Chapter 11: Learning

- Introduction

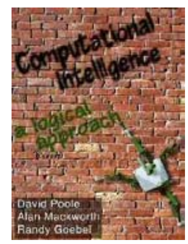

D. Poole, A. Mackworth, and R. Goebel, *Computational Intelligence: A Logical Approach*, Oxford University Press, January 1998

Learning

- Learning is the ability to improve one's behavior based on experience.
  - The range of behaviors is expanded: the agent can do more.
  - The accuracy on tasks is improved: the agent can do things better.
  - The speed is improved: the agent can do things faster.

Components of a learning problem

- The following components are part of any learning problem:
  - Task. The behavior or task that's being improved.
    For example: classification, acting in an environment
  - Data. The experiences that are being used to improve performance in the task.
  - Measure of improvement. How can the improvement be measured?
    For example: increasing accuracy in prediction, new skills that were not present initially, improved speed.

Learning Task

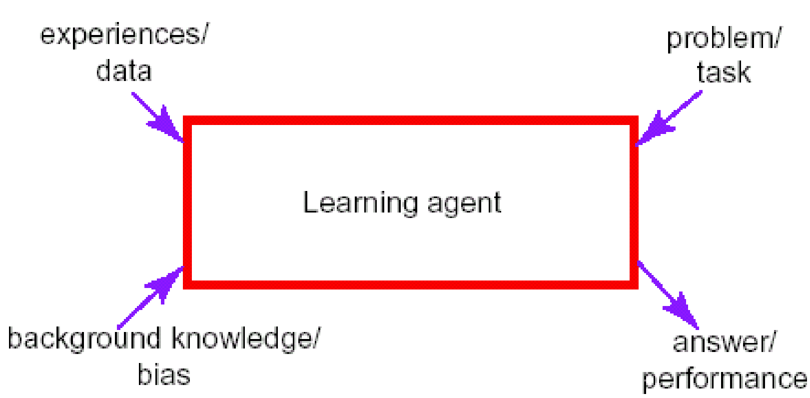

Learning Architecture

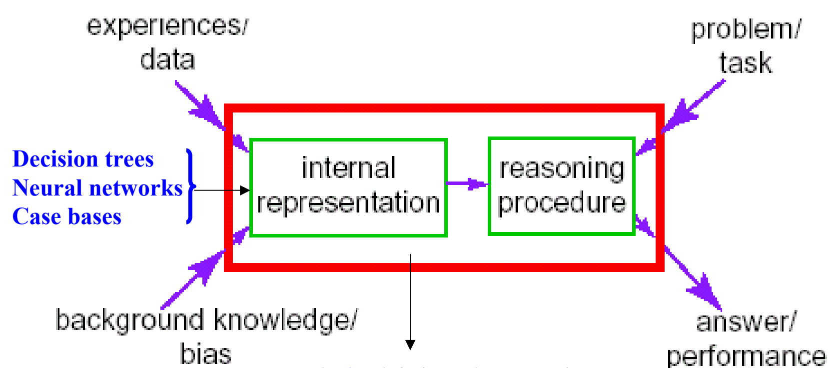

Induction: infer internal representation based on examples: classify unseen objects
Deduction: deriving consequences of a knowledge base
Abduction: hypothesizing what may be true about a particular case

Choosing a representation

- The richer the representation, the more useful it is for subsequent problem solving.
- The richer the representation, the more difficult it is to learn.

Common Learning Tasks

- **Supervised classification.** Given a set of pre-classified training examples, classify a new instance.
- **Unsupervised learning.** Find natural classes for examples.
- **Reinforcement learning.** Determine what to do based on rewards and punishments (feedback after actions: Credit Assignment Problem).

Example of Classification Data

Set of properties or features

Data that the infobot obtained from observing a user deciding whether to read/skip articles depending on whether the author was known or not, whether the article started a new thread or was a follow-up, the length of the article or whether it was read at home or at work.

examples	Action	Author	Thread	Length	Where
e1	skips	known	new	long	home
e2	reads	unknown	new	short	work
e3	skips	unknown	follow_up	long	work
e4	skips	known	follow_up	long	home
e5	reads	known	new	short	home
e6	skips	known	follow_up	long	work

- We want to classify new examples on property "*Action*" based on the examples' "*Author*", "*Thread*", "*Length*", and "*Where*".

Feedback

- Learning tasks can be characterized by the feedback given to the learner.
  - **Supervised learning.** What has to be learned is specified for each example.
  - **Unsupervised learning.** No classifications are given; the learner has to discover categories and regularities in the data.
  - **Reinforcement learning.** Feedback occurs after a sequence of actions.

Measuring Success

- The measure of success is not how well the agent performs on the training examples, but how well the agent performs for new examples (apparent versus true error).
- Consider two agents:
  - P claims the negative examples seen (during training) are the only negative examples. Every other instance is positive.
  - N claims the positive examples seen are the only positive examples. Every other instance is negative.
- Both agents correctly classify every training example, but disagree on every other new example: success in learning should not be judged on correctly classifying the training set but being able to classify unseen examples.

Bias

- The tendency to prefer one hypothesis over another is called a bias.
- Saying a hypothesis is better than N's or P's hypothesis isn't something that's obtained from the training data: but is external to the data.
- To have any inductive process make predictions on unseen data, you need a bias.
- What constitutes a good bias is an empirical question about which biases work best in practice.

Learning as Search

- Given a representation and a bias, the problem of learning can be reduced to one of search.
- Learning is search through the space of possible representations looking for the representation or representations that best fits the data, given the bias.
- These search spaces are typically prohibitively large for systematic search. Use of a hill climbing search.
- A learning algorithm is made of a search space, an evaluation function (heuristic), and a search method.

Noise

- Data aren't perfect:
  - some of the attributes (features) are assigned the wrong value
  - the attributes given are inadequate to predict the classification (poor features→ feature extraction)
  - there are examples with missing attributes
- Apparent error rates tend to be biased optimistically. The true error rate is almost invariably higher than the apparent error rate. This happens because of overfitting (overlearning) which occurs because of random correlations in the training set.