

# AITA : Overview and Revision

© John A. Bullinaria, 2004

1. Aims and Learning Outcomes
2. Artificial Intelligence – Goals, Roots, Sub-fields
3. Biological Intelligence and Neural Networks
4. Agents – Rational/Reflex/Intelligent Agents
5. Knowledge Representations – Semantic Nets and Frames
6. Production Systems – Recognize-Act Cycle
7. Uninformed Search
8. Expert Systems, Uncertainty Treatment, Machine Learning
9. Philosophical Issues and Misconceptions

## W2S1 : Module Aims and Learning Outcomes

### *Aims:*

1. Provide a general introduction to artificial intelligence and its techniques.
2. Give an overview of the key ideas such as knowledge representation, search, rule based systems, and learning that underlie the main sub-fields of artificial intelligence.
3. Demonstrate the need for different approaches for different problems.

### *Learning Outcomes:*

1. Structure the field of artificial intelligence into its main sub-fields, and outline the important features of AI systems.
2. Explain some of the most important knowledge representation formalisms and understand why there are different ones, discuss their advantages and drawbacks, and represent knowledge in unseen easy examples in any of them.
3. Apply simple uniformed search algorithms.
4. Understand the processes involved in Expert Systems and in building such systems.
5. Discuss the importance of learning for intelligent systems.
6. Provide examples of different types of AI systems, and explain their differences, common techniques, and limitations.

## W2S2 : What Exactly is AI?

“Artificial Intelligence (AI) is the part of computer science concerned with designing intelligent computer systems, that is, systems that exhibit characteristics we associate with intelligence in human behaviour – understanding language, learning, reasoning, solving problems, and so on.” (Barr & Feigenbaum, 1981)

**Engineering Goal** To solve real world problems using AI techniques such as knowledge representation, learning, rule systems, search, and so on.


**Scientific Goal** To determine which ideas about knowledge representation, learning, rule systems, search, and so on, explain various sorts of real intelligence.

AI has identifiable **roots** in a number of older fields: Philosophy, Logic/Computation, Psychology/Cognitive Science, Biology/Neuroscience, Evolution.

AI has many **sub-fields** (such as Neural Networks, Evolutionary Computation, Expert Systems, Natural Language Processing, Planning, Robotics, Vision), but they employ common **techniques** (such as Representation, Learning, Rules, Search).

## W3S1 : Biological Intelligence and Neural Networks

The human nervous system that forms the basis of our natural intelligence has the form:



We can attempt to use *Artificial Neural Networks* as the basis of AI systems, because:

1. They are extremely powerful computational devices (Turing equivalent).
2. Massive parallelism makes them very efficient.
3. They can learn and generalize – so no need for enormous feats of programming.
4. They are very fault tolerant – like the “graceful degradation” of biological systems.
5. They are very noise tolerant – coping where normal symbolic systems have difficulty.

They can be used for both *Brain Modelling* and *Artificial System Building*.

## W4S1 : Rational Agents

“An *agent* is anything that can be viewed as perceiving its environment through *sensors* and *acting* upon that environment through *effectors*.” (Russell & Norvig, 1995)



A *rational agent* is one that acts in a manner that causes it to be as successful as it can. We need to determine appropriate *performance measures* to judge success in each case.

An *ideal rational agent* is one that takes whatever action is expected to maximise its performance measure on the basis of the evidence provided by its perceptual history and whatever built-in knowledge it has. An *autonomous agent* will supplement its built-in knowledge with its own acquired (or learned) knowledge in order to act appropriately.

We design agents according to their *Percepts*, *Actions*, *Goals* and *Environment*.

# Types of Intelligent Agents

We can classify four types of intelligent agents of increasing sophistication:

1. **Simple Reflex Agents** – use simple *condition-action rules* or *IF–THEN rules* to produce appropriate (intelligent looking) actions for given percepts.
2. **Reflex Agents with an Internal State** – can keep track of previous states and use knowledge of how the world evolves to produce better actions.
3. **Goal based agents** – can determine sequences of actions to reach its goal state(s). *Search* and *planning* techniques will usually be required.
4. **Utility based agents** – can use utility/quality measures to choose between alternative sequences of actions/states that lead to a goal state being obtained.

Percepts should not only be used for generating agent actions, but also to improve the ability to act in the future, i.e. *learning*. Learning can range from trivial memorisation to the creation of entire scientific theories. A **learning agent** has four basic components: a Performance Element, a Learning Element, a Critic, and a Problem Generator.

## W5S1 : Knowledge Representation

The object of a *knowledge representation* is to express knowledge in a computer tractable form, so that it can be used to enable our agents to perform well.

A *knowledge representation language* is defined by two aspects:

1. *Syntax* The syntax of a language defines which configurations of components of the language constitute valid sentences.
2. *Semantics* The semantics defines which facts in the world the sentences refer to, and hence the statement about the world that each sentence can make.

A good knowledge representation system for a particular domain should possess four *important properties*: Representational Adequacy, Inferential Adequacy, Inferential Efficiency, and Acquisitional Efficiency.

Some common knowledge representation systems are: Natural Language, Databases, First Order Logic, Rule Systems, Semantic Networks, and Frames.

## Practical Aspects for Good Representations


In addition to the general requirements for knowledge representation, there are a number of practical aspects to consider when formulating good representations:

1. They must be *computable* – to be created with standard computing procedures.
2. They should make the important *objects* and *relations* explicit – so it is easy to see what is going on.
3. They need to *bring together* the objects and relations – so everything you need can be seen at once.
4. They should *suppress irrelevant detail* – so that rarely used details can be kept out of sight, but are still available when needed.
5. They should be *transparent* – so you can easily understand what is being said.
6. They need to be *concise* and *fast* – so information is stored and retrieved rapidly.
7. They should expose any natural *constraints* – so it is easy to express how one object or relation influences another.
8. They must be *complete* – so they can represent everything that needs representing.



## W6S1 : Semantic Networks

A *semantic network* represents knowledge as a set of labelled nodes and links, e.g.



Two important features are *inheritance* and *defaults* (i.e. typical values). We can assign expected/default values of parameters (e.g. height, has nose) and inherit them from higher up the hierarchy. This is more efficient than listing all the details at each level. We can also *over-ride* the defaults. For example, baseball players are taller than average men, so we make their default height over-ride the default height for men.

# Components of a Semantic Network

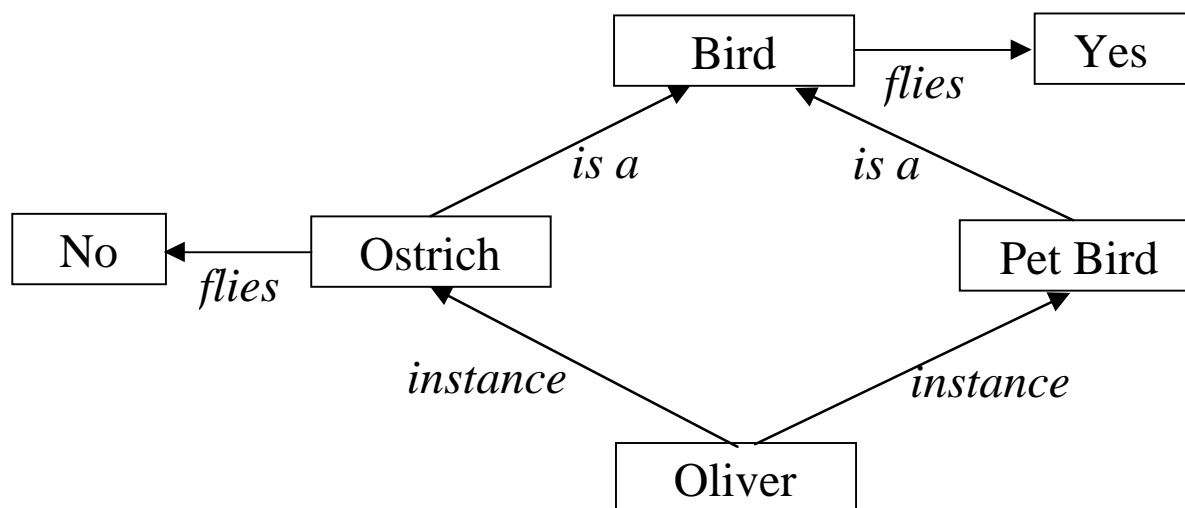
The formal components of semantic networks are quite straight-forward:

|                        |  |
|------------------------|--|
| <b>Lexical part</b>    | nodes – denoting objects<br>links – denoting relations between objects<br>labels – denoting particular objects and relations   |
| <b>Structural part</b> | the links and nodes form directed graphs<br>the labels are placed on the links and nodes   |
| <b>Semantic part</b>   | meanings are associated with the link and node labels<br>(the details will depend on the application domain)   |
| <b>Procedural part</b> | constructors allow creation of new links and nodes<br>destructors allow the deletion of links and nodes<br>writers allow the creation and alteration of labels<br>readers can extract answers to questions |

Some common examples include AND/OR Trees, IS-A and IS-PART Hierarchies, and Representations of Events and Natural Language Sentences.

## Multiple Inheritance and Tangled Hierarchies

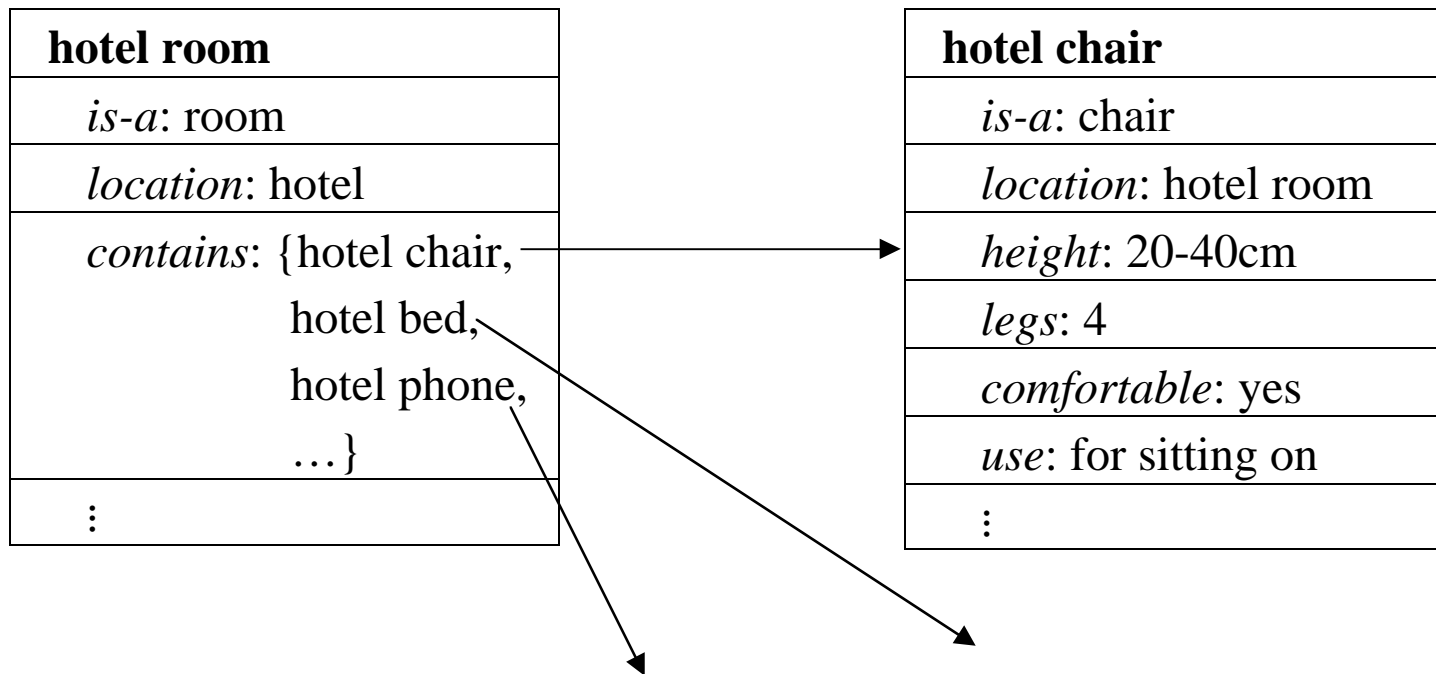
If *multiple inheritance* is allowed, we must avoid *inheritance conflicts*. Hierarchies that are not simple trees are called *tangled hierarchies*. For example, can Oliver fly in:



A better solution than having a specific “flies no” for all individual instances of an ostrich, would be to have an algorithm for traversing the hierarchy which guarantees that specific knowledge will always dominate over general knowledge. We can define the concept of *inferential distance*, that provides a *partial ordering* of closeness.

## W6S3 : Frame Based Systems

**Frames** are a natural extension of *Semantic Networks*. They consist of sets of slots filled by values, procedures for calculating values, or pointers to other frames. For example:



**Scripts** are frame based systems that describe stereotyped sequences of events and actions that enable an intelligent agent to perform appropriately in a particular context.

## W7S1 : Production Systems

A *production system* consists of four basic components:

1. A *set of rules* of the form  $C_i \rightarrow A_i$  where  $C_i$  is the condition part and  $A_i$  is the action part. The condition determines when a given rule is applied, and the action determines how it is applied.
2. One or more *knowledge databases* that contain whatever information is relevant for the given problem. Some parts of the database may be permanent, while others may temporary and only exist during the solution of the current problem. The information in the databases may be structured in any appropriate manner.
3. A *control strategy* that determines the order in which the rules are applied to the database, and provides a way of resolving any conflicts that can arise when several rules match at once.
4. A *rule applier* which is the computational system that implements the control strategy and applies the rules.

The whole system operates according to the deductive inference rule “modus ponens”.

## Recognize-Act Cycle

Typically, our production systems will have a rule interpreter that takes the form of a *Recognize-Act Cycle*. This cycle has four stages:

1. *Match* condition/premise patterns in the rules against the elements in the working memory to identify the set of applicable rules. This will usually involve *binding* specific values to the variables in the rules.
2. If there is more than one rule that can be ‘fired’ (i.e. that can be applied) at a given time, then use a *conflict resolution* strategy to choose which one from that *conflict set* to apply. If no rules are applicable, then stop.
3. *Apply* the chosen rule, which may involve adding a new item to the working memory or deleting an old one.
4. *Check* if the terminating condition is fulfilled. If it is, then stop. Otherwise, return to stage 1.

The *termination condition* can either be defined by a goal state, or by a cycle condition (e.g. a maximum number of cycles).

## Forward and Backward Chaining

*Forward chaining* or *data-driven inference* works by repeatedly: starting from the current state, matching at the premises of the rules (the IF parts), performing the corresponding actions (the THEN parts), and possibly updating the knowledge base or working memory.

*Backward chaining* or *goal-driven inference* works towards a final state by looking at the working memory to see if the sub-goal states already exist there. If not, then look at the actions (the THEN parts) of the rules that will establish the sub-goals, and set up new sub-goals for achieving the premises of those rules (the IF parts).

There are four major factors to help us choose between forward and backward reasoning:

1. Are there more possible start states or goal states? Start at the lowest.
2. Do we require the program to justify its reasoning? If so, follow human reasoning.
3. What kind of events trigger problem solving? New facts or queries?
4. In which direction is the branching factor greatest? Go with the lower.

## Conflict Resolution Strategies

Perhaps the five most common *general conflict resolution strategies* are:

1. Delete instantiations of rules that have already fired (to prevent repetitions).
2. Order instantiations by the generation age of all the elements. Prefer the younger ones, since new elements are more likely to describe the actual situation.
3. Compare the generation age of the elements in working memory which match the first condition of the rules. Prefer the younger ones. May be more efficient than 2.
4. Prefer the most specific rules (i.e. those with the most pre-conditions).
5. Random choice (which, if nothing else, is very easy to compute).

The easiest way to proceed in *problem specific cases* is to simply *add extra conditions* to the rules to avoid the conflicts. These extra conditions can be related to the inference strategies, e.g. to what is currently being searched for. However, we will end up with a mixture of heuristic and factual knowledge, and large knowledge bases will not be easily maintainable. It makes sense to separate the *object level knowledge* from the *meta-level knowledge*, and to supplement our rules with appropriate *meta-rules*.



## W8S1 : Uninformed Search

The *state space* is the space of all possible states, or configurations, our system may be in. Generally, we work with some convenient *representation* of that search space.

If the number of possible states of the system is small enough, we can represent them all, along with the transitions between them, in a *state space graph*.

The aim of our search algorithms is to find a route, or sequence of transitions, through the state space graph from our *initial state* to a *goal state*.

There are four important properties of search algorithms we need to consider:

1. *Completeness* – Is a solution guaranteed to be found if at least one solution exists?
2. *Optimality* – Is the solution found guaranteed to be best possible solution?
3. *Time Complexity* – The upper bound on the time to find a solution.
4. *Space Complexity* – The upper bound on the storage space (memory) required.

## Types of Uninformed Search

It is helpful to think of the search process as building up a *search tree* within the state space graph. The root of the search tree is the initial state. The leaf nodes correspond to states that have not yet been expanded, or have been expanded but generated no further nodes. A good strategy for avoiding *repeated states* will improve the search efficiency.

For tree branching factor  $b$ , maximum depth  $m$ , solution depth  $d$ , depth limit  $l$ , we find:

| Strategy | Complete      | Optimal | Time Complexity | Space Complexity |
|----------|---------------|---------|-----------------|------------------|
| BFS      | Yes           | Yes     | $O(b^d)$        | $O(b^d)$         |
| DFS      | No            | No      | $O(b^m)$        | $O(bm)$          |
| DLS      | If $l \geq d$ | No      | $O(b^l)$        | $O(bl)$          |
| DFIDS    | Yes           | Yes     | $O(b^d)$        | $O(bd)$          |
| BDS      | Yes           | Yes     | $O(b^{d/2})$    | $O(b^{d/2})$     |

The *best overall* is DFID which is complete, optimal and has low memory requirements.

## W9S1 : Expert Systems

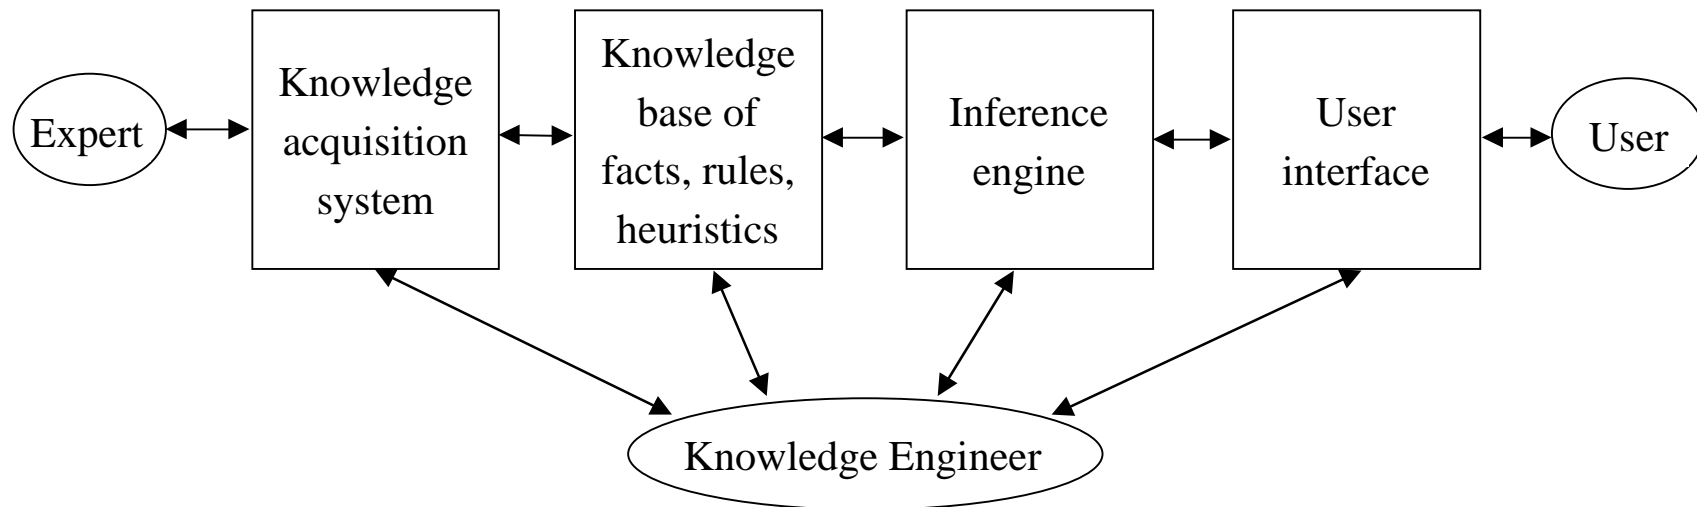
“An *expert system* is a computer program that represents and reasons with knowledge of some specialist subject with a view to solving problems or giving advice.” (Jackson, 1999)

There are no natural limits on what problem domains expert systems can be built to deal with. Expert systems can be distinguished from conventional computer systems in that:

1. They *simulate human reasoning* about the problem domain, rather than simulating the domain itself.
2. They perform *reasoning over representations of human knowledge*, in addition to doing numerical calculations or data retrieval. They have corresponding distinct modules referred to as the *inference engine* and the *knowledge base*.
3. Problems tend to be solved using *heuristics* (rules of thumb) or *approximate methods* or *probabilistic methods* which, unlike algorithmic solutions, are not guaranteed to result in a correct solution.
4. They usually have to provide *explanations* and *justifications* of their solutions or recommendations in order to convince the user that the reasoning is in fact correct.

# Building Expert Systems

Expert systems are typically built around a powerful production system:



The *knowledge acquisition* component allows the expert to enter knowledge or expertise into the system, and refine it later when required. The principal stages are: knowledge elicitation, intermediate representation, and compilation into executable form.

The major *technical problems* that need to be overcome are: the knowledge acquisition bottleneck, brittleness, lack of meta-knowledge, and the difficulty of reliable validation.

## W9S3 : Uncertainty Treatment

Potential *sources* of uncertainty for AI systems, such as expert systems, fall into two types: *imperfect domain knowledge* and *imperfect case data*.

Then there are three *types* of uncertainty: *randomness* (e.g. 3% of sensors fail in first year), *vagueness* (e.g. tall, little, etc.), and *inadequacy* (e.g. reliability of experts' rules).

The obvious way to treat uncertainty is to use the laws of probability (Bayes Rule, etc.) to perform *probabilistic reasoning*. However this will generally be intractable because of the number of joint/conditional probabilities that would be needed.

*Bayesian Networks* avoid this problem by using *conditional independencies* among the variables to ignore many interactions. If  $p(A|\{x_i\},\{y_i\}) = p(A|\{x_i\})$  then  $A$  is conditionally independent of  $\{y_i\}$  and, as far as  $A$  is concerned, if we know  $\{x_i\}$  we can ignore  $\{y_i\}$ .

Alternative uncertainty treatments include *MYCIN Style Certainty Factors* (but these are not necessarily self-consistent), *Dempster-Shafer Theory*, and *Fuzzy Logic*.

## W10S1 : Machine Learning

Strategies for learning can be classified according to the amount of inference the system has to perform on its training data. In increasing order we have

1. ***Rote learning*** – new knowledge is implanted directly with no inference at all.
2. ***Supervised learning*** – the system is given a set of training examples consisting of inputs and outputs and is required to discover the relation or mapping between them.
3. ***Unsupervised learning*** – the system is given a set of training examples consisting only of inputs and is required to discover for itself what the outputs should be.

When it comes to an expert system learning new rules there are two basic approaches:

1. ***Inductive rule learning*** methods create new rules about a domain not derivable from any previous rules. From ‘training data’ find rules that *generalize* to new situations.
2. ***Deductive rule learning*** enhances the efficiency of a system’s performance by deducing new rules from previously known domain rules and facts.

Useful machine learning procedures are: ***rule induction systems***, the ***version space*** approach to ***concept learning***, and ***decision tree algorithms*** such as ID3 and C4.5.

## W11S1 : Philosophical Issues and Misconceptions

We can distinguish *weak AI* (machines that act *as if* they are intelligent) and *strong AI* (machines which act intelligently have *real* conscious minds). We are interested in:

1. *Limitations* – What intelligent things are there that AI can never do?
2. *Misconceptions* – What can AI actually do, that some people think it can not do?

Four names that often crop up in “philosophical” debates are:

1. *Lucus and Penrose* – have tried to argue that *Gödel’s Incompleteness Theorem* (that any formal system rich enough to formulate arithmetic contains a true theorem that cannot be proved within that system) implies human reasoning can’t be mechanised.
2. *Drefus* – a philosopher who has written numerous books and papers about things he thinks computers and/or AI cannot do. Learning, uncertainty, common sense knowledge, control of sensors, and so on. He has been proven wrong on all counts.
3. *Searle* – his *Chinese Room* passes the Turing test in Chinese, without understanding Chinese or of how it passed the Turing test. Searle believes this to be of relevance to the possibility of strong AI. Compare this with the *Brain Prosthesis Experiment*.

## Overview and Reading

1. The module appears to have achieved its aims and learning outcomes.
2. We began with an overview of the roots, goals and sub-fields of AI, and a discussion of biological intelligence and its basis in neural network systems.
3. The general ideas of rational and intelligent agents were then covered.
4. This led to the need for a range of powerful knowledge representations such as semantic networks and frames, and production systems.
5. Procedures for performing uninformed search were then discussed.
6. The above ideas could then be used in building expert systems, for handling uncertainty in them, and in formulating machine learning systems.
7. We ended with a brief discussion of some relevant philosophical issues.

### Reading

Your lecture notes!