

Artificial Intelligence

COL333/COL671

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:
 - Yatin Nandwani (csz178057 at iitd.ac.in)
 - Keshav Sai Kolluru (csz178058 at iitd.ac.in)
 - Sachin Kumar Chauhan (csz188012 at iitd.ac.in)
 - Saransh Goyal (cs5150292 at iitd.ac.in)
 - Makkunda Sharma (cs5150459 at iitd.ac.in)
 - Phaneesh Barwaria (mcs182014 at iitd.ac.in)
 - Mehak (mcs182143 at iitd.ac.in)

Logistics

- Timings: Tue/Thu/Fri 11-12
- Office hours
 - By appointment
- Course Website:
www.cse.iitd.ac.in/~mausam/courses/col333/autumn2019
- Join class discussion group on Piazza (access code col333)
https://piazza.com/iit_delhi/fall2019/col333/home
- Textbook:
Artificial Intelligence: A Modern Approach (3rd edition), Russell and Norvig
- Communication: Piazza (no email)

Programming Assignments

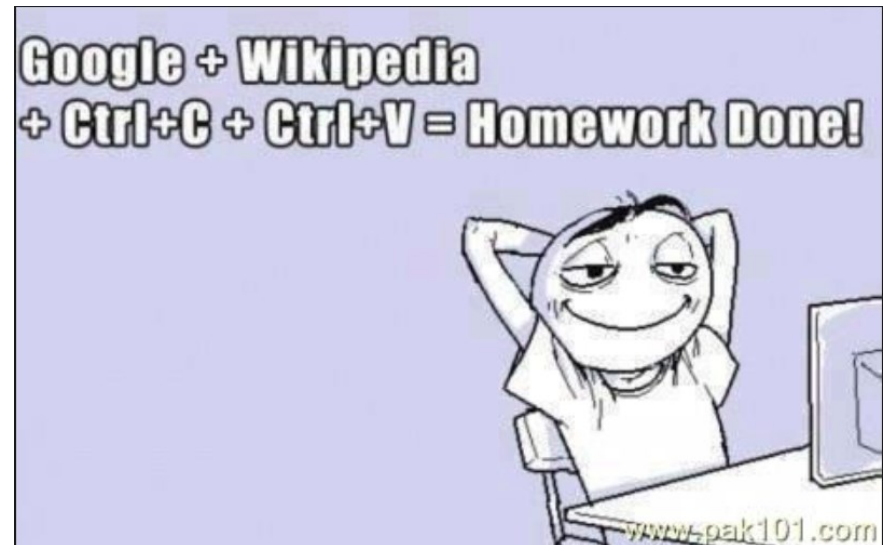
- 5 programming assignments; one of them in two parts
 - some assignments may be done in teams of two (as per instructions)
 - no team can be repeated for a second assignment (COL333)
 - no team can be repeated for a third assignment (COL671)
 - late policy (penalty of 10% every day)
 - I/O error (penalty of 20%)
 - Logical error (penalty of 50% only under special permission)

Grading

- Grading:
 - 50% assignments
 - 10% Minor 1
 - 10% Minor 2
 - 30% Major
 - Extra credit: constructive class participation, and discussion group participation
- Audit
 - 50% absolute in exams
 - C in course
- I believe in making course demanding rather than grading harsh
 - Make your life this semester miserable, rather than GPA beyond it low 😊

Academic Integrity

- Cheating → negative penalty (and possibly more)
 - Exception: if one person/team is identified as cheater
 - Non-cheater gets a zero
- 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!

Class Size

- Currently enrolled: don't know
- Expect to take: 150-160 students total.
- Others (fill the form on the course webpage)
 - <https://forms.gle/SZz7BEehuzKkQvAL8>
 - Grade in data structures
 - Total GPA
 - Minor degree: prev courses and grades
 - Prev relevant coursework (and grade)
 - Prev relevant project (and grade, if applicable)

Artificial Intelligence

Mausam

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

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!

HISTORY

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.




1956-1966

- 1950: Turing Test for Machine Intelligence
- 1956: AI born at Dartmouth College Workshop
- 1964: Eliza – the chatbot psychotherapist
- 1966: Shakey – general purpose mobile robot



AI Winters

- **1974 – 1980: Winter #1**
 - Failure of machine translation
 - Negative results in Neural nets
 - Poor speech understanding
- **1987 – 1993: Winter #2**
 - Decline of LISP
 - Decline of specialized hardware for expert systems
- **Lasting effects**
 - [Economist07] “Artificial Intelligence is associated with systems that have all too often failed to live up to their promises.”
 - [Pittsburgh BT06] “Some believe the word 'robotics' actually carries a stigma that hurts a company's chances at funding.”



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) [] $\neg(n(x + y) = n(x))$.

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])] $n(n(n(x) + x) \dots$

[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: 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.**

Success Story: Chess



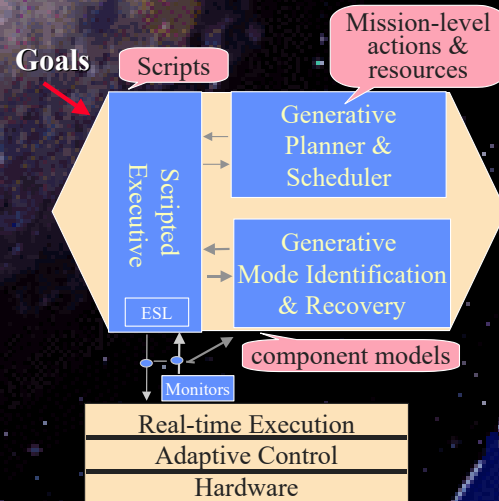
Does Deep Blue use AI?

“If it works, its not AI!”

Saying Deep Blue doesn't really think about chess is like saying an airplane doesn't really fly because it doesn't flap its wings.

– Drew McDermott

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



2005: Cars Drive Themselves

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



2005: Cars Drive Themselves

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




2011: IBM's Watson



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


2011: IBM's Watson




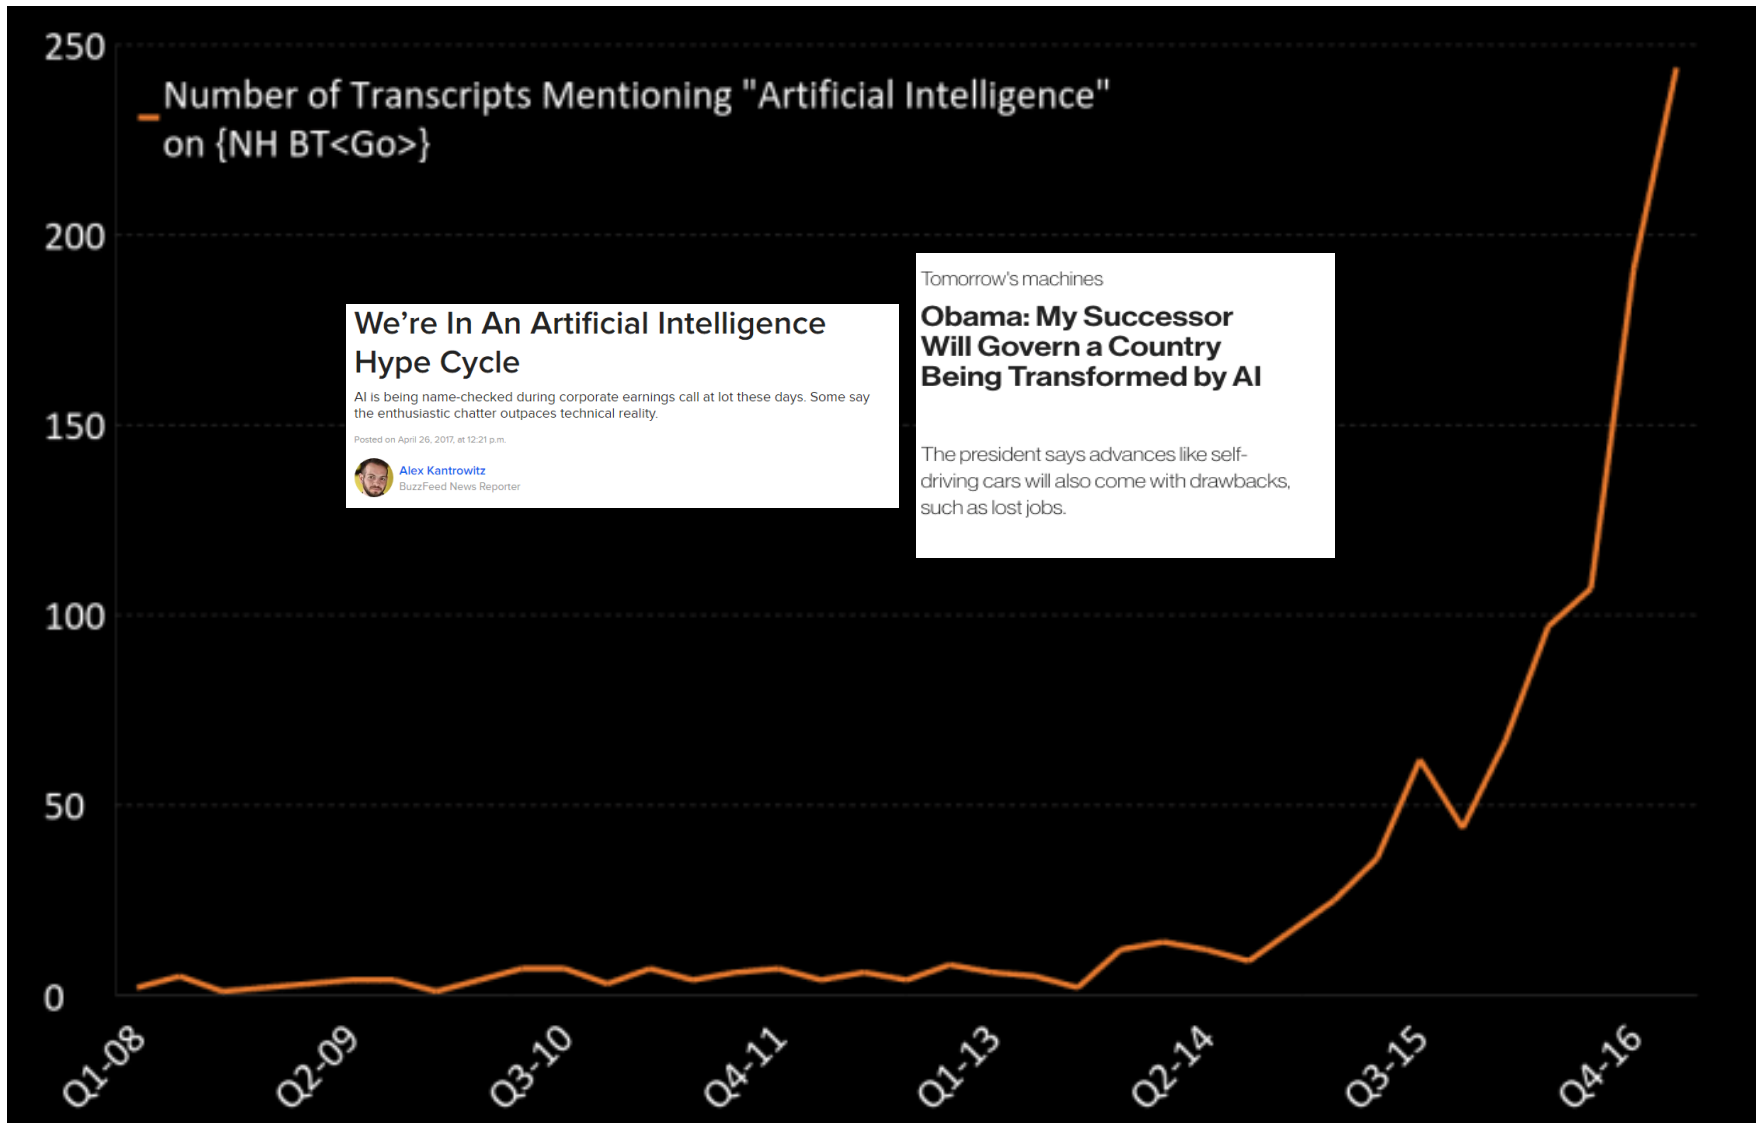
https://www.youtube.com/watch?v=WFR3IOm_xhE

PRESENT


2016: AlphaGo







What Changed?



Object Recognition




CIFAR

CANADIAN INSTITUTE
for ADVANCED RESEARCH




Network	Error	Layers
AlexNet	16.0%	8
ZFNet	11.2%	8
VGGNet	7.3%	19
GoogLeNet	6.7%	22
MS ResNet	3.6%	152!!

Artistic Applications



- Doodle to Painting!
- Style Transfer
- Image Colorization



<https://arxiv.org/pdf/1603.08511.pdf>
<https://arxiv.org/pdf/1603.01768.pdf>
<https://github.com/jcjohnson/fast-neural-style>

Image → Caption

Describes without errors



A person riding a motorcycle on a dirt road.

Describes with minor errors



Two dogs play in the grass.

Somewhat related to the image



A skateboarder does a trick on a ramp.



A group of young people playing a game of frisbee.




Two hockey players are fighting over the puck.



A little girl in a pink hat is blowing bubbles.

Automatic Speech Recognition



(c) <https://medium.com/@gaurav.sharma/voice-is-the-new-o-s-and-the-future-of-search-commerce-and-payments-64fc8cc848f6>

“if it works it is not AI” → “its all AI”

Conscious killer robots to

WIRED Opinion

Elon Musk is wrong. The AI singularity won't kill us all

Elon Musk has stirred up fear, yet again, over the threat of killer AI. But he's missing the point completely, argues professor Toby Walsh

And don't just take my word for it. A recent [survey of 50 Nobel Laureates](#) ranked the climate, population rise, nuclear war, disease, selfishness, ignorance, terrorism, fundamentalism, and Trump as bigger threats to humanity than AI.

**Jobs TAKING OVER:
 World into unemp
 Millions Of
 desp
 Will Create More Jobs**


Artificial Intelligence Is Creating New And Unconventional Career Paths

white-
 for
 By 2020, Artificial
 Than It Eliminates: Rep
 Jobs at
 Will Eliminate
 m AI —

Artificial intelligence can perform certain specific tasks just as well as humans, but it still has a long way to go before it can replace humans



- By 2050, develop a team of fully autonomous humanoid robots that can win against the human world champion team in soccer.



The Definition of AI

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?
 - replicating human behavior early hallmark of intelligence
- Are humans always intelligent?
- 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 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?

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 Cognitive Science 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

Examples: Formal Cognitive 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

Examples: Perceptual Tasks

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

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
- ...

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

- Logic vs. Probabilistic vs. Neural
 - In 1950s, logic dominates
 - attempts to extend logic
 - 1988 – Bayesian networks
 - efficient computational framework
 - 2013 – deep neural networks
 - powerful representation across modalities

Topics of this Course

- Phase 1: Search, Constraint Satisfaction, Logic, Games
- Phase 2: Uncertainty (decision theory, probabilistic knowledge representation), Learning (reinforcement)
- Phase 3: Deep Neural Networks