

## CS210: Artificial Intelligence

Dr. Chandra Prakash

Assistant Professor

Department of Computer Science and Engineering



(Slides adapted from Stuart J. Russell, B Ravindran, Mausam, Prof. Pallab Dasgupta, Prof. Partha Pratim Chakrabarti, Saikishor Jangiti, Shivaram Kalyanakrishnan)

1

## Course Objectives

- A brief idea to the philosophy and breadth of ideas in AI
- Basic ideas and techniques underlying the design of intelligent computer systems / agents.
- Learn the representation and use of knowledge in inference-based problem solving approaches
- Learn to apply probability theory to describe and model agents operating in uncertain environments
- Learn statistical and decision-theoretic modeling paradigm.
- Presentation practice



## Syllabus







- **Module 1:** Introduction to AI
- **Module 2:** Automated Problem Solving
- **Module 3:** Logic and Deduction
- **Module 4:** Planning in AI
- **Module 5:** Reasoning Under Uncertainty
- **Module 6:** Machine Learning
- **Module 7:** AI Applications

6

## Reading Material

For students who want to read more we recommend



- On line relevant material will be shared with you

7



## Course website

- **Website :**
  - tentative schedule update
  - lecture slides and notes
  - course policies, etc.
- <https://cprakash86.wordpress.com/csb-210-ai-2024/>
- **Discussion (lecture related, doubts)**
  - [cprakash@coed.svnit.ac.in](mailto:cprakash@coed.svnit.ac.in)
  - Teaching Assistant [TA]:

8



## Course Policies

- Lecture notes, programming assignments, and other useful information will be posted on the course web page.
- You should check the web page regularly.
- The lecture notes will be in power point.
- Discussion of the programming assignments is allowed and encouraged. However, each team is expected to do its own work.
- **Assignments which are similar will receive a zero.**
- Regular attendance is highly recommended. If you miss a class, you are responsible for all material covered or assigned in class. Late programming assignments will be penalized 10% of the points assigned per day (weekends count as one day).

9



## Course Components

- **Continuous Assessments [ 20 ] :**
  - Attendance
  - Surprise Quizzes
  - projects, groups of 4-5
    - Academic integrity!
    - Python
    - Give you hands-on experience with the algorithms
- **Mid Term [30]**
- **End Term [50]**
- **LAB:**
  - Programming assignments


10



## About 'the' Course

- An assignment based course
- More emphasis on developing an solution for a real time problems
- Peer learning through presentation
- Project:
  - Using Programming in Prolog and Python
- Pre-requisites
  - Data Structures
  - Probability

Which of these subjects you would like to learn




11


## Form Analysis




Where do you see yourself in the future, career wise?  
96 responses




How much you are comfortable with programming [ Python preferably ]  
96 responses



How would you prefer us to assess your learning?  
96 responses



Beside assignment what is your preferred means of learning how to do something ?  
96 responses



## Your Expectation

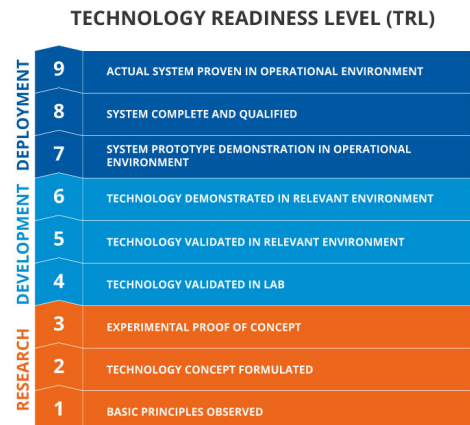


- Hope this course teaches a lot about the field of AI both theoretically and practically, so that by the end of this course we are at least able to make some mini projects on our own
- I should be able to develop ai solutions
- New learnings with innovations
- To have a good knowledge at ML / Able to build own machine learning model.
- Notes+worksheets+real life applications examples
- To provide course material or reference books for a particular unit after it has been taught
- Would love to have teacher student interactions, access to ask doubts, and receive notes, ppts or pdfs
- Team quizzes and assignments where groups of students work together on problems or even assignments building coordination and understanding amongst each other.
- Study Reinforcement Learning, a key aspect of AI involving decision-making and optimization.
- I do not want to attend boring lecture mean by completion of syllabus but it should be more about concepts.

## Project : Technology Readiness Levels (TRLs)



- Technology Readiness Levels (TRLs) are a method for understanding the technical maturity of a technology during its acquisition phase.
- TRLs allow engineers to have a consistent datum of reference for understanding technology evolution, regardless of their technical background.



## How to keep motivated in this course:



Don't Search for Numbers,  
Search for your queries/Answers

“I would rather have questions that can't be answered than answers that can't be questioned.” — Richard P. Feynman

## Module 1: Introduction to the Computers




- PART 1.1: What is Artificial Intelligence
- PART 1.2: History of AI
- PART 1.3: Possible Approaches in AI
- PART 1.4 : Application Domains and brief Overview of Modern AI
- PART 1.5: Areas Contributing to AI
- PART 1.6 : Core Capabilities covered in this course

16

## Human Sapiens



- Latin-
  - *wise man*
- Trying to understand How we
  - think
  - act
- Are humans the only intelligent species?
- Do we include all living beings as intelligent?
  - There do not exist standard and mathematically precise definitions of intelligence.




17

## Artificial Intelligence (AI)



- ?????
- Artificial = Machine
- But what is intelligence?
- AI is the **an attempt of** reproduction of human re intelligent behavior by computational methods



## Imagination / Reality



From the Mahabharata[1]



Modern-day videoconferencing[2]

19



## Sci-Fi AI?



20



## Today's Robot



21



## Lets Start



2403343781289312  
 + 2843033712837981  
 + 2362142787897881  
 + 3256541312323213  
 + 9864479802118978  
 + 8976677987987897  
 + 8981257890087988  
 = ?


We are naturally good at certain things.

○ Pattern recognition. ○ Languages and speech. ○ Reasoning. ○ Planning. ○ Learning. ○ Exploring.




## Learning / use Brain?

- Brains (human minds) are very **good** at making rational decisions, but **not perfect**
- Brains **aren't as modular** as software, so hard to reverse engineer!
- “Brains are to intelligence as wings are to flight”
- Lessons learned from the brain: **memory** and **simulation** are key to decision making



23

## Human Cognition Abilities





24

## What is Intelligence ?



- **Intelligence** (Oxford dictionary):
- Ability to
  - Learn
  - Understand and
  - Think.



25

## What is involved in INTELLIGENCE



- **Ability to interact with the real world**
  - to perceive, understand, and act
  - e.g., speech recognition and understanding and synthesis
  - e.g., image understanding
  - e.g., ability to take actions, have an effect
- **Reasoning and Planning**
  - modeling the external world, given input
  - solving new problems, planning, and making decisions
  - ability to deal with unexpected problems, uncertainties
- **Learning and Adaptation**
  - we are continuously learning and adapting
  - our internal models are always being "updated"
    - e.g., a baby learning to categorize and recognize animals




26

## Artificial + Intelligence



- **Artificial :**
  - fake, not real , man made
  - aeroplane= artificial flying
  - ships= artificial swimming
- **Intelligence:**
  - "the capacity to learn and solve problems"
  - in particular,
    - the ability to solve novel problems
    - the ability to act rationally
    - the ability to act like humans



27



## Computer Systems

- Common in almost all aspects of our daily lives.
  - Hard to imagine a world without them.



28



## What is “Artificial Intelligence”?



– Alan Turing [ 1950]

- I proposed to consider the question:  
**can machine think ?**

- [A. M. Turing (1950) **Computing Machinery and Intelligence**. *Mind* 49: 433-460.]

- Turing Test & **Total turing test** [1950]

- Operational test to determine an entity is intelligent / not
  - The computer is interrogated by a human via a teletype.
  - It passes if the human cannot tell if there is a computer or human at the other end



29



## 1955-56 Birth of the word AI

- Term coined by, John McCarthy (1955)
  - AI – “the science and engineering of making intelligent machines”
- **Dartmouth Summer Research Project** on Artificial Intelligence (1956)

On September 2, 1955, the project was formally proposed by McCarthy, Marvin Minsky, Nathaniel Rochester and Claude Shannon. The proposal is credited with introducing the term ‘artificial intelligence’.

The Proposal states<sup>[7]</sup>

“ 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. The study is to proceed on the basis of the conjecture that every aspect of learning or any other feature of intelligence can in principle be so precisely described that a machine can be made to simulate it. An attempt will be made to find how to make machines use language, form abstractions and concepts, solve kinds of problems now reserved for humans, and improve themselves. We think that a significant advance can be made in one or more of these problems if a carefully selected group of scientists work on it together for a summer.”

[https://en.wikipedia.org/wiki/Dartmouth\\_workshop](https://en.wikipedia.org/wiki/Dartmouth_workshop) [01 June, 2019]

Larger Intent, Dream, Overconfidence ...




30



## Some Early successes of Dartmouth

- Many key projects were initiated after dartmouth summer project.
  - a) Shakey robot [1966 - 1972]
    - Combined research in robotics, computer vision and natural language processing
    - First mobile robot to perceive environment
    - Could reason about its surroundings and actions!
      - Introduced **A\* algorithm** to find paths
      - **Hough Transform** for image analysis
      - Used **Lisp** for programming
      - **visibility graph** used for finding shortest paths in the presence of obstacles.



Source : <http://www.ai.sri.com/shakey/>

31

## Some Early successes of Dartmouth



- b) Dendral
  - attempted to encode the domain expertise in molecular biology as an **expert system**
  - determining 3D structures of complex chemical compounds
  - Led to the creation of expert systems for various other domain, including medical.
- A milestone in the history of AI !!!

32

## ARTIFICIAL INTELLIGENCE



There are **no clear agreement on the definition** of AI

- It is the science and engineering of making intelligent machines, especially intelligent computer programs.
- It is related to the similar task of using computers to understand human intelligence, but AI does not have to confine itself to methods that are biologically observable.
- *AI is the study of how to make computers just like humans. That means how to make computers to do things that people do better.*

33

## Other possible AI definitions



- AI is a collection of **hard problems** which can be solved by humans and other living things, but for which we don't have good algorithms for solving.
  - e. g., understanding spoken natural language, medical diagnosis, circuit design, learning, self-adaptation, reasoning, chess playing, proving math theories, etc.
- AI is a process of making a machine or a program that
  - Learn and understand like human
  - Acts like human (Turing test)
  - Thinks like human (human-like patterns of thinking steps)
  - Acts or thinks rationally (logically, correctly)

34

## What is AI?



The science of making machines that:

35





## Rational Decisions

We'll use the term rational in a very specific, technical way:


- **Rational:** maximally achieving pre-defined goals
- Rationality only concerns what decisions are made (not the thought process behind them)
- Goals are expressed in terms of the utility of outcomes
- Being rational means maximizing your expected utility



36



## Maximize Your Expected Utility





37



## Designing Rational Agents

- An **agent** is an entity that *perceives* and *acts*.
- A **rational agent** selects actions that maximize its (expected) **utility**.
- An agent is a function from percept histories to actions:  
[f: P\* → A]
- Characteristics of the **percepts**, **environment**, and **action space** dictate techniques for selecting rational actions
- **This course** is about:
  - General AI techniques for a variety of problem types
  - Learning to recognize when and how a new problem can be solved with an existing technique



38



## Cont...

AI is the study and design of intelligent agents

where,

an intelligent agent is a system that interact with its environment and takes actions that maximize its chances of success.

39



## Problems In AI

### Easy Problems in AI

- It's been **easier to mechanize** many of the **high level cognitive tasks** we usually associate with "intelligence" in people
  - e. g., **symbolic integration**, **proving theorems**, **playing chess**, some aspect of **medical diagnosis**, **Engineering tasks**, **Financial**, etc.



### Hard Problems in AI

- can be solved by humans and other living things, but for which we don't have good algorithms for solving
- It's been very **hard to mechanize tasks** that animals can do easily
  - walking around without running into things (**ASIMO**)
  - Perceptual task** :
    - interpreting complex sensory information (visual, aural, ...)
    - working as a team (ants, bees)
- Algorithmic view** such as NP hard , Search , Game playing, planning

40



## Humanoid robot



41



## AI pioneers

- Alan Turing (1912-1954)**
  - Father of computer science
  - Turing test for AI
- Marvin Minsky (MIT) – 1956**
  - Built first Neural network computer SNARC
- John McCarthy (Stanford University)**
  - Developed LISP, AI programming language
- 2018 Turning Award for Deep Learning**
  - Jeffrey Hinton, Yoshua Bengio, Yann LeCun




most prestigious technical award is given for major contributions of lasting importance to computing

42





## Foundations of AI

- 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/Statistics** : modeling uncertainty, learning from data
- Economics** : utility, decision theory, rational economic agents
- Neuroscience** : neurons as information processing units.
- Psychology / Cognitive Science** : how do people behave, perceive, process cognitive information, represent knowledge.
- Computer** : building fast computers engineering, GPU, TPU, *Wafer scale engine (WSE)* , Quantum computing
- Control theory**: design systems that maximize an objective function over time
- Linguistics** : knowledge representation, grammars



43

## Confusion [ AI, ML, DL ] ???

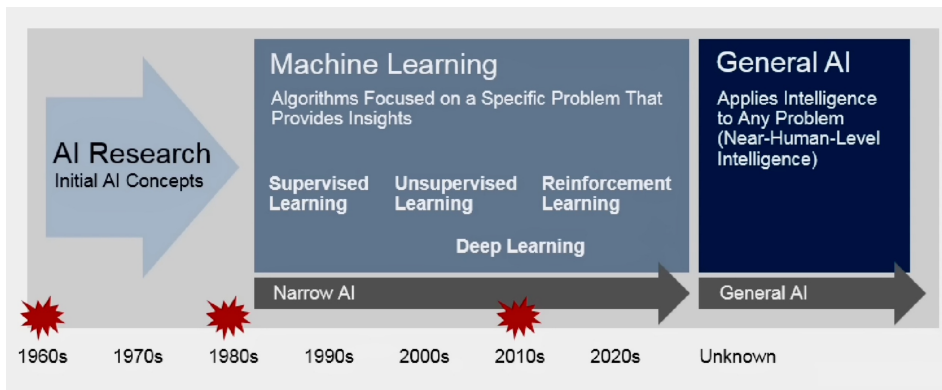


44

## Evolution of AI



## A brief history of AI



46

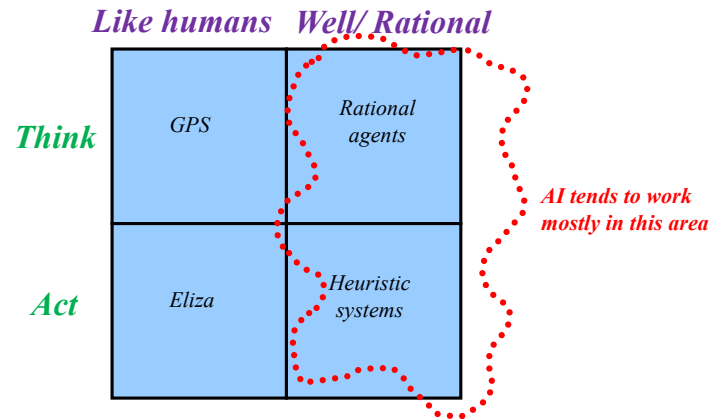
## Programming languages for AI



- The programs for AI problems can be written with on procedural languages like PASCAL or declaration languages like **PROLOG**.
- Generally relational languages like PROLOG or LISP are preferred for symbolic reasoning in AI.
- If the program requires much arithmetic computation (say for the purpose of uncertainty management), then procedural languages would be preferred.
  - Python.
  - Prolog.
  - LISP.
  - R.
  - C++
  - JavaScript.
  - Java.
  - Haskell.

47

## Dimensions / Possible Approaches in AI




48

## Thinking Humanly



- Cognitive Modelling Approach
  - very hard to understand how humans think
  - Requires scientific theories of internal activities of the human brain
  - How do we capture human thinking to implement ?
    - Cognitive Science, Neuroscience, Psychology
- Computational model should reflect “how” results were obtained.
- System : “**General Problem Solver (GPS)**” (Newell and Simon, 1961)
  - Designed to work as a universal problem solver
  - Problems represented by horn clauses
  - First AI Machine which has KB + Inference separation
  - Goal **produce a sequence of steps of the reasoning process** that was similar to the steps followed by a person in solving the same task.
- Growth of Cognitive science and AI supports each other
- A machine that thinks like human while solving a problem correctly





49

## Thinking well / rationally



- Law of thoughts Approach
  - Greek Philosopher Aristotle, Third century BC. – right thinking
  - Belief that “logic” governs the human thought process
- Develop formal models of knowledge representation, reasoning, learning, memory, problem solving, that can be result in **algorithms**.
  - There is often an emphasis on a systems that are provably correct, and guarantee finding an **optimal solution**.
- Thinking Rationally
  - Field of Logics gave rise to codifying rational thinking
  - When elements are “things”, we reason about things
- **Hurdles to the idea :**
  - Not everything can be logically coded : eg Hot Fire , reflex
  - No provably correct action at a moment
  - Exhaustive computational resources




50

## Acting Well / Rationally



- Acting Rationally- The Rational Agent Approach
  - rational behavior = **doing the right thing**
  - 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\* → A]
- Design best program for given machine resources
  - For a given set of inputs, generate an appropriate output that is not necessarily correct but gets the job done.
- A **heuristic** (heuristic rule/method) is a rule of thumb, strategy, trick, simplification, or any other kind of device which drastically limits search for solutions in large problem spaces.
  - Heuristics do not guarantee optimal solutions; in fact, they do not guarantee any solution at all
  - All that can be said for a useful heuristic is that it **offers solutions which are good enough most of the time**.



51



## Act like humans

- Behaviorist approach.
- Not interested in how you get results, just the similarity to what human results are.
- Exemplified by the Turing Test (Alan Turing, 1950).

|       | Like humans | Well              |
|-------|-------------|-------------------|
| Think | GPS         | Rational agents   |
| Act   | Eliza       | Heuristic systems |

52



## Eliza

- Joseph Weizenbaum, 1964
- ELIZA: A program that simulated a psychotherapist interacting with a patient and successfully passed the Turing Test.
- Coded at MIT during 1964-1966 by Joel Weizenbaum.
- Natural Language Processing Computer Program
- First Chatbot!
- First script was DOCTOR as Psychotherapist
  - The script was a simple collection of syntactic patterns not unlike regular expressions
  - Each pattern had an associated reply which might include bits of the input (after simple transformations (my → your))




Source : <https://news.mit.edu/2008/obit-weizenbaum-0310>

53



## Acting Humanly: The Turing Test

- Alan Turing (1912-1954)



- The interrogator is limited to using the responses to written questions in order to make the determination.
- **Skills necessary to pass these tests**
  - NLP, Knowledge Representation, Automated Reasoning, ML
  - +
  - Computer Vision & Robotics(for total turing test)

54



## The Loebner Contest




- A modern version of the Turing Test, held annually, with a \$100,000 cash prize.
  - [ <http://www.loebner.net/Prizef/loebner-prize.html> ]
- Restricted topic (removed in 1995) and limited time.
- Participants include a set of humans and a set of computers and a set of judges.
- Scoring
  - Rank from least human to most human.
  - Highest median rank wins \$25000.
  - If better than a human, win \$100,000. (Nobody yet...)

56



## Chinese Room argument (1980)

- Devised by John Searle
- Argument against the **possibility of true artificial intelligence**.
- The argument centers on a thought experiment in which someone who knows only English sits alone in a room following English instructions for manipulating strings of Chinese characters, such that to those outside the room it appears as if someone in the room understands Chinese.



57



## Cont...

- Some problems used to be thought of as AI but are now considered not
  - e. g., compiling Fortran (suited to numeric computation and scientific computing) in 1955,
  - symbolic mathematics (manipulate mathematical equations ) in 1965
  - proving math theories

58



## Some Revised definitions of artificial intelligence


|   |   |
|---|---|
| <p><b>Thinking Humanly</b></p> <p>“The exciting new effort to make computers think ... <i>machines with minds</i>, 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><b>Thinking Rationally</b></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><b>Acting Humanly</b></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><b>Acting Rationally</b></p> <p>“Computational Intelligence is the study of the design of intelligent agents.” (Poole <i>et al.</i>, 1998)</p> <p>“AI ... is concerned with intelligent behavior in artifacts.” (Nilsson, 1998)</p>                      |

59



## FOCUS AREAS FOR AI INTERVENTION in India?

- Preventive and affordable Healthcare
- Education and Skilling
- Agriculture and Rural Development
- Smart Mobility and Intelligent Transportation Systems
- Retail
- Manufacturing
- Energy management
- Smart Cities and Infrastructure




Source: Niti Aayog Discussion Paper on AI, June 2018

60

# Intelligent Systems in Your Everyday Life



- Natural Language
  - Speech technologies (e.g. Siri) – Automatic speech recognition (ASR)
    - Text-to-speech synthesis (TTS)
    - Dialog systems
  - Language processing technologies
    - Question answering : 2011 IBM's Watson
    - Machine translation

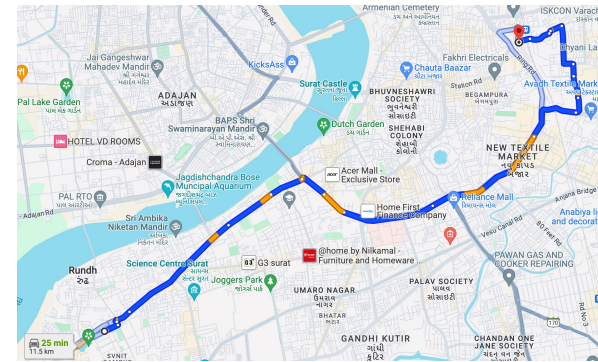


- Web search
- Text classification, spam filtering, etc...



61

# Maps and Self-driving cars



62

# Smart Systems



Google Now



Siri




Cortana

Google Now



Siri



Cortana



63

# Robot



Boston Dynam



64




## Business Intelligence

- Fraud Detection
- News generation



## Bionomics



66



## Sophia : Social humanoid robot

- Activated in 2015, Sophia is known for her human-like appearance and behavior.
  - Developed by Hong Kong-based company Hanson Robotics.
- In **October 2017**, the robot became a Saudi Arabian citizen, the first robot to receive citizenship of any country.
- Sophia uses artificial intelligence, facial and voice recognition, and visual data processing, is able to imitate human facial expressions and gestures, and can make conversation on predefined topics.



World's first robot citizen's India debut at IIT Bombay



## How is the AI of this Era is Different?





## Other Modern AI Factors

### Machine Learning

- The main driver of recent successes in AI
- Move from "code" to "data"




- Continued expansion of open source AI, especially in Python, aiding machine learning and big data ecosystems.
- Leading **deep learning libraries open-sourced**, allowing further adoption by industry.
- Open sourcing of large datasets of millions of labeled images, text datasets such as Wikipedia has also driven breakthroughs.

69



## Vision and Deep Learning

ImageNet Large Scale Visual Recognition Challenge (ILSVRC) winners



## Applications



### Deep Learning/AI APPLICATIONS

Image Classification  
Object Detection

**COMPUTER VISION**

Voice Recognition  
Language Translation

**SPEECH & AUDIO**

Recommendation Engines  
Sentiment Analysis

**NATURAL LANGUAGE PROCESSING**

*Few Popular Applications: Precision Agriculture, Learner Profiling, Video Captioning, Exploring Patterns from Satellite images, Image detection in Healthcare, Identifying specific markers in Genomes, Creating Art and Music, Recommendations, behavior prediction,*




## Computer Vision

**Face Recognition**


Image quality enhancement  
Beautification

**Gesture Recognition**

Classification

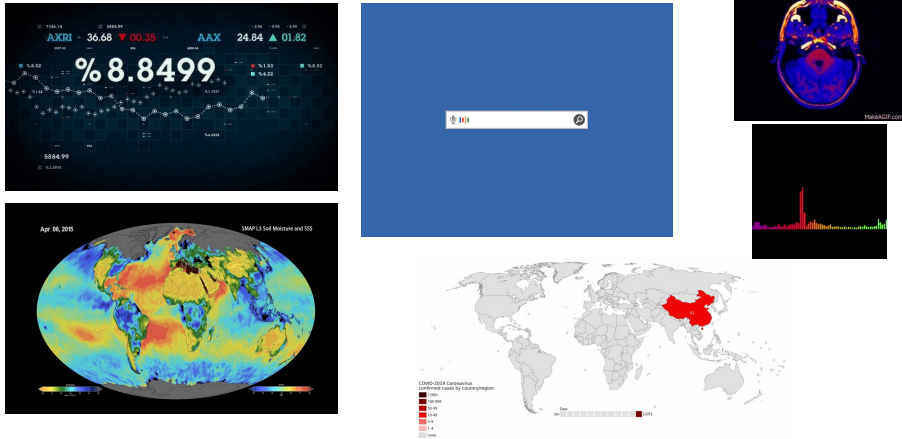


Object detection (Self driving car)



72

# Sequence Modeling Algorithms




# Game Agents



Pong



Enduro



Beamrider



# Game Agents



- Classic Moment: May, '97: Deep Blue vs. Kasparov
  - First match won against world champion
  - "Intelligent creative" play
  - 200 million board positions per second
  - Humans understood 99.9 of Deep Blue's moves
  - Can do about the same now with a PC cluster



- 1996: Kasparov Beats Deep Blue  
"I could feel --- I could smell --- a new kind of intelligence across the table."

- 1997: Deep Blue Beats Kasparov  
"Deep Blue hasn't proven anything."




Text from Bart Selman, image from IBM's Deep Blue pages

# AlphaGo



- AlphaGo is the first computer program to defeat a professional human Go player
- At the opening move in Chess there are 20 possible moves. In Go the first player has 361 possible moves
- Policy network -selects the next move to play.
- In late 2017, AlphaZero, a single system that taught itself from scratch how to master the games of chess, shogi, and Go, beating a world-champion program in each case.
- AlphaGo must restrict Breath and Depth of search among all board configurations with heuristics information supplied by training and winning policy for max reward.








## Motion Analysis



Deformable Part Model  
Felzenszwalb, McAllester, Ramanan, 2009



Histogram of Gradients (HoG)  
Dalal & Triggs, 2005



Move: Walk Around challenge, one of the official challenges in the NeurIPS 2019

**How to take care of seniors while keeping them safe?**

- Early Symptom Detection of COVID-19
- Monitor Patients with Mild Symptoms
- Manage Chronic Conditions

Versatile Scalable

Mobility Infection Live Feed Burn-Free  
Sleep Diet



## Robots: Human-AI Interaction



## Robotics

- Robotics
  - Part mech. eng.
  - Part AI
  - Reality much harder than simulations!
- Technologies
  - Vehicles
    - 132 miles DARPA Grand challenge
  - Rescue
  - Help in the home
  - Lots of automation...
- In this class:
  - We ignore mechanical aspects
  - Methods for control



Images from UC Berkeley, Boston Dynamics, RoboCup, Google




## What can't AI systems do yet?



## What Can AI Do?

Quiz: Which of the following can be done at present?

- ✓ Play a decent game of KBC?
- ✓ Win against any human at chess?
- ✓ Win against the best humans at Go?
- ✓ Play a decent game of tennis?
- ✓ Grab a particular cup and put it on a shelf?
- ✗ Unload any dishwasher in any home?
- ⊕ Drive safely along the highway?
- ✗ Drive safely along Narela market ?
- ✓ Buy a week's worth of groceries on the web?
- ✗ Buy a week's worth of groceries at D-Mart?
- ⊕ Discover and prove a new mathematical theorem?
- ✗ Perform a surgical operation?
- ✗ Unload a know dishwasher in collaboration with a person?
- ✓ Translate spoken Chinese into spoken English in real time?
- ✓ Write an intentionally funny story?



81




## What can't AI systems do yet?

- Common Sense Problem-Solving and Reasoning
- Emotional Intelligence
- Creativity and Innovation
- Generalization in Unseen Scenarios
- Ethical Decision-Making
- Understanding Natural Language at Human Level
- Dynamic Learning and Adaptation
- Cross-Domain Learning

82



## RoboCup 2050 challenge,

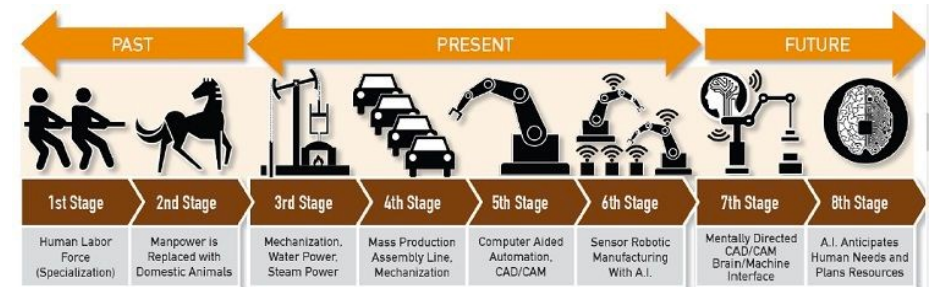


- RoboCup and Its Role in the History and Future of AI
  - Source : [<https://ai.sony/blog-006/> ]

83



## Stages of Industrial Revolution



Source : <https://www.textiletoday.com.bd/industry-4-0-a-typical-discussion-and-recommendation-for-bangladesh-spinning-industry/>

84



## AI Good or Bad



## AI Good or Bad ??

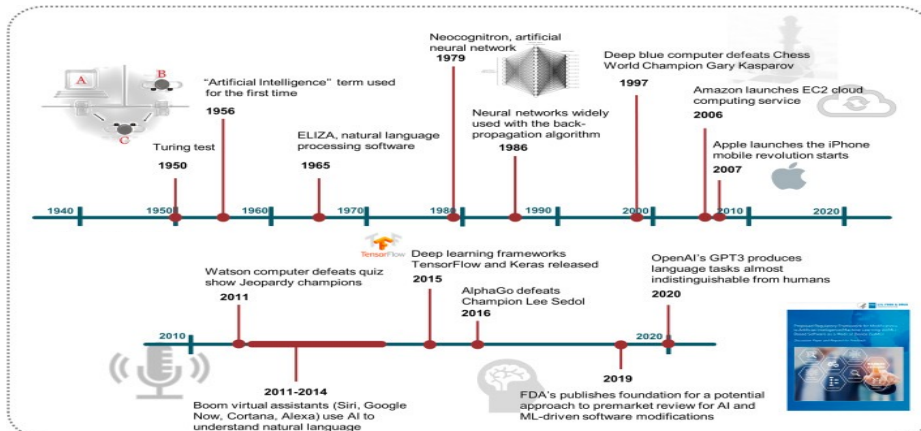


- Possible Risks :
  - Lethal Autonomous weapons
  - Surveillance and persuasion
  - Biased decision making
  - Impact on employment
  - Cybersecurity
- Robots waking up and deciding to revolt against humanity
  - Is it possible ???
  - Never happen if we give them the right objective, because there is no such thing as a robot deciding to change its objective
  - how to give them the right objective to begin with
- Colobration between AI and Human

86



## History of AI in one shot



## AI Problem Areas /Tasks



- **1<sup>st</sup> Generation of AI :** Fomal cognitive Tasks
  - Game
    - Tic-Tac-Toe
    - Chess
    - Checkers
    - Go
  - Mathematics
    - Logic
    - Geometry
    - Calculus
    - Proving properties of programs
- **2<sup>nd</sup> Generation : Expert Tasks**
  - Knowledge Reprेशन
    - Design
    - fault finding
    - Manufacturing planning
  - Medical
    - Diagnosis
    - Medical Image Analysis
  - Financial
    - Stock market prdeictions
- **3<sup>rd</sup> Generation of AI :** Perceptual Tasks
  - Perception
    - Vision
    - Speech
  - Natual Language
    - Understanding
    - Generation
    - Translation
  - Robot Control

89



## Course Plan


- A brief idea to the philosophy of AI
- A brief idea to 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

90


## Bridging the gap



Modeling



Inference



- **Theory vs. Modeling vs. Algorithm vs. Applications**
  - Lecture focused towards modeling
  - Assignment tilted towards applications
  - Few theorems

91



we can not cover everything

92

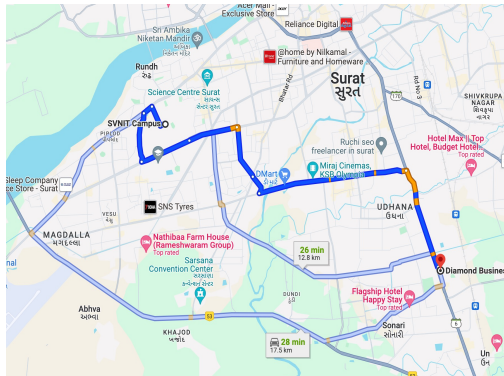
## Core Capabilities: Covered in this course

- The ability to solve problems
  - Searching Algo, Constraint satisfaction, Optimization
- The ability to plan
  - Abstraction
- The ability to deduce
  - Logic, Reasoning algorithms
- The ability to learn
  - Models, Data, Learning algorithms
- The ability to handle uncertainty
  - Probabilistic and Neural
- The ability to interface with the real world




93

## Constraints and Optimization



- Path Finding
  - I wish to find a shortest path
  - I wish to find a path with minimum congestion
  - I wish to find a path with combination of transportation options (metro, bus, taxi)
  - I wish to find a path which goes past a medicine shop
  - I wish to find a path which minimizes energy consumption from my battery in a e-vehicle
- When the size and complexity becomes too big we use “heuristic functions” to cut out unnecessary parts.
- In the lack of domain knowledge, we can statistically learn the best way (reinforcement learning) by exploration.
- Modern AI aims to combine learning from data with structured use of domain knowledge. 94

## AI Planning





- Elements of a Planning Problem
  - A set of **states** (worlds) described in terms of predicates
  - A set of **actions** which transforms some parts of one world to take us to another world
  - An **initial world**
  - A **goal** in terms of the predicates that must hold in the final world
- Planning is widely used in robotics and automated control
- Modern AI explores techniques that combine planning with machine learning
  - Autonomous driving is one of many areas where such combinations are highly relevant

95

## Logical Reasoning



- Automated ways to use what is known to reason about something which is not explicitly known.
- Automated Reasoning:
  - Deduction
    - Rule: All the marbles in this bag are blue
    - Case: These marbles are from this bag
    - Inference: These marbles are blue
  - Abduction
    - Rule: All the marbles in this bag are blue
    - Observation: These marbles are blue
    - Case: These marbles are from this bag
  - Induction
    - Case: These marbles are from this bag
    - Observation: These marbles are blue
    - Rule: All the marbles in this bag are blue



- Five color theorem: All maps can be colored with five colors, where neighboring countries get different colors [Proved in 1800s]
- Four color theorem: All maps can be colored with four colors, where neighboring countries get different colors [Proved in 1976 with help of computers]
- Applications of Logical Reasoning
  - Automated Theorem Proving
  - Rule-based Systems
  - Complexity Analysis


96

## Logical Reasoning with Other Fun Things



- Facts
  - grandfather, grandmother,
  - maternalgrandfather, maternalgrandmother ,
- Query :
  - maternalgranduncle

- Who is the maternal great uncle of Freya?



*We need that a social media platform to suggests Freya to post a picture of Fergus on the Maternal-Great-Uncle day*

97


## Logical Reasoning with Other Fun Things



- grandfather, grandmother, maternalgrandfather, maternalgrandmother, maternalgranduncle

- $father(x, z), father(z, y) \Rightarrow grandfather(x, y)$
- $father(x, z), mother(z, y) \Rightarrow maternalgrandfather(x, y)$
- $mother(x, z), father(z, y) \Rightarrow grandmother(x, y)$
- $mother(x, z), mother(z, y) \Rightarrow maternalgrandmother(x, y)$
- $maternalgrandmother(x, z), mother(z, p), son(p, y) \Rightarrow maternalgreatuncle(x, y)$

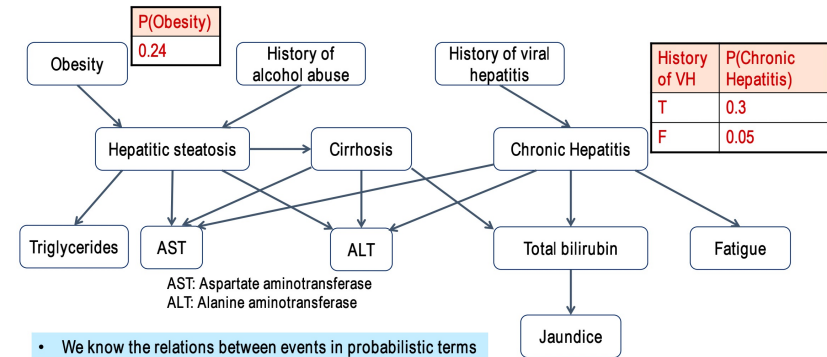
- Who is the maternal great uncle of Freya?



- maternalgrandmother(Freya, Charlotte), mother(Charlotte, Lindsey), son(Lindsey, Fergus)  $\Rightarrow$  maternalgreatuncle(Freya, Fergus)

98

## Reasoning under Uncertainty

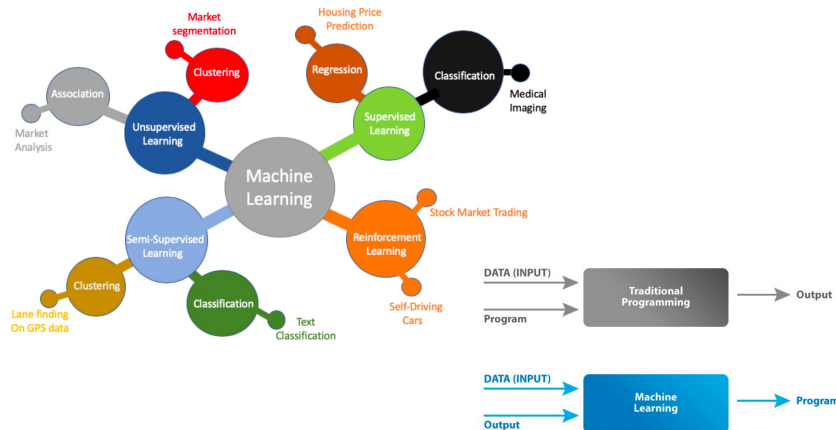


- We know the relations between events in probabilistic terms
- We are given a set of known events
- Goal: Find the probability of some other event

Challenge: Sometimes we do not know which is the cause and which is the effect.

99

## Machine Learning



## Course Plan



- A brief idea to the philosophy of AI
- A brief idea to breadth of ideas in AI
- General computer scientist
  - general tools to aid in tackling a new problem
- Serious AI enthusiast
  - a primer from which to launch advanced study
- Theory vs. Modeling vs. Algorithm vs. Applications**
  - Lecture focused towards modeling
  - Assignment tilted towards applications
  - Few theorems

102





## Takeaways

- Artificial Intelligence is a very broad and flexible concept.
- Learning AI at a practical level is about methods associated with AI goals.
- Arguably, everything a computer does is AI, at a conceptual level. Traditionally, AI goals have centered around definitions of intelligence that go beyond “useful.”
- (Arguably, all or almost all statements are arguable.)

103



## Conclusion

- Definition : Intelligence and Artificial Intelligence (AI )
  - AI Paradox: Once we understand how X works, X is no longer AI!
- History of AI : Turing Test, Chines Room Argument
- Approches in AI
- Application and domains and overview of Modern AI
- Core capabilities of AI
- Robots waking up and deciding to revolt against humanity
- Colobration between AI and Human
- The subject of AI deals more with symbolic reasoning that conventional number crusting problems.
- Common areas covered under AI
  - Knowledge representation, learning, speech and uncertainty management of data and knowledge.
- Python and PROLOG are the used for programming AI problems.



## Module 2: Automated Problem Solving

- PART 2.1: Intelligent Agent & Environment
- PART 2.2: Problem solving Agent
- PART 2.3: Problem Solving Methods
- PART 2.4: Search Strategies
- PART 2.5: Adversarial Search
- PART 2.6: Constraint Satisfaction Problems

105



## Home work 1

### Due date : 20-Jan

1. Read the Niti Aayog Discussion Paper on AI, June 2018. Find the key areas for AI intervention in India. Consider one area and suggest how would you contribute in the selected area (at least 2 pages ) .
2. Read Turing's original paper on AI (Turing, 1950). In the paper, he discusses several potential objections to his proposed enterprise and his test for intelligence.
  - a) Discuss the "Heads in the Sand" Objection mention in his article.
  - b) According to you which objections still carry some importance?
  - c) Can you think of new objections arising from developments since he wrote the paper for the next 50 years ?
3. Are reflex actions (such as flinching from a hot fire) rational? Are they intelligent? Justify your answer.

**For more questions refer Course webpage**

106

## References

---



- Artificial Intelligence – A Modern Approach, Stuart J Russell, Peter Norvig, Pearson Education India
- Slides adapted from CS188 Instructor: Anca Dragan, University of California, Berkeley
- Slides adapted from CS60045 ARTIFICIAL INTELLIGENCE