ARTIFICIAL INTELLIGENCE & MACHINE LEARNING

UNIT I

Learning Objectives

AI Fundamentals:

1. Defining Artificial Intelligence
2. Types of AI
3. Turing test
4. Defining AI techniques
5. Comparison - AI, ML and Deep Learning
6. Problem characterization and reduction
7. Defining State Space
8. AI Search Techniques (Hill Climbing, Breadth first and depth first search, Best first search, A*, AO*, Constraint Satisfaction)
9. MEA
10. Ethics of AI

Artificial Intelligence

- Artificial Intelligence (AI) is a **set of technologies** that enable computers to perform a variety of advanced functions, including the ability to see, understand and translate spoken and written language, analyze data, make recommendations, and more.
- Alan Turing (Father of Computer Science) asks the question

"Can Machine Think?"

- He Also offers a famous test "Turing Test" where a human interrogator would try to distinguish between computer and human text response.

Artificial Intelligence (contd.)

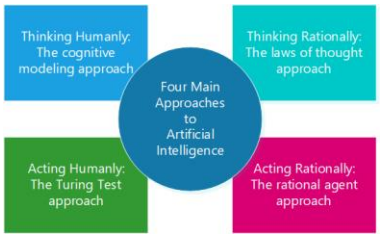

Types of AI

- Based on the **capability and functionality**, there are two types of AI:
  - Type-I
  - Type-II

Type-I	Type-II
• Narrow • General • Strong	• Reactive Machines • Limited Memory • Theory of mind • Self-awareness

Types of AI- Based on Capabilities

- Narrow/Weak AI:
  - Able to perform dedicated task with intelligence.
  - This is the most common AI.
  - It is **only trained for one specific task** only.
  - e.g. Apple Siri, self-driving cars
- General AI:
  - It can perform any intellectual task with efficiency like a human.
  - **No such system exist** till now which can be smarter and think like a human by its own.
- Super AI:
  - A level of intelligence of a system where the machine could surpass human intelligence and can perform task better than human.
  - **It can be outcome of general AI**.
  - Commute by its own, make judgement, solve the puzzle.

Types of AI- Based on Capabilities (contd.)

We are here

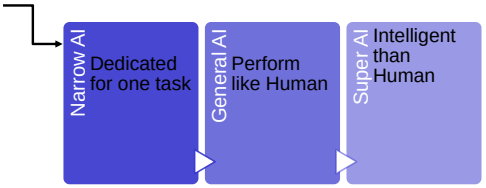

Types of AI- Based on Functionality

- Reactive Machines
  - Most basic types of AI
  - They **do not store memories or past experiences** for future actions.
  - Only focuses on current scenarios
  - e.g. IBM Deep Blue (a chess playing super computer)
- Limited Memory
  - This **can store past experience** for short period of time.
  - e.g. self driving cars→ can store speed and distance of nearby cars to navigate the roads.
- Theory of mind
  - This understand the human emotions, people, beliefs and **able to interact socially like humans**.
- Self-Awareness
  - **It is the future of AI**.
  - These machines will be super intelligent.
  - They will have their own sentiments and self awareness.

Turing Test

- Developed by **Alan Turing** in 1950.
- He proposed that whether or not a computer (machine) can think intelligently like humans?
- He engages both a human and a machine in a **text based conversation** and then decides which of the two they believe to be human.
- If the judge (human) unable to distinguish between the human and the machines based on the conversation, then the machine is said to have passed the Turing test.

This tests evaluate the machine intelligence for over six decades.

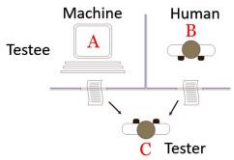

Turing Test today

- The Turing Test has its detractors, but it remains a measure of the success of artificial intelligence projects. An updated version of the Turing Test has more than one human judge interrogating and chatting with both subjects. The project is considered a success **if more than 30% of the judges**, after five minutes of conversation, conclude that the computer is a human.
- **The Loebner Prize** is an annual Turing Test competition that was launched in 1991 by Hugh Loebner, an American inventor and activist. Loebner created additional rules requiring the human and the computer program to have 25-minute conversations with each of four judges. The winner is the computer whose program receives the most votes and the highest ranking from the judges.
- In 2014, Kevin Warwick of the University of Reading organized a **Turing Test competition to mark the 60th anniversary of Alan Turing's death**. A computer chatbot called Eugene Goostman, who had the persona of a 13-year-old boy, technically passing the Turing Test in that event. He secured the votes of 33% of the judges who were convinced that he was human.
- In 2018, Google Duplex revealed the capability to performing tasks via the telephone. In various demonstrations, Duplex schedule a hair appointment as well as called a restaurant, with the human on the other end of the line not realizing they were interacting with a machine. However, critics point out that the interaction does not conform to the actual Turing test and claim the **test has still yet to be beaten by a machine.**

Limitations of Turing Test

- It focuses primarily on **language-based conversation** and not taking into account other important aspects such as problem-solving, perception, decision-making etc.
- The results of the test can be **influenced by the biases**, making it difficult to obtain objectives and reliable results.
- May **not be suitable to test for intelligence** as different computing systems are structured differently.

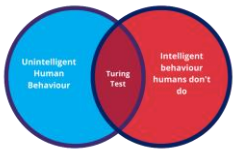

Defining AI Techniques

- AI techniques revolutionizing the way humans interact with technology.
- Depending on the **machine's ability** to utilize **past experiences** to anticipate future judgments, memory, and self-awareness, artificial intelligence can be classified into a variety of subcategories.
- Following are the four broad types of AI Techniques:
  - Machine Learning (ML)
  - Natural Language Processing (NLP)
  - Machine Vision/Computer Vision
  - Automation and Robotics

Machine Learning (ML)

- A subset of AI that uses **statistical methods** to enable machines to learn from data.
- It involves the creation of algorithms that can identify patterns, make predictions and improve their **performance over time without explicit programming**.

Supervised Learning
- The algorithm is trained on a labelled data set.
- Learn relationship between input-output.
- e.g. Linear Regression, Support Vector Machine

Unsupervised Learning
- The algorithm is trained on a unlabelled dataset, and it identified patterns, or structures in the data without guidance.
- e.g. K-means, PCA and associate rule learning.

Reinforcement Learning
- The algorithms learn from its actions and interactions with an environment to maximize a reward signal.
- e.g. Q-learning, Deep Q-network.

Natural Language Processing (NLP)

- A subfield of AI that focuses how machine understand, **interpret and generate human language**.
- e.g. Speech Recognition, Chatbots, Sentiment analysis
- It involves the use of computational techniques to process and analyze natural language data, such as text and speech, with the goal of understanding the meaning behind the language.
- Some commonly used techniques in NLP are:
  - Machine translation
  - Text classification
  - Part-of-speech tagging
  - tokenization
- **Challenge** : ambiguity and complexity of human language; similar word can have different meaning in multiple language.

Machine Vision/Computer Vision

- A field of AI that involves **training machines** to interpret and understand **visual data** from the world around us.
- Analyze visual data, such as images or videos and **identify patterns** and objects within them.
- Machine vision, paired with AI and deep learning, expands the **role of robots** in performing production-line tasks, such as picking, sorting, placing and performing a manufacturing line scan

Automation and Robotics

- Involves the use of technology to perform tasks that would otherwise be performed by humans.
- **Automation** is a process that uses software, machines, or other technology to carry out tasks in place of human workers.
- **Robotics** is a field that combines engineering and computer science to design and build robots to perform tasks.

Differences Between Automation and Robotics!!!

AI, ML and Deep Learning

- Deep learning and machine learning are **subsets of AI** wherein AI is the umbrella term. Each of these technologies can create smart applications. Companies can use machine learning, deep learning, and artificial intelligence for several projects.
- **Goal:**
  - AI: To make machines as smart as humans.
  - ML: To make machine learn without explicit programming
  - DL: To mimic the human brains and make accurate predictions
- **Data Handling:**
  - AI: deals with structured, semi-structured and unstructured data.
  - ML: Structured and semi-structured data
  - DL: Structured and unstructured data

AI, ML and Deep Learning (Contd.)

- **Types of algorithm**
  - AI: uses types of algorithm that reflects intelligence through decision-making
  - ML: employs an algorithm that enables the system to learn from data
  - DL: algorithm that includes a neural network to analyze data and outcomes.
- **Applications**
  - AI: ride sharing, auto pilot and online games.
  - ML: pattern recognition, fraud detection, content recommender system.
  - DL: stock market trading applications, medical diagnosis, customer support applications.

Problem characterization

- AI is a way of making a computer or machine or a software thinks intelligently in the similar manner the intelligent humans think.
- We must have a method to choose the **best possible solution**.
- To choose appropriate method for a particular problem we need to **categorize a problem**.

Problem characterization and Reduction (Contd.)

1. Decomposability of the problem into a set of independent smaller subproblems
2. Possibility of undoing solution steps, if they are found to be unwise
3. Predictability of the problem universe
4. Possibility of obtaining an obvious solution to a problem without comparison of all other possible solutions
5. Type of the solution: whether it is a state or a path to the goal state
6. Role of knowledge in problem solving
7. Nature of solution process: with or without interacting with the user

Problem characterization 1

1. Is the problem decomposable into small sub-problems which are then easy to solve.

Solving polynomial $2x^2 + x + 2$

- This problem can be solved by **breaking it into smaller problems**, each of which we can solve by using a small collection of specific rules.
- Using this technique of problem decomposition, we can solve very large problems very easily. This can be considered as an intelligent behaviour.

Problem characterization 2

2. Can Solution steps be ignored or undone?

Suppose we are trying to prove a mathematical theorem: first we proceed considering that proving a lemma will be useful. Later we realize that it is not at all useful. We start with another one to prove the theorem. Here we simply ignore the first method.

Consider the 8-puzzle problem to solve: we make a wrong move and realize that mistake. But here, the control strategy must keep track of all the moves, so that we can **backtrack to the initial state** and start with some new move.

Problem characterization 2 (Contd.)

2. Can Solution steps be ignored or undone?

Consider the problem of playing chess. Here, once we make a move we never recover from that step. These problems are illustrated in the three important classes of problems mentioned below:

1. Ignorable, in which solution steps can be ignored. e.g: Theorem Proving
2. Recoverable, in which solution steps can be undone. e.g: 8-Puzzle
3. Irrecoverable, in which solution steps cannot be undone. e.g: Chess

Problem characterization 3

3. Is the universe of the problem predictable?

- Consider the 8-Puzzle problem. Every time we make a move, we know exactly what will happen.
- This means that it **is possible to plan an entire sequence** of moves and be confident what the resulting state will be.
- We can **backtrack to earlier moves** if they prove unwise.
- Suppose we want to play Bridge. We need to plan before the first play, but we cannot play with certainty. So, the outcome of this game is very uncertain.
- In case of 8-Puzzle, the outcome is very certain. To solve uncertain outcome problems, we follow the process of plan revision as the plan is carried out and the necessary feedback is provided. The **disadvantage** is that the planning in this case is often very expensive.

Problem characterization 4

4. Is a good solution to the problem a absolute or relative?

Consider the problem of answering questions based on a database of simple facts such as the following:

1. X was a man.
2. X was a worker in a company.
3. X was born in 1905.
4. All men are mortal.
5. All workers in a factory died when there was an accident in 1952.
6. No mortal lives longer than 100 years

Suppose we ask a question: 'Is X alive?'

Problem characterization 4 (Contd.)

There are two ways to answer the question shown below:

Method I:

1. X was a man.
2. X was born in 1905.
3. All men are mortal.
4. Now it is 2008, so X's age is 103 years.
5. No mortal lives longer than 100 years.

Method II:

1. X is a worker in the company.
2. All workers in the company died in 1952.

Answer: So X is not alive. It is the answer from the above methods

Problem characterization 5

5. Is the solution to the problem a state or a path?

A path-solution problem can be reformulated as a state-solution problem by describing a state as a partial path to a solution. **The question is whether that is natural or not.**

Problem characterization 6

6. What is the role of knowledge in solving a problem using AI?

- Playing a chess: Consider again the problem of playing chess. Suppose you had unlimited computing power available. **How much knowledge would be required by a perfect program**?
- Reading Newspaper: Now consider the problem of scanning daily newspapers to decide which are supporting the Democrats and which are supporting the Republicans in some upcoming election. Again assuming unlimited computing power, **how much knowledge would be required by a computer trying to solve this problem?**

Problem characterization 7

7. Does the task of solving a problem require human interactions?

- Sometimes it is useful to program computers to solve problems in ways that the majority of people would not be able to understand.
- This is fine if the level of the interaction between the computer and its human users is problem-in solution-out. But increasingly we are building programs that require intermediate interaction with people, both to provide additional input to the program and to provide additional reassurance to the user.
- The solitary problem, in which there is **no intermediate communication** and no demand for an explanation of the reasoning process.
- The conversational problem, in which **intermediate communication** is to provide either additional assistance to the computer or additional information to the user.

Question

Analyze each of them with respect to the seven problem characteristics:

- 8-puzzle
- Tower of Hanoi

Problem characteristic	Satisfied	Reasons

Problem Reduction

- In problem reduction space, the **nodes** represents problems to be solved or goals to be achieved, and the **edges** represents the decomposition of the problem into **subproblems.**
- Via **And and Or graphs.**

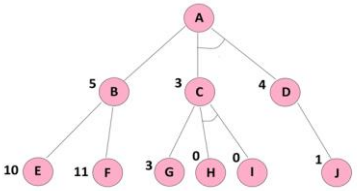

State Space Search

- Different state to solve a problem and achieve goals.
- Define problem precisely.
- $S: \{s, a, action(s), result(s,a), cost(s,a)\}$
- Action: up, down, right, left

3	4	2
7		6
5	1	8

Start State

1	2	3
8		4
7	6	5

Goal State

Searching types

- Uniformed: A **blind search**, the search algorithm produces the search tree without using any domain knowledge, which is a brute force in nature. e.g. BFS, DFS.
  - Only start and goal state is known
  - Time consuming
  - More time and space complexity
- Informed: Heuristic Approaches
  - Quick solution
  - Less time and space complexity
  - E.g. A*

Hill Climbing algorithm

- Considered greedy approach.
- The algorithm never reverses the earlier decision (**no backtracking**)
- State Space in hill climbing:
  - Local maxima
  - Global maximum
  - Current state
  - Flat local maximum
  - Shoulder

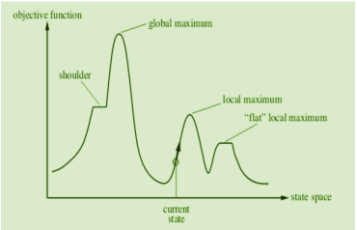

Types of Hill Climbing algorithm

- Simple Hill Climbing: It only checks its **one successor state** if found best it stops.
- Steepest-Ascent hill-climbing: It examines all the neighboring nodes and select **one neighbours node** which is closet to the goal state.
- Stochastic hill climbing: It selects one neighbours node **randomly** and decides whether to choose it as a current-state or examine another state.

Breadth First Search

- Uniformed search technique
- Use data structure **Queue** (FIFO).
- Shallowest node.

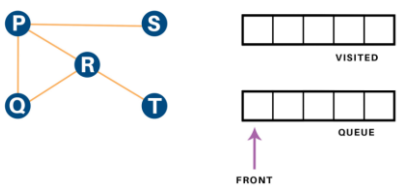

Depth First Search

- Uniformed search technique
- Use data structure **Stack** (LIFO).
- Deepest node.

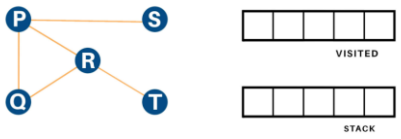

A* algorithm

- Informed searching.
- Finds the shortest path between two points.
- It finds the **optimal path**.
- It uses weighted graphs for its implementation.
- It is said to be a **optimal and complete algorithm**.
  - Optimal : As it finds least cost outcome of a problem.
  - Complete: As it finds all the possible outcome of the problem.

A* algorithm (contd.)

- Applications of A*: Robotics, video games, Logistics.
- It is an extension of Dijkstra's algorithm.
  - Instead of using priority queue to store all elements, we use **heaps** to store them.
  - It takes **more space** to store all the paths and **takes more time** to find them out.
  - It also uses **heuristic function** to provide additional information regarding how far are we from the goal.

A* algorithm (contd.)

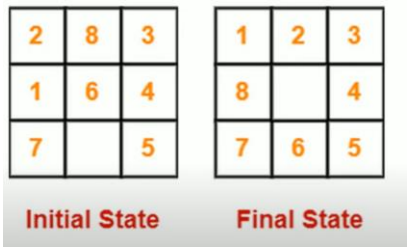

Best First search

- Informed Search, based upon **heuristic approaches.**
- Uses **Priority Queue** to store cost of nodes.
- A **greedy search** algorithm.
- Priority queue containing initial states.

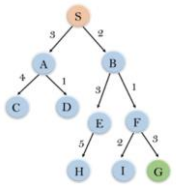

node	H (n)
A	12
B	4
C	7
D	3
E	8
F	2
H	4
I	9
S	13
G	0

Mean End Analysis (MEA)

- Means-Ends Analysis is problem-solving techniques used in Artificial intelligence for limiting search in AI programs.
- Means end analysis uses **actions (means)** to tell you how to achieve your **goal (the end)**.
- Means-ends analysis is a version of **divide-and-conquer.**
- The difference between the two is that divide-and-conquer is purely **recursive**: the subproblems that are solved are always of the same type.
- Means-ends analysis is more flexible, and less obviously recursive, because the **subproblems** that are defined for it need **not all be of the same type**.

Mean End Analysis (MEA) contd.

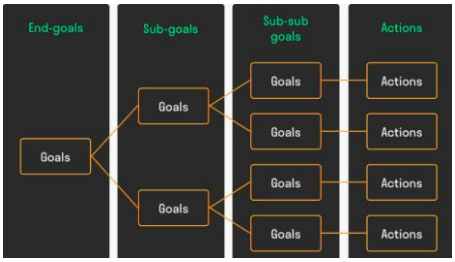

Ethics of AI

- AI ethics are the study of the **moral and ethical considerations** involved in developing and using AI.
- This field focuses on:
  - How can we make sure that autonomous machines act following our values?
  - How can we ensure that they have less probability of harming humans than other technologies?
  - What is our responsibility as designers and users of ethical AI systems?

It is also termed as **machine ethics/computational ethics.**

Ethics of AI (Contd.)

- AI ethics revolves around four main areas:
  - Safety: This refers to how well AI can avoid harming humans. (Not causing human harm or using offensive language)
  - Security: How well AI can prevent other systems from attacking it or taking advantage.
  - Privacy: How much information an AI knows about you, where it gets data from, how it stores that information.
  - Fairness: This refers to whether or not your rights as a consumer are being protected when interacting with a company.

As a new field of AI, ethics is still in development phase.

It is a **challenge in AI** to know whether a developed program is acting ethically or not.

Closing Remarks

- General questions you might have
  - Do we want to build a computer that will be like us?
  - If so, what do we need them for?
  - What will the human-computers do for humanity?
- No answers to these questions…yet research and achievement continues to progress each year - We must wait and see what the future holds

Thanks