Artificial Intelligence
COL333/CSL671

Mausam

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

Personnel

- Instructor: Mausam, SIT 402, mausam@cse.iitd.ac.in

- TAs:
  - Yashoteja Prabhu (csz138234 at cse.iitd.ac.in)
    Himanshu Jain (csz138519 at cse.iitd.ac.in)
    Neetu Jindal (neetu at cse.iitd.ac.in)
    Prajna Upadhyay (csz138110 at cse.iitd.ac.in)
    Dilpreet Kaur (csz158041 at cse.iitd.ac.in)
    Harinder Pal (mcs142123 at cse.iitd.ac.in)
    Kapil Thakkar (mcs142124 at cse.iitd.ac.in)
    Madhur Gupta (cs5110282 at cse.iitd.ac.in)
    Mayank Raj (cs5110284 at cse.iitd.ac.in)
    Shiva Chandra (cs5110296 at cse.iitd.ac.in)

Logistics

- Timings: Tue/Thu/Fri 11-12
- Office hours
  - By appointment
- Course Website: www.cse.iitd.ac.in/~mausam/courses/csl333/autumn2015
- Join class discussion group on Piazza (access code csl333) https://piazza.com/iit_delhi/fall2015/csl333csl671/home
- Textbook:
  Artificial Intelligence: A Modern Approach (3rd edition), Russell and Norvig

Programming Assignments

- 5 programming assignments; one of them in two parts
  - some assignments may be done in teams of two (as per instructions)
  - no team can be repeated for a second assignment
  - late policy (penalty of 10% every day)
  - I/O error (penalty of 20%)
  - Logical error (penalty of 50% only under special permission)

Grading and Academic Integrity

- Grading:
  - 50% assignments
  - 10% Minor 1
  - 10% Minor 2
  - 30% Major
  - Extra credit: constructive class participation, and discussion group participation

- Academic Integrity
  - Cheating → negative penalty (and possibly more)
    - Exception: if one person/team is identified as cheater
      - Non-cheater gets a zero
  - http://www.willa.me/2013/12/the-top-five-unsanctioned-software.html

- Collaboration is good!!! Cheating is bad!!! Who is a cheater?
  - No sharing of part-code
  - No written/soft copy notes
  - Right to information rule
  - Kyunki saas bhi kabhi bahu thi Rule

Class Requirements & Prereqs

- Class requirements
  - Uses a variety of skills / knowledge:
    - Probability and statistics
    - Boolean Logic
    - Algorithms
    - Above average coding skills
  - You will often have to work to fill the gaps
- Official Prerequisites
  - Data structures
- Unofficial Prerequisites
  - A willingness to learn whatever background you are missing

Languages

- English ☺
- C++/Java/Python
  - Coding efficiency : python
  - Program efficiency : C++
- Your choice of language may give unfair disadvantage to you!

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

Theory vs. Modeling vs. Applications

- Lecture balance tilted towards modeling
- Assignment balance tilted towards applications
- Relatively few theorems and even fewer proofs
- Desired work – lots!

MOTIVATION

1946: ENIAC heralds the dawn of Computing

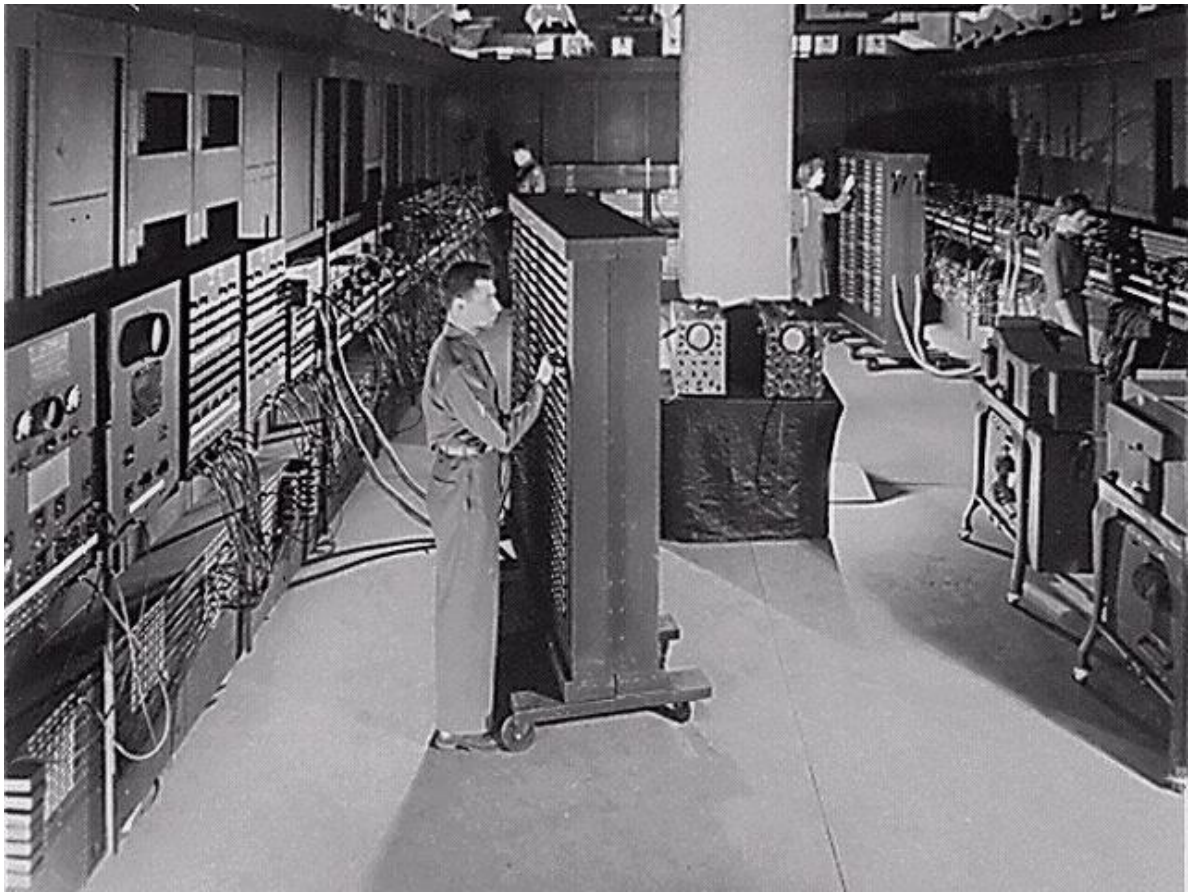

1950: Turing asks the question….

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.

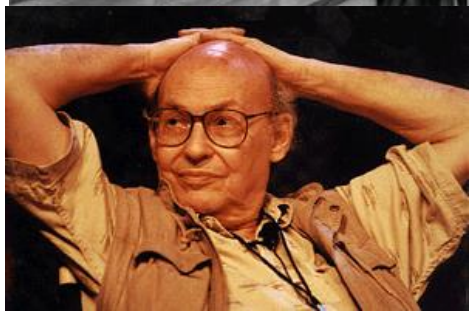

1996: EQP proves that Robbin's Algebras are all boolean

```
----- 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

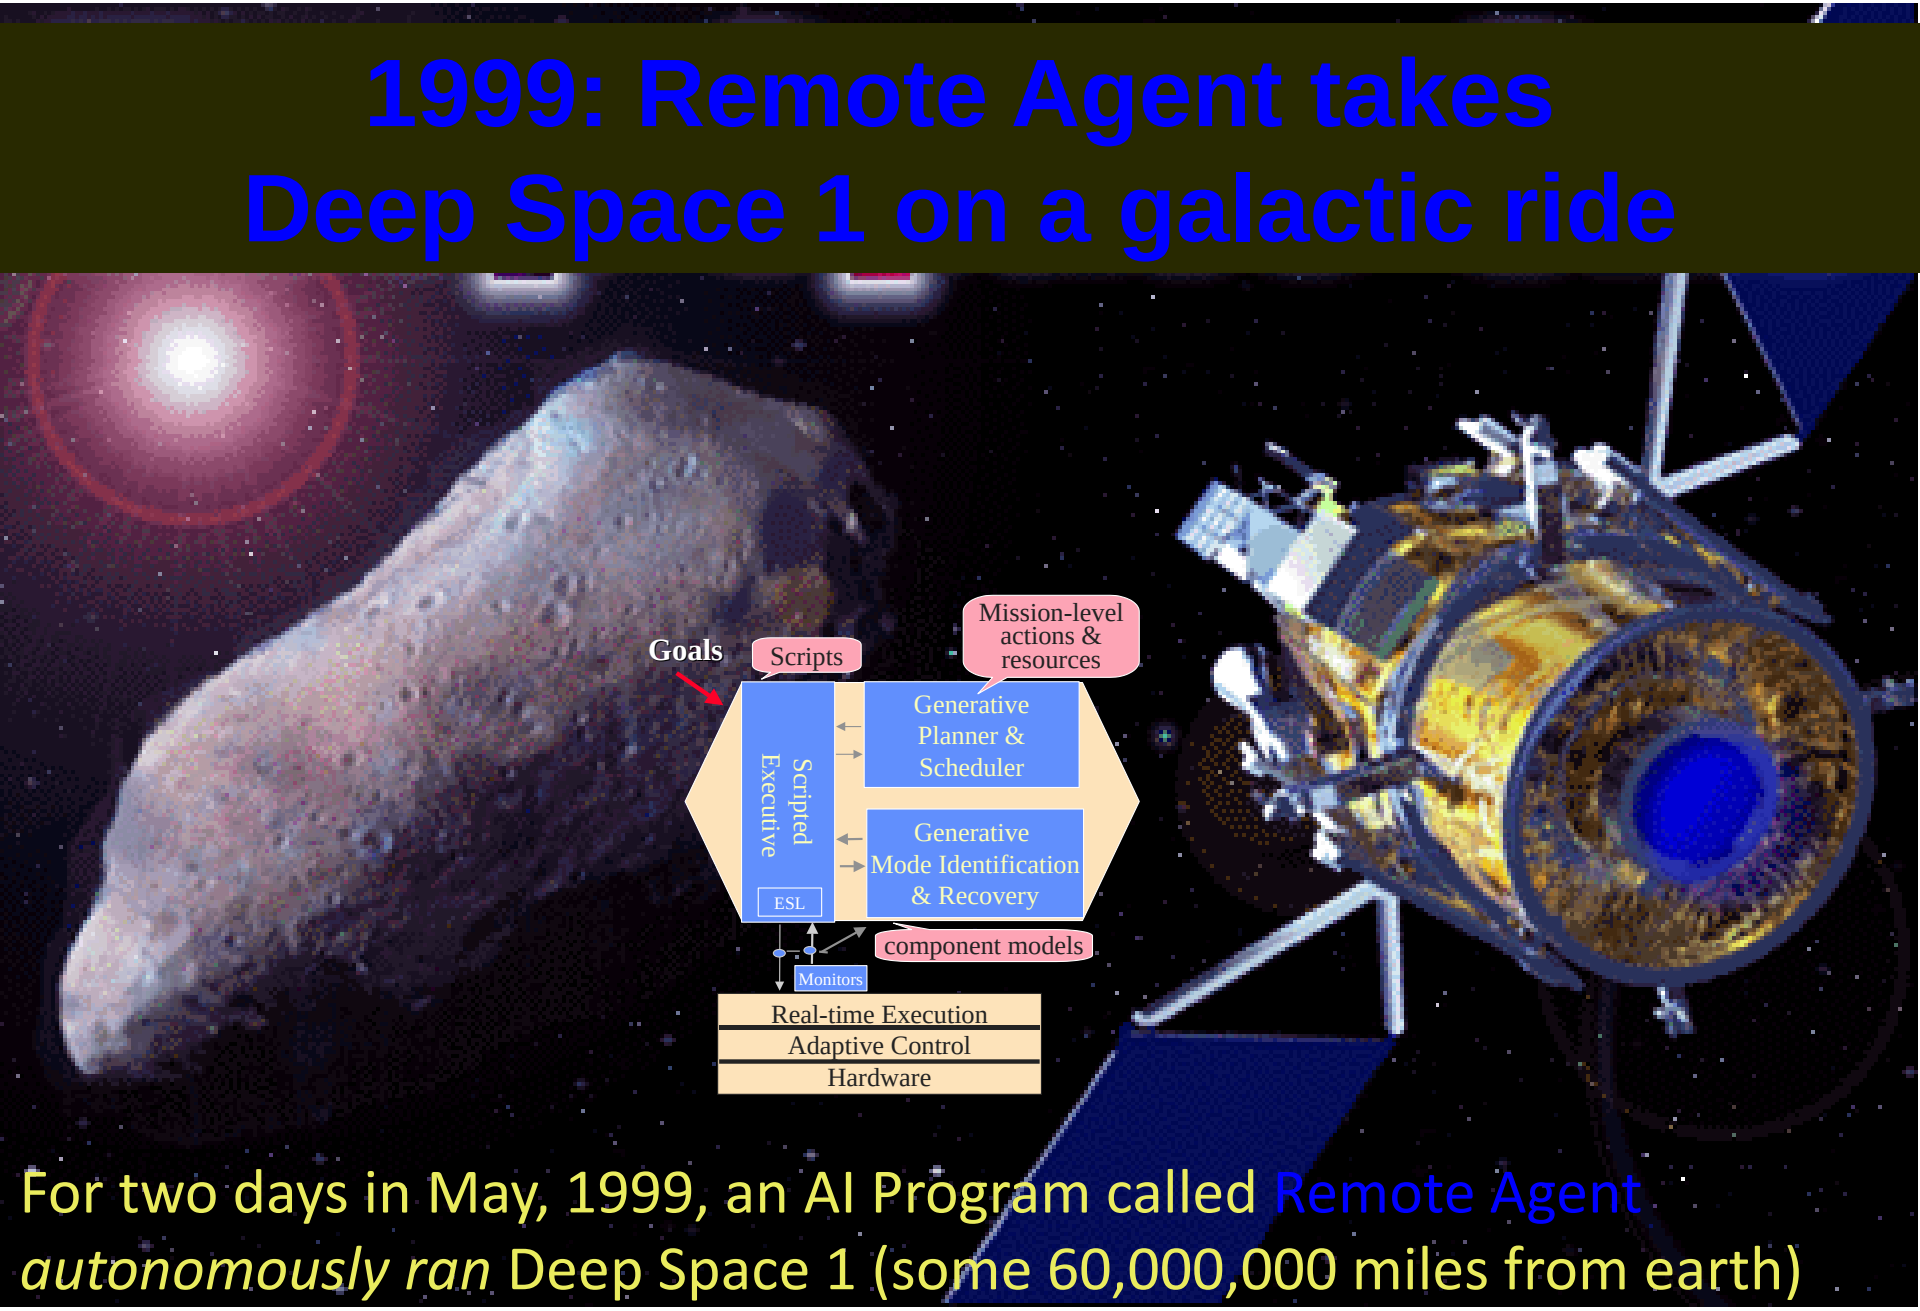

For two days in May, 1999, an AI Program called Remote Agent *autonomously ran* Deep Space 1 (some 60,000,000 miles from earth)

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

DARPA GRAND CHALLENGE

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

2015: Robots Threaten to Take all your jobs

..and thankfully

You step in to thwart them

by taking COL333

Welcome to the Holy War!

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)

PHILOSOPHY and THEMES

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?

thought *vs.* behavior	human-like *vs.* rational	
	"[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 & McDertmott 1985)
	"The study of how to make computers do 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?

thought *vs.* behavior \ human-like *vs.* rational	human-like *vs.* rational
thought	Systems that think like humans	Systems that think rationally
behavior	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.

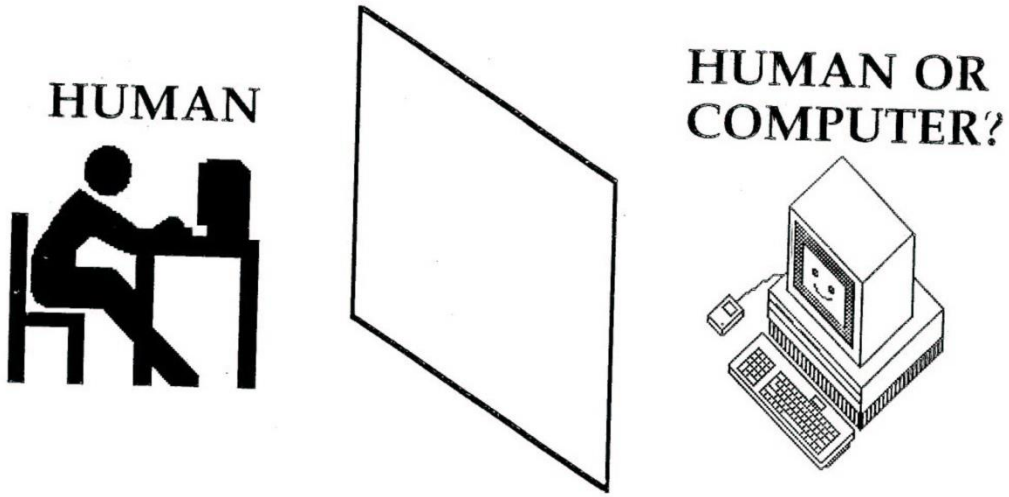

Acting Humanly

- Loebner Prize
  - Every year in Boston
  - Expertise-dependent tests: limited conversation
- What if people call a human a machine?
  - Shakespeare expert
  - Make human-like errors
- 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.

Acting → Thinking?

- Weak AI Hypothesis vs. Strong AI hypothesis
  - Weak Hyp: machines could act as if they are intelligent
  - Strong Hyp: machines that act intelligent have to think intelligently too

Rational Agents

- An agent should strive to do the right thing, based on what it can perceive and the actions it can perform. The right action is the one that will cause the agent to be most successful

- Performance measure: An objective criterion for success of an agent's behavior

- E.g., performance measure of a vacuum-cleaner agent could be amount of dirt cleaned up, amount of time taken, amount of electricity consumed, amount of noise generated, etc.

Ideal Rational Agent

"For each possible percept sequence, does whatever action is expected to maximize its performance measure on the basis of evidence perceived so far and built-in knowledge."

- Rationality vs omniscience?
- Acting in order to obtain valuable information

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 1990s 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

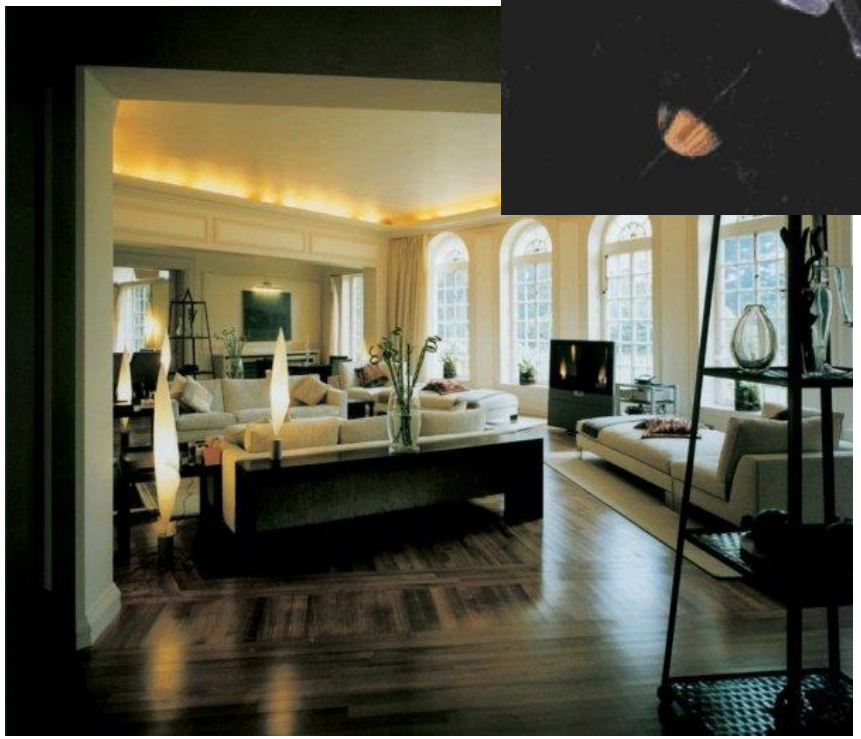

AI as Engineering

- 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 1950s, 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.* Knowledge-based Methods
  - Weak – general search methods (e.g., A* search)
    - primarily for problem solving
    - not motivated by achieving human-level performance
  - Strong AI -- 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

- Phase 1: Search, Planning, Constraint Satisfaction, Logic, Games

- Phase 2: Uncertainty (decision theory, probabilistic knowledge representation), Learning (supervised, unsupervised, reinforcement)

- Phase 3: Advanced topics and research talks