CS 343: Artificial Intelligence

Introduction

Definition of AI

- "The art of creating machines that perform functions that require intelligence when performed by people" (Kurzweil, 1990).
- "The branch of computer science that is concerned with the automation of intelligent behavior." (Luger and Stublefield, 1993)

Systems that think like humans.	Systems that think rationally.
Systems that act like humans.	Systems that act rationally.

Acting Humanly: The Turing Test

- If the response of a computer to an unrestricted textual natural-language conversation cannot be distinguished from that of a human being then it can be said to be intelligent.

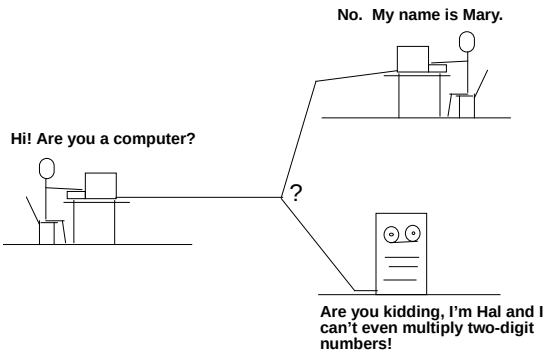

- Loebner Prize: Current contest for restricted form of the Turing test.

Thinking Humanly: Cognitive Modelling

- Method must not just exhibit behavior sufficient to fool a human judge but must do it in a way demonstrably analogous to human cognition.
- Requires detailed matching of computer behavior and timing to detailed measurements of human subjects gathered in psychological experiments.
- **Cognitive Science**: Interdisiplinary field (AI, psychology, linguistics, philosophy, anthropology) that tries to form computational theories of human cognition.

Thinking Rationally: Laws of Thought

- Formalize “correct” reasoning using a mathematical model (e.g. of deductive reasoning).
- **Logicist Program**: Encode knowledge in formal logical statements and use mathematical deduction to perform reasoning:

  Problems:

  - Formalizing common sense knowledge is difficult.
  - General deductive inference is computationally intractable.

Acting Rationally: Rational Agents

- An agent is an entity that perceives its environment and is able to execute actions to change it.
- Agents have inherent goals that they want to achieve (e.g. survive, reproduce).
- A rational agent acts in a way to maximize the achievement of its goals.
- True maximization of goals requires omniscience and unlimited computational abilities.
- Limited rationality involves maximizing goals within the computational and other resources available.

Foundations of AI

- Many older disciplines contribute to a foundation for artificial intelligence:
  - Philosophy: logic, philosophy of mind, philosophy of science, philosophy of mathematics
  - Mathematics: logic, probability theory, theory of computability
  - Psychology: behaviorism, cognitive psychology
  - Computer Science & Engineering: hardware, algorithms, computational complexity theory
  - Linguistics: theory of grammar, syntax, semantics

Birth

- McCullouch and Pitts (1943) theory of neurons as logical computing circuits.
- Work in early 50’s by Claude Shannon and Turing on game playing and Marvin Minsky on neural networks.
- Dartmouth conference (1956)
  - Organized by John McCarthy attended by Marvin Minsky, Allen Newell, Herb Simon, and a few others.
  - Coined term “artificial intelligence.”
  - Presentation of game playing programs and Logic Theorist.

Early Years

- Development of General Problem Solver by Newell and Simon in early sixties.
- Arthur Samuel's late fifties work on learning to play checkers.
- Frank Rosenblatt's Perceptron (1962) for training simple neural networks
- Work in the sixties at MIT lead by Marvin Minsky and John McCarthy
  - Development of LISP symbolic programming language
  - SAINT: Solved freshman calculus problems
  - ANALOGY: Solved IQ test analogy problems
  - SIR: Answered simple questions in English
  - STUDENT: Solved algebra story problems
  - SHRDLU: Obeyed simple English commands in the blocks world

Early Limitations

- Hard to scale solutions to toy problems to more realistic ones due to difficulty of formalizing knowledge and combinatorial explosion of search space of potential solutions.
- Limitations of Perceptron demonstrated by Minsky and Papert (1969).

Knowledge is Power: Expert Systems

- Discovery that detailed knowledge of the specific domain can help control search and lead to expert level performance for restricted tasks.
- First expert system DENDRAL for interpreting mass spectrogram data to determine molecular structure by Buchanan, Feigenbaum, and Lederberg (1969).
- Early expert systems developed for other tasks:
  - MYCIN: diagnosis of bacterial infection (1975)
  - PROSPECTOR: Found molybendum deposit based on geological data (1979)
  - R1: Configure computers for DEC (1982)

AI Industry

- Development of numerous expert systems in early eighties.
- Estimated $2 billion industry by 1988.
- Japanese start "Fifth Generation" project in 1981 to build intelligent computers based on Prolog logic programming.
- MCC established in Austin in 1984 to counter Japanese project.
- Limitations become apparent, prediction of AI Winter
  - Brittleness and domain specificity
  - Knowledge acquisition bottleneck

Rebirth of Neural Networks

- New algorithms discovered for training more complex neural networks (1986).
- Cognitive modelling of many psychological processes using neural networks, e.g. learning language.
- Industrial applications:
  - Character and hand-writing recognition
  - Speech recognition
  - Processing credit card applications
  - Financial prediction
  - Chemical process control

Recent Times

- General focus on learning and training methods to address knowledge-acquisition bottleneck.
- Shift of focus from rule-based and logical methods to probabilistic and statistical methods (e.g. Bayes nets, Hidden Markov Models).
- Increased interest in particular tasks and applications
  - Data mining
  - Intelligent agents and Internet applications (softbots, believable agents, intelligent information access)
  - Scheduling/configuration applications (Successful companies: I2, Red Pepper, Trilogy)