Overview of AI

In this chapter we look at what artificial intelligence is and what it tries to do. We have already seen some famous techniques such as search and expert systems, and some applications of AI, such as game playing. The subject has come a long way from the work of Turing, Shannon, McCarthy, Minsky, Newell and Simon, and others

1.1 What is AI?

Many different definitions have been proposed. Most talk about perception and reasoning:

- *The study of computations that make it possible to perceive, reason and act. (Winston)*
- *The study of how to make computers do things which currently people do better. (Rich & Knight)*
- *The study of mental faculties through the use of computational models. (Charniak & McDermott)*
- *AI is what AI people do.*

1.2 What Do We Study in AI?

We study some general methods and many applications. Some knowledge and methods are specific to a particular domain, but Computer Science tries to distill some general approaches. Examples of methods are:

- search
- logic & deduction
- statistical reasoning
- knowledge representation
- AI programming languages

Examples of applications are even more numerous:

- planning
- learning
- reasoning (common sense & specialized)

- vision
- natural language understanding
- expert systems
- game playing
- robotics
- mathematics
- scientific, medical & financial analysis

You should consider how one might write a program to do the following:

1. Solve IQ test questions e.g. those requiring geometric analogy
2. A geometry theorem prover: Starting with Euclid's axioms prove theorems of geometry
3. Play tictactoe; and play it perfectly
4. Automated animal classification
5. Reconstruction of simple 3D objects from 2-dimensional views
6. Understanding pictures: e.g. recognizing the two pieces of tiger as part of the same tiger
7. Planning in block-world
8. Calculus expert: automatic calculation of derivatives

1.3 Methods of AI

A general approach to an AI problem is the following:

1. define problem precisely

2. analyze

3. isolate and recognize task knowledge

4. choose appropriate problem-solving techniques

Three fundamental problem-solving techniques are:

- exhaustive search—generate each possibility in turn
- reduction/decomposition—break into pieces

- representations

Winston proclaims that "Good representations are the key to good problem solving". Good representations

- make the important objects and relations explicit
- expose natural constraints
- are transparent, complete and concise.

The game tree is a good example. The important ***objects*** are the possible positions of the game, and the ***relation*** is which positions come from which—clear from the descendant relationship. Another example is the graph for the fox-goose-wheat problem: these lie in the category of representations known as semantic nets. A good description is nice. Winston claims that "Once a problem is described using an appropriate representation, the problem is almost solved."

1.4 Measuring Success

The concept of knowledge representation and searching is fundamental to artificial intelligence. We can measure an AI solution by its effectiveness: a modified Turing test, for example.

But others argue that there are features of a good solution such as:

- generalization from existing knowledge
- clarity of knowledge
- extensibility of solution
- ability to use incomplete or inaccurate data.

And then there's still the question of intelligence, both human and artificial. Cognitive Science studies how humans reason and perceive and other aspects of intelligence. A lot of work in AI has been stimulated by theories of the human computational processes; a lot of work in cognitive science has been stimulated by the success or lack thereof of computational models.

In particular, there is a debate between the appropriateness of a purely symbolic model (as in, for example, a traditional computer) and a sub-symbolic model (as in, for example, an analog model of a neuron and neural networks).