting with a different attribute



Seems like a much better starting point than address
 Almost completely predicts whether we will be paid back

Decision Tree Learning

- ❑ Decision tree is a class of learning models that are more robust to noise as well as more powerful as compared to concept learning.
- ❑ Decision tree can be seen as rules for performing a categorisation
 E.g., "will we be paid back?"
- ❑ We're learning from examples
 Not turning thought processes into decision trees

Decision Tree Learning

- ❑ A decision-tree learning algorithm approximates a target concept using a tree representation, where each internal node corresponds to an attribute, and every terminal node corresponds to a class.
- ❑ Attributes describe examples (background knowledge)
 Each attribute takes only a finite set of values
- ❑ They classify instances or examples by starting at the root of the tree and moving through it until a leaf node.

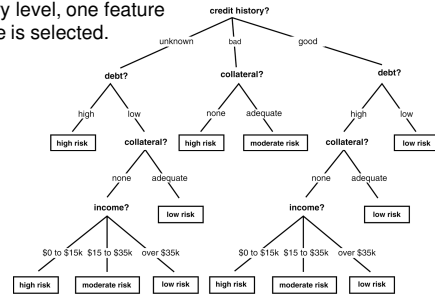
Decision Tree

Example (from the book): The problem of estimating credit risk by considering four features of a potential creditor. Such data can be derived from the history of credit applications.

NO.	RISK	CREDIT HISTORY	DEBT	COLLATERAL	INCOME
1.	high	bad	high	none	\$0 to \$15k
2.	high	unknown	high	none	\$15 to \$35k
3.	moderate	unknown	low	none	\$15 to \$35k
4.	high	unknown	low	none	\$0 to \$15k
5.	low	unknown	low	none	over \$35k
6.	low	unknown	low	adequate	over \$35k
7.	high	bad	low	none	\$0 to \$15k
8.	moderate	bad	low	adequate	over \$35k
9.	low	good	low	none	over \$35k
10.	low	good	high	adequate	over \$35k
11.	high	good	high	none	\$0 to \$15k
12.	moderate	good	high	none	\$15 to \$35k
13.	low	good	high	none	over \$35k
14.	high	bad	high	none	\$15 to \$35k


Decision Tree

At every level, one feature value is selected.



Decision Tree

Another possible decision tree:



The ID3 Algorithm

- The major question in decision tree learning Which nodes to put in which positions
- ID3 uses a measure called Information Gain Based on the notion of *entropy*
- Used to choose which node to put in next
- Node with the highest information gain is chosen
- When there are no choices, a leaf node is put on

information gain

measures how well a given attribute separates the training examples according to their target classification

This measure is used to select among the candidate attributes at each step while growing the tree