
Introduction

CS 461 – Artificial Intelligence

Pinar Duygulu

Bilkent University, Spring 2008

Slides are mostly adapted from AIMA

AI is exiting

- For thousands of years it is tried to be understood how humans think
- The field of Artificial Intelligence goes further and attempts not just to understand but also to build intelligent systems
- It is one of the newest science – started after World War 2 and the name given in 1956
- It is one of the most preferred science – together with molecular biology
- Still has openings for several full time Einsteins
- Huge variety of subfields – ranging from general purpose areas such as learning and perception to such specific tasks as playing chess, proving mathematical theorems, writing poetry, and diagnosing diseases
- AI systematizes and automates intellectual tasks and therefore it is relevant to any sphere of intellectual activity

What is AI?


<p>Thinking humanly</p> <p>The exciting new effort to make computers think...machines with minds, in the full and literal sense (Haugeland, 1985)</p> <p>[The automation of activities that we associate with human thinking, activities such as decision making, problem solving, learning... (Bellman, 1978)</p>	<p>Thinking rationally</p> <p>The study of mental faculties through the use of computational models (Charniak and McDermott, 1985)</p> <p>The study of the computations that make it possible to perceive, reason and act (Winston, 1992)</p>
<p>Acting humanly</p> <p>The art of creating machines that perform functions that require intelligence when performed by people (Kurzweil, 1990)</p> <p>The study of how to make computers do things at which, at the moment, people are better (Rich and Knight, 1991)</p>	<p>Acting rationally</p> <p>Computational Intelligence is the study of the design of intelligent agents (Poole et. al , 1998)</p> <p>AI ... is concerned with intelligent behaviour in artifacts (Nilson, 1998)</p>

* A system is rational if it does the right thing given what it knows

- Human-centered approaches use empirical science, involving hypothesis and experimental confirmation,
- Rationalist approaches involves a combination of mathematics and engineering

Acting humanly: Turing Test

- Proposed by Alan Turing (1950)
- Operational definition for intelligent behaviour: the Imitation Game



- The computer passes the test if a human interrogator, after posing some written questions, cannot tell whether the written responses come from a person or not
- Suggested major components of AI: natural language processing, knowledge representation, automated reasoning, machine learning
- Total Turing test also requires computer vision and robotics

- “Artificial flight” is succeeded by not imitating the birds, but by learning aerodynamics. The goal is not to fool pigeons.

Thinking humanly: cognitive modeling

- In order to say that a given program thinks like a human, we must have some way of determining how humans think
- Requires scientific theories of internal activities of the brain
- How to validate? requires
 - 1) Predicting and testing behavior of human subjects (top-down)
 - 2) Direct identification from neurological data (bottom-up)
- Both approaches (roughly, Cognitive Science and Cognitive Neuroscience) are now distinct from AI

Thinking rationally: "laws of thought"

- Aristotle: what are correct arguments/thought processes?
- “Socrates is a man, all men are mortal, therefore Socrates is mortal”
- Several Greek schools developed various forms of *logic*: *notation* and *rules of derivation* for thoughts;
- Direct line through mathematics and philosophy to modern AI

- Problems:
 1. Not all intelligent behavior is mediated by logical deliberation
 2. What is the purpose of thinking? What thoughts should I have?

Acting rationally: rational agent

- Rational behavior: doing the right thing
- The right thing: that which is expected to maximize goal achievement, given the available information
- Doesn't necessarily involve thinking – e.g., blinking reflex – but thinking should be in the service of rational action

Rational agents

- An agent is an entity that perceives and acts
- This course is about designing rational agents
- Abstractly, an agent is a function from percept histories to actions:

$$[f: P^* \rightarrow A]$$

- For any given class of environments and tasks, we seek the agent (or class of agents) with the best performance
- Caveat: computational limitations make perfect rationality unachievable
 - design best program for given machine resources

AI prehistory

- Philosophy Logic, methods of reasoning, mind as physical system foundations of learning, language, rationality
- Mathematics Formal representation and proof algorithms, computation, (un)decidability, (in)tractability, probability
- Economics utility, decision theory
- Neuroscience physical substrate for mental activity
- Psychology phenomena of perception and motor control, experimental techniques
- Computer engineering building fast computers
- Control theory design systems that maximize an objective function over time
- Linguistics knowledge representation, grammar

The gestation of AI (1943-1955)

- The first work that is now generally recognized as AI was done by Warren McCulloch and Walter Pitts (1943)
- They proposed a model of artificial neurons in which each neuron is characterized as being “on” or ”off”, with a switch to “on” occurring in response to stimulation by a sufficient number of neighboring neurons
- They showed that any computable function could be computed by some network of connected neurons, and all the logical connectives (and, or, not, etc) could be implemented by simple net structures
- Later (1949), Donald Hebb demonstrated a simple updating rule for modifying the connection strengths between neurons (Hebbian rule)
- In 1950 Marvin Minsky and Dean Edmond (two undergrads at Harvard) built the first neural network computer (SNARC)
- Then Minsky studied universal computation in neural networks during his PhD at Princeton
- Later, Minsky proved influential theorems showing the limitations of NN research

Alan Turing (1950)

- It was Alan Turing who first articulated a complete vision of AI in his 1950 article
- “Computing Machinery and Intelligence”
- He introduced the Turing test, machine learning, genetic algorithms, and reinforcement learning

The birth of AI (1956)

- In 1956 John McCarthy organized a two month workshop at Dartmouth College bringing together researchers interested in automata theory, neural nets, and the study of intelligence
- This workshop did not lead to any new breakthroughs but introduced the major figures to each other. For the next 20 years the field is dominated by these people and their students and colleagues at MIT, CMU, Stanford and IBM
- The name Artificial Intelligence is also given by McCarthy during this workshop (computational rationality would be the other alternative)
- The proposal of the workshop also explains why AI becomes a separate field rather than a subfield of operational research, control theory or mathematics:
 - AI duplicates human faculties like creativity, self-improvement, and language use
 - Also the methodology is different. AI is clearly a branch of computer science and AI is the only one filed to attempt to build machines that will function autonomously in complex, changing environments

Early enthusiasm, great expectations (1952-1969)

- The early years of AI were full of successes – in a limited way
- Given the primitive computers and programming tools of the time, and the fact that only a few years earlier computers were seen as things that could do arithmetic and no more, it was astonishing whether a computer did anything remotely clever
- The intellectual establishment, by and large, preferred to believe that “a machine can never do X”
- AI researchers naturally responded by demonstrating one X after another
- John McCarthy referred to this period as the “Look, Ma, no hands!” era

Early enthusiasm, great expectations (1952-1969)

- Allen Newell and Herbert Simon created a reasoning program – Logic Theorist (LT)
- The program was able to prove most of the theorems in the book Principia Mathematica by Russell and Whitehead and the one of the proofs was even shorter than the one in Principia
- This success is followed by General Problem Solver (GPS) which was designed from the start to imitate human problem-solving protocols.
- Within the limited class of puzzles it could handle it turned out that the order in which the program considered subgoals and possible actions was similar to that in which humans approached the same problems
- Based on that Newell and Simon formulated the famous physical symbol system hypothesis – “ a physical symbol system has the necessary and sufficient means for general intelligent action” – any system (human or machine) exhibiting intelligence must operate by manipulating data structures composed of symbols
- In 1959 at IBM, Geometry Theorem Prover is constructed – prove theorems that many students of mathematics would find tricky
- Arthur Samuel wrote a series of programs for checker (draughts) that eventually learned to play at a strong amateur level. His program quickly learned to play a better game than his creator

Early enthusiasm, great expectations (1952-1969)

- John McCarthy made three crucial contributions in one year, 1958, at MIT
 - Defined LISP which became the dominant AI programming language
 - Time sharing is invented
 - In “Programs with Common Sense” he described Advice Taker, a hypothetical program that can be seen as the first complete AI system
 - For example, he showed how some simple axioms would enable the program to generate a plan to drive to the airport to catch a plane
 - It embodied central principles of knowledge representation and reasoning. That is useful to have a formal, explicit representation of the world and of the way an agent’s actions affect the world and to be able to manipulate these representations with deductive processes
 - most of that work remains relevant today

Early enthusiasm, great expectations (1952-1969)

- While McCarthy stressed representation and reasoning in formal logic, Minsky was more interested in getting programs to work and eventually developed an anti-logical outlook
- His group chose limited problems that appeared to require intelligence to solve.
- These limited domains are known as microworlds (e.g. closed form calculus integration problems, geometric analogy problems that appear in IQ tests)
- The most famous microworld was the blocks world, which consists of a set of solid blocks placed on a table. A typical task is to rearrange the blocks in a certain way, using a robot hand that can pick one block at a time

A dose of reality (1966-1973)

- Hebert Simon – 1957
- “It is not my aim to surprise or shock you – but the simplest way I can summarize is to say that are now in the world machines that think, that learn and that create. Moreover, their ability to do these things is going to increase rapidly until – in a visible future – the range of problems they can handle will be coextensive with the range to which the human mind has been applied”

A dose of reality (1966-1973)

- Simons's over-confidence was due to the promising performance of the early AI systems on simple examples. In almost all cases, however, the early systems turned out to fail miserably when tried out on wider selections of problems and on more difficult problems
- Problems:
- Most of early programs contained little or no knowledge of their subject matter, they succeeded by means of simple syntactic manipulations (e.g. translation)
- Intractability of problems that AI was attempting to solve (genetic algorithms)
- Fundamental limitations on the basic structures being used to generate intelligent behavior (perceptrons – although they learn anything that they could represent, they could represent very little)

Knowledge based systems: The key to power (1969-1979)

- general purpose search mechanisms trying to string together elementary reasoning steps to find complete solutions
- Weak methods – general but cannot scale up
- Alternative – more powerful, domain specific knowledge
- DENDRAL system – inferring molecular structure from the information provided by a mass spectrometer
- Expert systems – medical diagnosis

Abridged history of AI

- 1943 McCulloch & Pitts: Boolean circuit model of brain
- 1950 Turing's "Computing Machinery and Intelligence"
- 1956 Dartmouth meeting: "Artificial Intelligence" adopted
- 1952—69 Look, Ma, no hands!
- 1950s Early AI programs, including Samuel's checkers program, Newell & Simon's Logic Theorist, Gelernter's Geometry Engine
- 1965 Robinson's complete algorithm for logical reasoning
- 1966—73 AI discovers computational complexity
Neural network research almost disappears
- 1969—79 Early development of knowledge-based systems
- 1980-- AI becomes an industry
- 1986-- Neural networks return to popularity
- 1987-- AI becomes a science (speech recognition)
- 1995-- The emergence of intelligent agents

State of the art

- Deep Blue defeated the reigning world chess champion Garry Kasparov in 1997
- Proved a mathematical conjecture (Robbins conjecture) unsolved for decades
- No hands across America (driving autonomously 98% of the time from Pittsburgh to San Diego)
- During the 1991 Gulf War, US forces deployed an AI logistics planning and scheduling program that involved up to 50,000 vehicles, cargo, and people
- NASA's on-board autonomous planning program controlled the scheduling of operations for a spacecraft
- `Proverb` solves crossword puzzles better than most humans

Deep Blue

(taken from Wikipedia)



- The computer system dubbed "Deep Blue" was the first machine to win a chess game against a reigning world champion (Garry Kasparov) under regular time controls. This first win occurred on February 10, 1996. Deep Blue - Kasparov, 1996, Game 1 is a famous chess game. However, Kasparov won 3 games and drew 2 of the following games, beating Deep Blue by a score of 4–2. The match concluded on February 17, 1996.
- Deep Blue was then heavily upgraded (unofficially nicknamed "Deeper Blue") and played Kasparov again in May 1997, winning the six-game rematch 3.5–2.5, ending on May 11th, finally ending in game six. Deep Blue thus became the first computer system to defeat a reigning world champion in a match under standard chess tournament time controls.
- The project was started as "ChipTest" at Carnegie Mellon University by Feng-hsiung Hsu; the computer system produced was named Deep Thought after the fictional computer of the same name from The Hitchhiker's Guide to the Galaxy. Hsu joined IBM (Research division) in 1989 and worked with Murray Campbell on parallel computing problems. Deep Blue was developed out of this. The name is a play on Deep Thought and Big Blue, IBM's nickname.
- The system derived its playing strength mainly out of brute force computing power. It was a massively parallel, 30-node, RS/6000, SP-based computer system enhanced with 480 special purpose VLSI chess chips. Its chess playing program was written in C and ran under the AIX operating system. It was capable of evaluating 200,000,000 positions per second, twice as fast as the 1996 version. In June 1997, Deep Blue was the 259th most powerful supercomputer, capable of calculating 11.38 gigaflops, although this did not take into account Deep Blue's special-purpose hardware for chess.