

Artificial Intelligence

CSE P573

Mausam


(Based on Slides by Stuart Russell, Henry Kautz,
Subbarao Kambhampati, and UW-AI faculty)

Logistics

- Instructor: Mausam, CSE 454, mausam@cs.washington.edu
- TA: Dvijotham Krishnamurthy, dvij@cs.washington.edu
- Course Website: www.cs.washington.edu/csep573
- Join class mailing list (instructions on website)
- Text: Artificial Intelligence: A Modern Approach (3rd edition), Russell and Norvig
- 4 programming assignments
- 6-9 mini-assignments/quizzes
- Grading:
 - 50% programming assignments
 - 20% short written assignments
 - 30% final

MOTIVATION

1946: ENIAC heralds the dawn of Computing



1950: Turing asks the question....




I propose to consider the question:

“Can machines think?”

--Alan Turing, 1950

1956: A new field is born

- We propose that a 2 month, 10 man study of artificial intelligence be carried out during the summer of 1956 at Dartmouth College in Hanover, New Hampshire.
- - [Dartmouth AI Project Proposal](#); J. McCarthy et al.; Aug. 31, 1955.



1996: EQP proves that Robbin's Algebras are all boolean



```
----- EQP 0.9, June 1996 -----  
The job began on eyas09.mcs.anl.gov, Wed Oct 2 12:25:37 1996  
UNIT CONFLICT from 17666 and 2 at 678232.20 seconds.  
----- PROOF -----  
2 (wt=7) [] -(n(x + y) = n(x)).  
3 (wt=13) [] n(n(n(x) + y) + n(x + y)) = y.  
5 (wt=18) [para(3,3)] n(n(n(x + y) + n(x) + y) + y) = n(x + y).  
6 (wt=19) [para(3,3)] n(n(n(n(x) + y) + x + y) + y) = n(n(x) + y).  
.....  
17666 (wt=33) [para(24,16426),demod([17547])] n(n(n(x) + x) ...
```

[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

1997: HAL 9000 becomes operational in fictional Urbana, Illinois



...by now, every intelligent person knew that
H-A-L is derived from Heuristic Algorithmic

-Dr. Chandra, 2010: Odyssey Two

1997: Deep Blue ends Human Supremacy in Chess



Deep Blue had Kasparov in deep thought
(CNN)

vs.



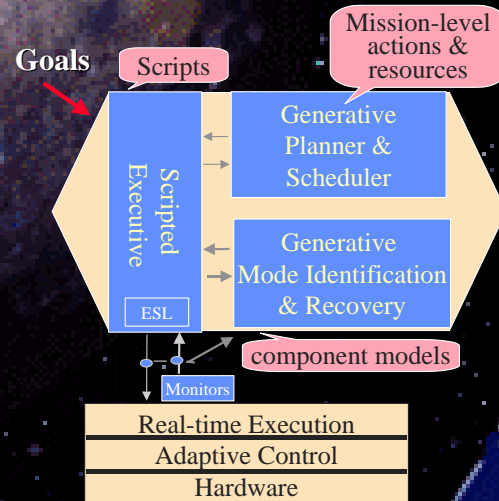
The cunning Deep Blue (CNN)

I could feel human-level intelligence across the room

-Gary Kasparov, World Chess Champion (human)

In a few years, even a single victory
in a long series of games would be the triumph of human genius.

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



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

2004 & 2009





- Provide a standard problem where a wide range of technologies can be integrated and examined
- By 2050, develop a team of fully autonomous humanoid robots that can win against the human world champion team in soccer.



<http://www.youtube.com/watch?v=Cv7333wHFMM>

2005: Cars Drive Themselves

- Stanley and three other cars drive themselves over a 132 mile mountain road



http://www.youtube.com/watch?v=XOgkNh_IPjU



2007: Robots Drive on Urban Roads



11 cars drove themselves on urban streets (for DARPA Urban Challenge)



Recentmost Success 2011



And Ken Jennings pledges obeisance to the new Computer Overlords..

Europa Mission ~ 2018



About the only thing Microsoft & Google can agree on these days...

- “If you invent a breakthrough in artificial intelligence, so machines can learn,” Mr. Gates responded, “that is worth 10 Microsofts.”
- No. 1: AI at human level in 10-20 year time frame
 - Sergey Brin &
 - Larry Page
 - (independently, when asked to name the top 5 areas needing research. Google Faculty Summit, July 2007)

Goals of this course

- A brief intro to the philosophy of AI
- A brief intro to the breadth of ideas in AI
- General computer scientist
 - general tools to aid in attacking a new problem
- Serious AI enthusiast
 - A primer from which to launch advanced study

Science of AI

Physics: Where did the *physical universe* come from?
And what laws guide its dynamics?

Biology: How did *biological life* evolve?
And how do living organisms function?

AI: What is the nature of *intelligent thought*?

What is intelligence?

- Dictionary.com: *capacity for learning, reasoning, understanding, and similar forms of mental activity*
- Ability to perceive and act in the world
- Reasoning: proving theorems, medical diagnosis
- Planning: take decisions
- Learning and Adaptation: recommend movies, learn traffic patterns
- Understanding: text, speech, visual scene

Intelligence vs. humans

- Are humans intelligent?
- Are humans rational?
- Can non-human behavior be intelligent?

What is *artificial* intelligence?

human-like vs. rational

thought
vs.
behavior

“[automation of] activities that we associate with human thinking, activities such as decision making, problem solving, learning...” (Bellman 1978)	“The study of mental faculties through the use of computational models” (Charniak & McDermott 1985)
“The study of how to make computers to things at which, at the moment, people are better” (Rich & Knight 1991)	“The branch of computer science that is concerned with the automation of intelligent behavior” (Luger & Stubblefield 1993)

What is *artificial* intelligence?

human-like vs. rational

thought
vs.
behavior

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

Thinking Humanly

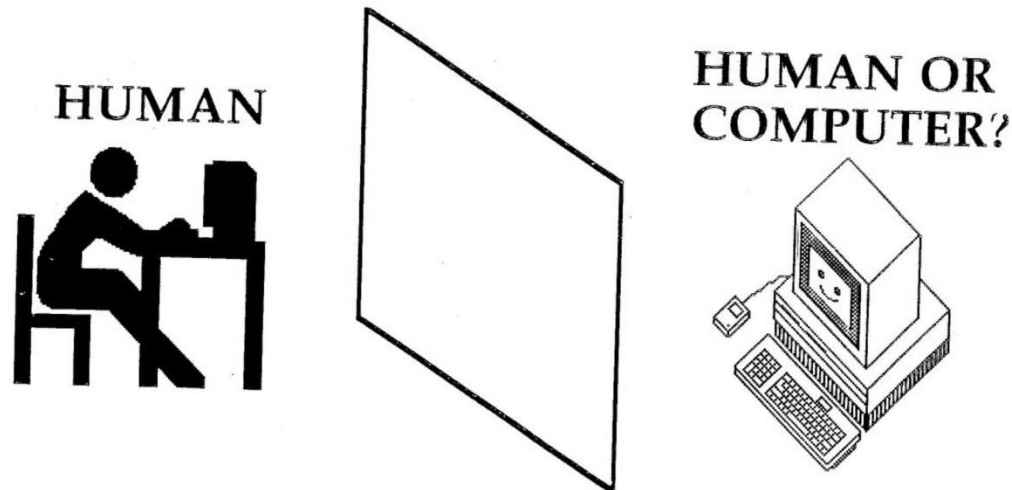
- Cognitive Science
 - Very hard to understand how humans think
 - Post-facto rationalizations, irrationality of human thinking
- Do we want a machine that beats humans in chess or a machine that *thinks like humans* while beating humans in chess?
 - Deep Blue supposedly DOESN'T think like humans..
- Thinking like humans important in applications
 - Intelligent tutoring
 - Expressing emotions in interfaces... HCI
- The goal of aeronautical engg is not to fool pigeons in flying!

Thinking Rationally: laws of thought

- Aristotle: what are correct arguments/thought processes?
 - Logic
- Problems
 - Not all intelligent behavior is mediated by logical deliberation (reflexes)
 - What is the purpose of thinking?

Acting Humanly: Turing's Test

- If the human cannot tell whether the responses from the other side of a wall are coming from a human or computer, then the computer is intelligent.



Acting Humanly

- Loebner Prize
 - Every year in Boston
 - Expertise-dependent tests: limited conversation
- What if people call a human a machine?
 - Shakespeare expert
- Problems
 - Not reproducible, constructive or mathematically analyzable

Acting rationally

- Rational behavior: doing the right thing
- Need not always be deliberative
 - Reflexive
- Aristotle (Nicomachean ethics)
 - Every art and every inquiry, and similarly every action and every pursuit is thought to aim at some good.

What is *artificial* intelligence (agent view)


- An **agent** is anything that can be viewed as **perceiving** its **environment** through **sensors** and **acting** upon that environment through **actuators**
- Human agent:
 - eyes, ears, and other organs for sensors
 - hands, legs, mouth, and other body parts for actuators
- Robotic agent:
 - cameras and laser range finders for sensors
 - various motors for actuators
- We will revisit this view in detail later in the course

Autonomous Systems

- In the 1990's there was a growing concern that work in classical AI ignored crucial scientific questions:
 - How do we **integrate the components** of intelligence (*e.g.* learning & planning)?
 - How does **perception** interact with reasoning?
 - How does the demand for **real-time performance** in a complex, changing environment affect the **architecture** of intelligence?

Examples of Agents

- Robots
- Intelligent buildings
- Autonomous spacecraft
- Web agents



AI as Engineering

- How can we make software systems more powerful and easier to use?
 - Speech & intelligent user interfaces
 - Autonomic computing
 - Mobile robots, softbots & immobots
 - Data mining
 - Medical expert systems
 - ...

What is *artificial* intelligence (algorithmic view)

- A large number of problems are NP hard
- AI develops a set of tools, heuristics, ...
 - to solve such problems in practice
 - for naturally occurring instances
- Search
- Game Playing
- Planning
- ...

Examples: Mundane Tasks

- Perception
 - Vision
 - Speech
- Natural Language
 - Understanding
 - Generation
 - Translation
- Reasoning
- Robot Control

Examples: Formal Tasks

- Games
 - Chess
 - Checkers
 - Othello
- Mathematics
 - Logic
 - Geometry
 - Calculus
 - Proving properties of programs

Examples: Expert Tasks

- Engineering
 - Design
 - Fault Finding
 - Manufacturing planning
- Medical
 - Diagnosis
 - Medical Image Analysis
- Financial
 - Stock market predictions

Recurrent Themes

- **Logic vs. Probability**

- In 1950's, logic dominates (McCarthy, ...

- attempts to extend logic

- 1988 – Bayesian networks (Pearl)

- efficient computational framework

- Today, no longer rivals

- Hot topic: combining probability & FOL

Recurrent Themes

- **Weak vs. Strong Methods**
 - Weak – general search methods (e.g., A^* search)
 - primarily for problem solving
 - not motivated by achieving human-level performance
 - Strong -- knowledge intensive (e.g., expert systems)
 - more knowledge \Rightarrow less computation
 - achieve better performance in specific tasks
- How to combine weak & strong methods seamlessly?

Recurrent Themes

- **Knowledge Representation**
 - “In knowledge lies the power”
 - Feature engineering in Machine Learning
 - Reformulation
- **Combinatorial Explosion**
- **Micro-world successes are hard to scale up.**
- **How to organize and accumulate large amounts of knowledge?**

Limits of AI Today

- Most of today's successful AI systems
 - operate in well-defined domains
 - employ narrow, specialized knowledge
- *Exceptions:*
 - *Watson???*
 - *Self-driving cars???*
- *Commonsense Knowledge*
 - needed in complex, open-ended worlds
 - Your kitchen vs. GM factory floor
 - understand unconstrained natural language

Role of Knowledge in Natural Language Understanding

- WWW Information Extraction
- Speech Recognition
 - “word spotting” feasible today
 - continuous speech – rapid progress
- Translation / Understanding
 - limited progress

The spirit is willing but the flesh is weak. (English)

The vodka is good but the meat is rotten. (Russian)

How the heck do *we* understand?

- John **gave** Pete a book.
- John **gave** Pete a hard time.
- John **gave** Pete a black eye.
- John **gave** in.
- John **gave** up.
- John's legs **gave** out beneath him.
- It is 300 miles, **give** or take 10.

How to Get Commonsense?

- CYC Project (Doug Lenat, Cycorp)
 - Encoding 1,000,000 commonsense facts about the world by hand
 - Coverage still too spotty for use!
- Machine Learning
- Open Mind
- Mining from Wikipedia & the Web

Topics of this Course

- Breadth
 - Search
 - Planning
 - Constraint Satisfaction
 - Logic
 - Uncertainty
 - Learning
- Depth
 - ???