

AI Methodology

Theoretical aspects

- Mathematical formalizations, properties, algorithms

Engineering aspects

- The act of building (useful) machines

Empirical science

- Experiments

What's involved in Intelligence?

A) Ability to interact with the real world

to perceive, understand, and act

speech recognition and understanding (natural language)

image understanding (computer vision)

B) Reasoning and Planning

modeling the external world

problem solving, planning, and decision making

ability to deal with unexpected problems, uncertainties

CS4700

C) Learning and Adaptation

Lots of data. Use to train statistical models.

We are continuously learning and adapting.

We want systems that adapt to us!

AI Leverages from different disciplines

philosophy

e.g., foundational issues (can a machine think?), issues of knowledge and believe, mutual knowledge

psychology and cognitive science

e.g., problem solving skills

neuro-science

e.g., brain architecture

computer science and engineering

e.g., complexity theory, algorithms, logic and inference, programming languages, and system building.

mathematics and physics

e.g., statistical modeling, continuous mathematics, statistical physics, and complex systems.

Historical Perspective

Obtaining an understanding of the human mind is one of the final frontiers of modern science.

Founders:

George Boole, Gottlob Frege, and Alfred Tarski

- formalizing the laws of human thought

Alan Turing, John von Neumann, and Claude Shannon

- thinking as computation

John McCarthy (Stanford), Marvin Minsky (MIT),
Herbert Simon and Allen Newell (CMU)

- the start of the field of AI (1959)

History of AI: The gestation of AI 1943-1956

1943 McCulloch and Pitts

- McCulloch and Pitts' model of artificial neurons
- Minsky's 40-neuron network

1950 Turing's "Computing machinery and intelligence"

1950s Early AI programs, including Samuel's checkers program, Newell and Simon's Logic theorist

1956 Dartmouth meeting : Birth of "Artificial Intelligence"

- 2-month Dartmouth workshop; 10 attendees
- Name was chosen. AI

History of AI:

(1952-1969)

Early enthusiasm, great expectations

1957 Herb Simon:

It is not my aim to surprise or shock you – but the simplest way I can summarize is to say that there are now in the world machines that think, that learn, and that create. 😊

1958 John McCarthy's LISP (symbol processing at core)


1965 J.A. Robinson invents the resolution principle, basis for automated theorem. **General reasoning procedure.**

Limited intelligent reasoning in **microworlds**


(such as the “blocks world” --- a toy robotics domain)

The Blocks World

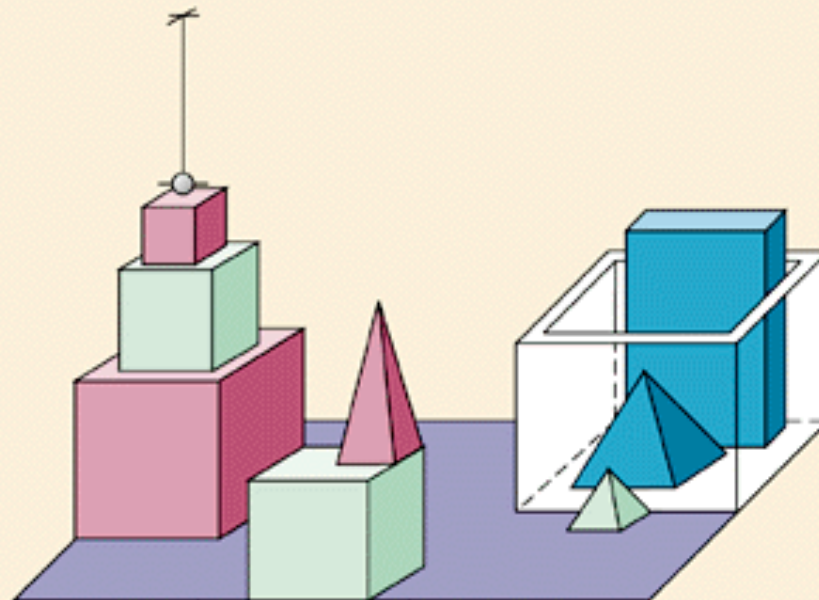
gripper



(a) "Pick up a big red block."



(b) "Find a block which is taller than the one you are holding and put it into the box."



(c) "Will you please stack up both of the red blocks and either a green cube or a pyramid?"

History of AI

A dose of reality (1965 - 1978)

1) Weizenbaum's ELIZA (“fools” users)

Capturing general knowledge is hard.

2) Difficulties in automated translation

See Babelfish

Syntax and dictionaries are not enough

Consider going from **English** to **Russian** back to **English**.

“The spirit is willing but the flesh is weak.”

“The vodka is good but the meat is rotten.”

*Natural language processing (NLP) is hard.
(Ambiguity! Context! Anaphora resolution.)*

History of AI

A dose of reality, cont. (1965 - 1978)

3) Cars climbing up trees (at CMU)...

Road sides look like parallel lines.

But, unfortunately, so do trees!

Computer vision is hard.
(Ambiguity! Context! Noisy pixels.)

4) Limitations of perceptrons discovered Minsky and Papert (1969)

Can only represent linearly separable functions

Neural network research almost disappears

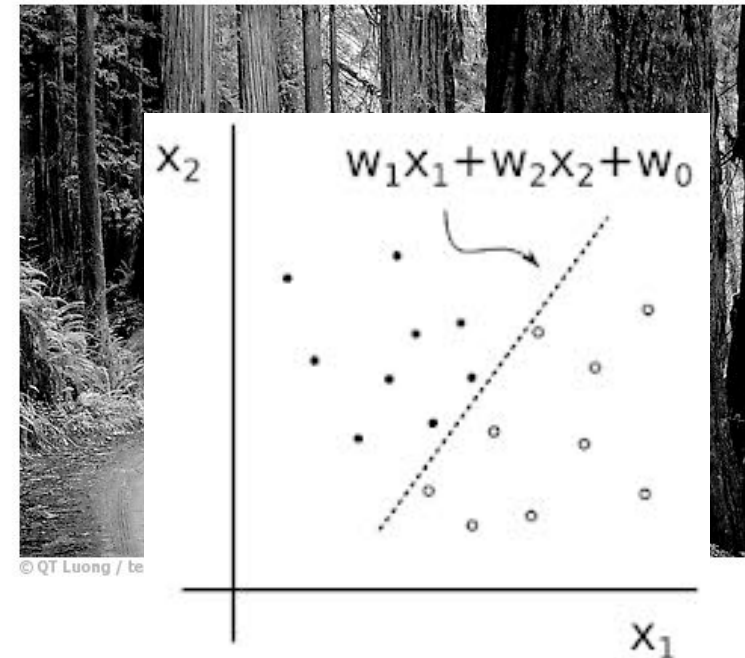
Machine learning is hard.

5) Intractability of inference. NP-Completeness (Cook 72)

Intractability of many problems attempted in AI.

Worst-case result....

Machine reasoning is hard.



History of AI

Knowledge based systems (1969-79)

Intelligence requires knowledge

Knowledge-based systems (lots of knowledge with limited but fast reasoning)

(Feigenbaum)

versus

general “weak” methods (a few basic principles with general reasoning)

(Simon and Newell)

Some success: *Expert Systems*

- Mycin: diagnose blood infections (medical domain)
- R1 : configuring computer systems
- AT&T phone switch configuration

**Surprising insight:
Modeling medical
expert easier than
modeling
language / vision /
reasoning of
3 year old.
(not foreseen)**

Expert Systems

Very expensive to code. (\$1M+)

Response: Try to learn knowledge from data.

Weak with uncertain inputs / noisy data / partial information

Response: Incorporate probabilistic reasoning

Brittle! (fail drastically outside domain)

Leads to 1980 -- 1995:

--- **General foundations reconsidered**

--- **Foundations of machine learning established (e.g. computational learning theory; PAC learning; statistical learning)**

--- **Foundations of probabilistic formalisms: Bayesian reasoning; graphical models; mixed logical and probabilistic formalisms.**

From 1995 onward:


--- **Data revolution combined with statistical methods**

--- **Building actual systems**

--- **Human world expert performance matched (and exceeded) in certain domains**

History of AI: 1995 - present

Several success stories with high impact ...



Machine Learning

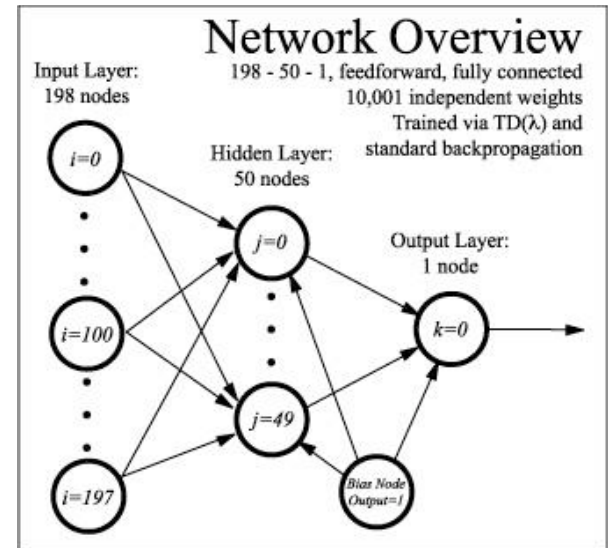
In '95, TD-Gammon.

World-champion level play by Neural Network that **learned from scratch** by playing millions and millions of *games against itself!* (about 4 months of training. Temporal-Difference learning.)

(initial games hundreds of moves)


Has changed human play.

*Remaining open question: Why does this **NOT** work for, e.g., chess??*



1996 --- EQP: Robbin's Algebras are all Boolean

A mathematical conjecture (Robbins conjecture) unsolved for 60 years!



**First creative mathematical
proof by computer.**

**Contrast with brute-force based proofs
such as the 4-color theorem.**

The Robbins problem was to determine whether one particular set of rules is powerful enough to capture all of the laws of Boolean algebra.

Mathematically:

Can the equation $\text{not}(\text{not}(P)) = P$ be derived from the following three equations?

[1] $(P \text{ or } Q) = (Q \text{ or } P)$

[2] $(P \text{ or } Q) \text{ or } R = P \text{ or } (Q \text{ or } R),$

[3] $\text{not}(\text{not}(P \text{ or } Q) \text{ or } \text{not}(P \text{ or } \text{not}(Q))) = P.$

*[An Argonne lab program] has come up with a major mathematical proof that would have been called **creative if a human had thought of it.***
New York Times, December, 1996

<http://www-unix.mcs.anl.gov/~mccune/papers/robbins/>

ROBBINS CONJECTURE

THE PROOF

7	$\overline{\overline{p+q+p+q}} = q$	[Robbins axiom]
10	$\overline{\overline{p+q+p+q+q}} = \overline{p+q}$	[7 → 7]
11	$\overline{\overline{\overline{p+q+p+q+q}}} = \overline{\overline{p+q}}$	[7 → 7]
29	$\overline{\overline{\overline{\overline{p+q+p+2q+p+q}}} = q$	[11 → 7]
54	$\overline{\overline{\overline{\overline{\overline{p+q+p+2q+p+q+r+q+r}}} = r$	[29 → 7]
217	$\overline{\overline{\overline{\overline{\overline{\overline{p+q+p+2q+p+q+r+q+r+r+q+r}}} = \overline{q+r}$	[54 → 7]
674	$\overline{\overline{\overline{\overline{\overline{\overline{\overline{p+q+p+2q+p+q+r+q+r+r+r+s+q+r+s}}} = s$	[217 → 7]
6736	$\overline{\overline{\overline{\overline{\overline{\overline{\overline{\overline{\overline{3p+p+3p+3p+p+5p}}} = \overline{3p+p}$	[10 → 674]
8855	$\overline{\overline{\overline{\overline{\overline{\overline{\overline{\overline{\overline{\overline{3p+p+5p}}} = \overline{3p}$	[6736 → 7, simp : 54]
8865	$\overline{\overline{\overline{\overline{\overline{\overline{\overline{\overline{\overline{\overline{\overline{3p+p+3p+2p+3p}}} = \overline{3p+p+2p}$	[8855 → 7]
8866	$\overline{\overline{\overline{\overline{\overline{\overline{\overline{\overline{\overline{\overline{\overline{3p+p+3p}}} = p$	[8855 → 7, simp : 11]
8870	$\overline{\overline{\overline{\overline{\overline{\overline{\overline{\overline{\overline{\overline{\overline{3p+p+3p+q+p+q}}} = q$	[8866 → 7]
8871	$\overline{\overline{\overline{\overline{\overline{\overline{\overline{\overline{\overline{\overline{\overline{3p+p+2p}}} = 2p$	[8865, simp : 8870]

A Baker's Dozen. The key steps in proving the Robbins conjecture, as reported by EQP, an automated theorem-proving program developed by William McCune and colleagues at Argonne National Laboratory. (See Box, "Substitute Teacher," page 63 for details.)

As Easy

The phrase *il est* repeatedly in the French mathematician phrase is common in spelling out details th mathematical techniques. words. What Laplace intense mathematical

Oddly enough, La see how a particular was easy to see. Th right combination of Could a computer have.

William McCune, Laboratory in Illinois proving programs.) dictis on a range of pr His most powerful pi can be every bit as program considers e: the program's creato

Cryptic or not, E announced a solution bolic logic that hat posed in the 1930's. has been solved by described as reasoni

"It's a clear lan Stanley Burris, a lo; Canada. "Now tha problem, it opens th

The Robbins co which are named (1815-1864), who into algebraic expr expressed as $p + p + N(q)$; a logic

Note: Same order of search complexity as performed by Deep Blue per move. Quantative threshold for creativity?

1997: Deep Blue beats the World Chess Champion



Deep Blue had Kasparov in deep thought
(CNN)

vs.



I could feel human-level intelligence across the room
Gary Kasparov, World Chess Champion (human...)

Note: when training in self-play,
be careful to randomize!

Deep Blue vs. Kasparov



Game 1: 5/3/97:
Kasparov wins

Game 2: 5/4/97:
Deep Blue wins

Game 3: 5/6/97:
Draw

Game 4: 5/7/97:
Draw

Game 5: 5/10/97:
Draw

Game 6: 5/11/97:
Deep Blue wins

Game 3:
Why did
Kasparov not
simply repeat
moves from
game 1?

*The value of IBM's stock
increased by \$18 Billion!*

We'll discuss Deep Blue's architecture, when we
cover *multi-agent search*.

On Game 2

Game 2 - Deep Blue took an early lead. Kasparov resigned, but it turned out he could have forced a draw by perpetual check.

Interestingly, if Kasparov had been playing a human he would most likely not have resigned.

This was real chess. This was a game any human grandmaster would have been proud of.

Joel Benjamin

grandmaster, member Deep Blue team

Kasparov on Deep Blue

1996: Kasparov Beats Deep Blue

“I could feel --- I could smell --- a new kind of intelligence across the table.” (CNN)

1997: Deep Blue Beats Kasparov

“Deep Blue hasn't proven anything.” 😊

Current strongest play: Computer-Human hybrid

May, '97 --- Deep Blue vs. Kasparov. First match won against world-champion. ``intelligent creative" play.
200 million board positions per second!

Kasparov: ... still understood 99.9 of Deep Blue's moves.

Deep Blue considers *60 billion boards per move!* Human?


Around 10 to 20 lines of play. Hmm...

Intriguing issue: How does human cognition deal with the search space explosion of chess?

Or how can humans compete with computers at all?? (What does human cognition do? Truly unknown...)

1999: Remote Agent takes Deep Space 1 on a galactic ride

Deep Space One



For two days in May, 1999, an AI Program called Remote Agent *autonomously ran* Deep Space 1 (some 60,000,000 miles from earth)

NASA: Autonomous Intelligent Systems.

Engine control next generation spacecrafts.

Automatic planning and execution model.

Fast real-time, on-line performance.

Compiled into 2,000 variable logical reasoning problem.

Contrast: current approach customized software with ground control team. (e.g., Mars mission \$50M.)

Remote Agent:

1999 Winner of NASA's Software of the Year Award

NASA engineers were at first very reluctant. Multi-million dollar mission on the line. On-board computer hardware was very limited till mid-nineties.



It's one small step in the history of space flight. But it was one giant leap for computer-kind, with a state-of-the-art AI system being given primary command of a spacecraft.


Known as Remote Agent, the software operated NASA's Deep Space 1 spacecraft and its futuristic ion engine during two experiments that started on Monday, May 17, 1999. For two days, Remote Agent ran on the on-board computer of Deep Space 1, more than 60,000,000 miles (96,500,000 kilometers) from Earth. The tests were a step toward robotic explorers of the 21st century that are less costly, more capable and more independent from ground control.

<http://ic.arc.nasa.gov/projects/remote-agent/index.html>

Robocup @ Cornell --- Raff D'Andrea 2000




RoboCup Japan open 2013



The Amazon logo, featuring the word "amazon" in a bold, black, sans-serif font with a yellow curved arrow underneath it, pointing from the letter 'a' to the letter 'z'.

Kiva Systems \$700M

**From Robocup to
Warehouse Automation**



2005 Autonomous Control: DARPA GRAND CHALLENGE



October 9, 2005

Stanley and the Stanford Racing Team were awarded 2 million dollars for being the first team to complete the 132 mile DARPA Grand Challenge course (Mojave Desert). Stanley finished in just under 6 hours 54 minutes and averaged over 19 miles per hour on the course.

**Sebastian Thrun:
Google's driverless car (2011)**

**Cornell team stuck ☹
due to malfunctioning GPS.**

<http://www.youtube.com/watch?v=bp9KBrH8H04>

path planning

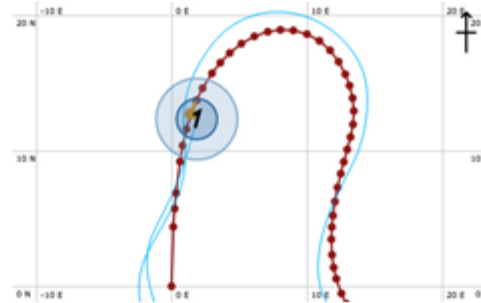
A* algorithm
Covered in
search and
problem solving.

Path Planning Overview

Path planning is the basic process by which our vehicle decides on what path to take through the world. The A.I. uses the world model created by the sensors, the GPS waypoints provided during the race by DARPA, and a road following algorithm to pick a best path.

Road Following

The road following algorithm uses color differences, shadowing, and edge-detection algorithms to detect the sides of a road (if there is a road) and then decides if the road is turning, going straight, which direction, how sharply, etc. The road following algorithm uses input from most of the vehicle sensors, and provides the A.I. with probable road characteristics.



The red line represents the ideal path picked by our A.I., and the turquoise path represents the actual path traveled by our vehicle. The differences arise because 1) we did not start our vehicle on the ideal path, and 2) our vehicle must, without exceeding its performance limits, avoid small obstacles such as boulders.

Cornell: 4th!

Also, in historic

1st autonomous driverless car

collision. Rear-ended by MIT car!

2007 Darpa Urban Challenge

Winner: CMU Tartan Racing's Boss

2007 Darpa Urban Challenge

The **Urban Challenge** will pit driverless vehicles against one another on city streets. Robots will have to handle traffic, intersections, rules of the road and other robots. The challenge is a high-stakes competition that plays out on a world stage. The prize is \$2M, but the payoff for driver safety is much greater. This competition will be held November 3, 2007.

The Urban Challenge is third in a series of autonomous vehicle competitions designed to catalyze robotic technology development. On October 8, 2005, Carnegie Mellon's "Sandstorm" and "Highlander" crossed the finish line of DARPA Grand Challenge after successfully completing a 132-mile course through the Nevada desert, coming in second and third place respectively.



URBAN CHALLENGE



<http://www.tartanracing.org/blog/index.html#26>


Watson: Question-Answering system, 2011



Watson defeats the
two greatest Jeopardy!
champions

<http://www.youtube.com/watch?v=dr7IxQeXr7g>

WATSON



Neural Networks --- Deep Learning, 2012.

New York Times: “Scientists See Promise in Deep-Learning Programs,” Saturday, Nov. 24, 2012.

<http://www.nytimes.com/2012/11/24/science/scientists-see-advances-in-deep-learning-a-part-of-artificial-intelligence.html?hpw>

Multi-layer neural networks, a resurgence!

- a) Winner one of the most recent learning competitions**
- b) Automatic (unsupervised) learning of “cat” and “human face” from 10 million of Google images; 16,000 cores 3 days; multi-layer neural network (Stanford & Google).**
- c) Speech recognition and real-time translation (Microsoft Research, China).**

Aside: see web site for great survey article “A Few Useful Things to Know About Machine Learning” by Domingos, CACM, 2012.



Start at min. 3:00. Deep Neural Nets in speech recognition.₃₃

Other promising ongoing efforts

- 1) Intelligent autonomous assistants, e.g., iPhone's Siri
(still a long way to go 😊) Integrated, autonomous agents.
Google Glass will be the next step. Location / context aware;
rich sensing, vision and speech understanding and generation.
- 2) Fully self-driving car (Google; assisted driving Mercedes and BMW
--- the cost of a car is becoming software and sensors Incredibly
more lines of code in a Mercedes than in a Boeing 747.)
- 2) Google translate. Reaches around 70% of human translator
performance. Almost fully a purely statistical approach.

Not clear yet how far one can go without a real understanding of the semantics (meaning). But with Big Data, statistical methods already went much further than many researchers had considered possible only 10 years ago.

Course Administration



What is Artificial Intelligence?



Course Themes, Goals, and Syllabus

Setting expectations for this course

Are you going to build real systems and robots?

NO...

Goal:

Introduce the conceptual framework and computational techniques that serve as a foundation for the field of artificial intelligence (AI).

Syllabus

- **Structure of intelligent agents and environments.**
- **Problem solving by search: principles of search, uninformed (“blind”) search, informed (“heuristic”) search, and local search.**
- **Constraint satisfaction problems: definition, search and inference, and study of structure.**
- **Adversarial search: games, optimal strategies, imperfect, real-time decisions.**
- **Logical agents: propositional and first order logic, knowledge bases and inference.**
- **Uncertainty and probabilistic reasoning: probability concepts, Bayesian networks, probabilistic reasoning over time, and decision making.**
- **Learning: inductive learning, concept formation, decision tree learning, statistical approaches, neural networks, reinforcement learning.**

So far, we discussed

Artificial Intelligence and characteristics of *intelligent* systems.

Brief history of AI

Major recent AI achievements

Reading: Chapter 1 Russell & Norvig