

# Foundations of Artificial Intelligence

## Draft only

Phil Grant  
Department of Computer Science  
University of Wales Swansea

October 14, 1998

# Chapter 1

## Introduction

The aim of artificial intelligence (AI) is in some way to try to make computers perform tasks that humans tend to be good at. In particular problem solving. So in a sense it is an investigation into the attempt to make computers behave in a smarter manner. The actual name *Artificial Intelligence* was coined by John McCarthy in the 60's – he was the designer of the language LISP.

There is also often a two way interaction - by observing how humans perform a particular task it may be possible to imitate this by means of a program - from the program (especially a learning program) it may be possible to improve the human performance. A good illustration of this is the Dendral system for recognising chemical structures from mass spectrograms — rules were derived which were not known by the experts. It should be pointed out that Dendral is a very successful program and is in fact the world expert on such analysis. Now what does it mean for a machine to be intelligent? Turing thought about this problem in the 50's and came up with the following test.

### 1.1 The Turing Test for Intelligence

A human interrogator can ask questions, via say a vdu, of a machine in one room and a human in another. If she is unable to decide from the answers received which is connected to the machine then Turing would say that it should be deemed intelligent. A very early attempt to try to build a program which would pass the Turing test was the Eliza program of Weizenbaum — MIT 1966. This is supposed to mimic a conversation with a psychologist. When one first sees it, it does seem quite effective but once one has played with it for a little time it is quite easy to trick it and in fact infer how it works. We shall see an implementation of Eliza later.

Clearly in order to pass the Turing test the system must understand natural language. This is an extremely difficult task to say the least! For humans make all sorts of assumptions and inferences in order to make sense of sentences and utterances. It is generally believed that to really handle natural language in a manner comparable to humans then a system must incorporate common sense reasoning and have access to the kind of general knowledge

a human would expect to know.

One attempt to provide this background knowledge is the **CYC** project of Lenat's. He is building an enormous system at MCC Texas which is trying to capture the full breadth of human knowledge. A company, *CYCORG* has now been set up to market the technology. The team examines text from encyclopedias, newspapers etc. and codes this information (making use of what the writer assumed the reader knew of the world) into the database. There are currently several million entries in the knowledge base.

## 1.2 Main Areas of AI

Here are a number of the topics which might be considered as part of the general subject of Artificial Intelligence.

- Pattern Recognition and Image Analysis
- Natural Language Understanding
- Representation of Knowledge
- Theorem Proving and Problem Solving
- Game Playing - Chess, Draughts etc.
- Automatic Programming - Program Synthesis
- Program Transformation
- Expert Systems
- Conjecture and Hypothesis Formation
- Robotics
- Vision
- Neural Networks
- Machine Learning
- Genetic Programming & Algorithms
- Fuzzy Logic

We shall only consider a few of these topics. In particular we shall not be dealing with the important area of expert systems or neural networks as these will be dealt with in other courses — some in the follow on course *Applications of AI*.

To put AI into perspective, it is instructive to see how two well known researchers in this field, Lenat and Feigenbaum, view the ultimate goals of AI. In a recent article they have listed nine goals.

### 1.2.1 Nine Ultimate Goals of AI (Lenat & Feigenbaum)

**Understand human cognition** How do humans solve problems? Try to obtain deep knowledge of human memory, problem-solving abilities, learning, decision making etc.

**Cost-effective automation** Replace humans in intelligent tasks. Have programs which perform as well as humans currently doing the job.

**Cost-effective intelligent amplification** Build systems to help humans think better, faster, deeper . . . . For example system to help GP diagnose disease.

**Superhuman intelligence** Build programs to exceed human intelligence.

**General problem solving** Solve broad range of problems. Systems with breadth of mind.

**Coherent discourse** Communicate with people using natural language. Carry out an intelligent dialogue. (c.f. Turing test).

**Autonomy** Have intelligent systems acting on own initiative. Must react to the real world.

**Learning (induction)** The system should be able to gather own data and how. Generalise, hypothesise, apply/learn heuristics; reason by analogy.

**Information** Store information and know how to retrieve it.

Having presented this list they also offer some strategies for achieving these goals.

## 1.3 References

In this course we will concentrate on general approaches to AI. Most of the examples to illustrate techniques and methods will use Prolog as the programming language. Both *Sterling & Shapiro* and *Bratko* are good references.

The paper of Lenat and Feigenbaum is *On the Threshold of Knowledge* in special issue of the AI journal *Foundations of AI* edited by D Kirsh, MIT Press, 1992.