

Artificial Intelligence

Chapter 1: Introduction

Edited by Neng-Fa Zhou

Thinking Humanly

"The exciting new effort to make computers think ... *machines with minds*, in the full and literal sense." (Haugeland, 1985)

"[The automation of] activities that we associate with human thinking, activities such as decision-making, problem solving, learning .. ." (Hellman, 1978)

Thinking Rationally

"The study of mental faculties through the use of computational models."
(Charniak and McDermott, 1985)

"The study of the computations that make it possible to perceive, reason, and **act**."
(Winston, 1992)

Acting Humanly

"The art of creating machines that perform functions that require intelligence when performed by people." (Kurzweil, 1990)

"The study of how to make computers do things at which, at the moment, people are better." (Rich and Knight, 1991)

Acting Rationally

"Computational Intelligence is the study of the design of intelligent agents." (Poole *et al*, 1998)

"AI ... is concerned with intelligent behavior in artifacts." (Nilsson, 1998)

- **The Turing Test**

- A computer passes the test if a human interrogator, after posing some written questions, cannot tell whether the written responses come from a person or from a computer.

- **natural language processing**
 - **knowledge representation**
 - **automated reasoning**
 - **machine learning**

- **Total Turing Test**

- **computer vision**
 - **robotics**

- We need to get *inside* the actual workings of human minds
 - through introspection
 - through psychological experiments
 - through brain imaging
- **Cognitive Science**
 - Interdisciplinary 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)
- 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

- A rational agent is one that acts so as to achieve the best outcome or, when there is uncertainty, the best expected outcome.
 - 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.
 - Limited rationality involves maximizing goals within the computational and other resources available.

- **Philosophy**

- Can formal rules be used to draw valid conclusions?
- How does the mind arise from a physical brain?
- Where does knowledge come from?
- How does knowledge lead to action?

- **Mathematics**

- What are the formal rules to draw valid conclusions?
- What can be computed?
- How do we reason with uncertain information?

- **Economics**
 - How should we make decisions so as to maximize payoff?
 - How should we do this when others may not go along?
 - How should we do this when the payoff may be far in the future?
- **Neuroscience**
 - How do brains process information?
- **Psychology**
 - How do humans and animals think and act?

- **Computer engineering**
 - How can we build an efficient computer?
- **Control theory and cybernetics**
 - How can artifacts operate under their own control?
- **Linguistics**
 - How does language relate to thought?

The birth of artificial intelligence

- 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.

- 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
- Development of LISP symbolic programming language
- J.A. Robinson's resolution method (1965)

- Work in the sixties at MIT lead by Marvin Minsky and John McCarthy
 - 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

- Hard to scale solutions from toy problems to more realistic ones
 - Most early programs knew nothing of their subject matter. They succeeded by means of simple syntactic manipulations.
 - Many of the problems that AI was attempting to solve are intractable.
 - 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 molybdenum deposit based on geological data (1979)
 - R1: Configure computers for DEC (1982)

- Development of numerous expert systems in early eighties.
- 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

The return of neural networks

- New algorithms, such as back-propagation, discovered for training more complex neural networks (1986).
- Connectionist models vs. symbolic models
- Neural network research
 - create effective network architectures and algorithms and understanding their mathematical properties
 - careful modeling of the empirical properties of actual neurons and ensembles of neurons
- Industrial applications (e.g., OCR)

- 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, Data mining).
- Increased interest in particular tasks and applications
 - Intelligent agents and Internet applications
 - Self-driving cars
 - recognition, machine translation, language understanding
 - AI planning