COMP210: Artificial Intelligence
Lecture 1. Introduction

Trevor Bench-Capon
Room 215, Ashton Building

http://www.csc.liv.ac.uk/~tbc/COMP210

Course Outline

- The course consists of:
  - 30 lectures slots (may use some for tutorials or class tests);
  - tutorial exercises;
  - lab exercises;
    - Not assessed
  - enough self study to understand the material;
- Assessment
  - two class tests;
  - a two hour exam.
- Course materials, syllabus, the course guide, lecture slides, tutorial and lab exercises etc can be obtained from http://www.csc.liv.ac.uk/~tbc/COMP210

Timetable

- Tuesday, 13:00, Maths 209
- Thursday, 12:00, Lecture Theatre C, University Lecture Rooms Block
- Friday, 10:00, Lecture Theatre D, University Lecture Rooms Block
  - PROLOG on Friday for first 5 weeks
  - Non-assessed lab exercises
  - *Open book class test*

Reading

- Good AI books include:-
  - S. Russell and P. Norvig. *AI A Modern Approach*. Second Edition Prentice Hall, 2003
  - E. Rich and K. Knight. *Artificial Intelligence*, McGraw-Hill, 1991 (2nd edition)
- The following is a (cheap) text (not as good as the above) which covers standard material.
  - A. Cawsey. *The Essence of Artificial Intelligence.* Prentice-Hall, 1998.

Out of date and out of print, but perhaps interesting:
T. Bench-Capon: *Knowledge representation: an approach to artificial intelligence,* Academic Press, 1991
Still maybe copies in the Harold Cohen

Prolog Books

- Best book
  - I. Bratko. *Prolog Programming for Artificial Intelligence*. Addison Wesley 1990.
- Straightforward book:
  - Danny Crookes *Introduction to programming in Prolog* Prentice-Hall,1988

Course Contents

- Introduction to Artificial Intelligence
- Prolog - an AI programming language
- Search
- Knowledge Representation
- Propositional Logic
- First-Order Logic
- Resolution Based Proof for Propositional and First-Order Logics
- Expert Systems
- AI Applications

Aims

- To introduce students to ***knowledge representatio***n, common ***knowledge representation paradigms*** and the issues involved in knowledge representation.
- To introduce students to the sorts of ***systems*** that can be built using artificial intelligence techniques, in particular ***knowledge based systems***.
- To give students an awareness of the issues involved in building such systems.
- To provide a grounding in ***Prolog***.

Learning Outcomes

- An awareness of the principles of knowledge representation.
- An understanding of search techniques and logic, particularly as related to knowledge representation.
- An understanding of the major knowledge representation paradigms: production rules, propositional and first order predicate calculus and structured objects.
- An understanding of how these representations can be manipulated to solve problems in a knowledge based systems context.
- An appreciation of knowledge based systems.
- Awareness of other applications of AI.
- Familiarity with the essentials of Prolog so as to enable exploration of the above in practice.

What I expect from you.

- To attend lectures.
- To be punctual.
- To turn mobile phones off and not to chat in lectures.
- To do whatever reading and self study is required to understand the material.
- To attempt the tutorial and laboratory exercises.
- To sit the class tests – it is ***essential*** that you attend these
- To pass the exam

Credits

- This set of slides is based on the materials provided by people who have previously taught this course in the University of Liverpool
  - Boris Konev
  - Clare Dixon
  - Simon Parsons
  - Trevor Bench-Capon
- AiMA supplement slides

What is AI?

- AI attempts to build intelligent entities
  - Intelligence - must relate to tasks involving "higher mental processes": so ***not*** simple response to some stimulus (e.g. a thermostat)
  - .***not*** mechanical performance of some algorithm
- "Hard" AI view
  - "The ultimate goal of AI research ... is to build a person, or more humbly an animal"- *Charniak and McDermott*
- "Soft" AI view
  - "AI is the study of how to make computers do things, at which, at the moment, people are better" - *Elaine Rich*

AI is both science and engineering:

- the *science* of understanding intelligent entities — of developing theories which attempt to explain and predict the nature of such entities;
- the *engineering* of intelligent entities

Four Views of AI

- **Systems that think like humans**
  - cognitive science, expert systems
- **Systems that act like humans**
  - The Turing Test, chess programs
- **Systems that think rationally**
  - Approaches based on logic and mathematics
- **Systems that act rationally**
  - Contemporary agent based approaches

Acting Humanly

- Emphasis on how to *tell* if a machine is intelligent*,* not on how to *make* it intelligent
- when can we count a machine as being intelligent?
- "Can machines think?" ="Can machines behave intelligently?"

Turing Test

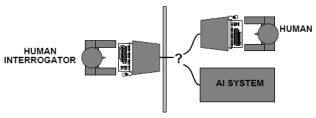

System passes if the questioner cannot tell the difference

No program has yet passed the test: most successful ones rely on ***tricks*** rather than intelligence

But can obtain human level performance (or better) in some ***specific*** areas like chess

Thinking Humanly

- Try to understand how the mind works — how do we think?
- Two possible routes to find answers:
  - by *introspection* — we figure it out ourselves!
  - by *experiment* — draw upon techniques of psychology (or neuroscience) to conduct controlled experiments.
- The discipline of *cognitive science:* at one time influential in vision, natural language processing, and learning.

Human vs Machine Thinking

- Expert systems — "AI success story in early 80's"
  - Human expert's knowledge and experience is represented as a computer program
  - Rule-based representation of knowledge
  - Typical domains are:
    - medicine (INTERNIST, MYCIN, . . .)
    - geology (PROSPECTOR)
    - chemical analysis (DENDRAL)
    - configuration of computers (R1)
    - Law (British Nationality Act)
- Thinking humanly can work (sometimes)

Human vs Machine Thinking

- Computer program playing chess
  - "Human way"
    - Tried by World champion M.Botvinnik (who also was a programmer)
    - Poor performance
  - "Computer way"
    - Sophisticated search algorithms
    - Vast databases
    - Immense computing power
    - Human world champion beaten
- True also for vision, robotics, speech recognition etc

Thinking Rationally

- Trying to understand how we actually think is one route to AI — but how about how we should think.
- Use logic to capture the laws of rational thought as symbols.
- Reasoning involves shifting symbols according to well-defined rules (like algebra).
- Result is idealised reasoning.

Logic and AI

- Logical approach theoretically attractive.
- Lots of problems:
  - transduction — how to map the environment to symbolic representation;
  - representation — how to represent real world phenomena (time, space, . . .) symbolically;
  - reasoning — how to do symbolic manipulation tractably

Acting Rationally

- Acting rationally = acting to achieve one's goals, given one's beliefs.
- An agent is a system that perceives and acts; an intelligent agent is one that acts rationally w.r.t. the goals we delegate to it.
- Emphasis shifts from designing theoretically best decision making procedure to the best decision making procedure possible in circumstances.
- Logic may be used in the service of finding the best action — not as an end in itself.

Acting Rationally

- Achieving *perfect rationality* — making the best decision theoretically possible — is not usually achievable, because of
  - limited resources:
  - limited time;
  - limited computational power;
  - limited memory;
  - limited or uncertain information about environment.
- The aim is to *do the best with what you've got*

- Today
  - General Course Information
  - Overview of What AI is
- Next time
  - overview to some common AI techniques we will study during the course
  - Typical AI applications
  - A brief history of AI