USN | | | | | | | | | |

CMRIT — CMR Institute of Technology, Bengaluru. Accredited with A+ Grade by NAAC

INTERNAL ASSESSMENT TEST – I

Sub:	Intelligent Systems and Machine Learning Algorithms							Code:	BEC515A
Date:	07/ 11 / 2024	Duration:	90 mins	Max Marks:	50	Sem:	V	Branch:	ECE

Answer any 5 full questions

		Marks	CO	RBT
1	What are the different ways of defining AI based on the four perspectives namely thinking humanly, acting humanly, thinking rationally, acting rationally.	[10]	CO1	L1

AI is a branch of computer science that focuses on building intelligent machines that function and respond just like people. Machines, particularly computer systems, simulate human intelligence processes in this process. These include learning (acquiring knowledge and rules for using it), reasoning (using rules to arrive at approximate conclusions), and self-correction.

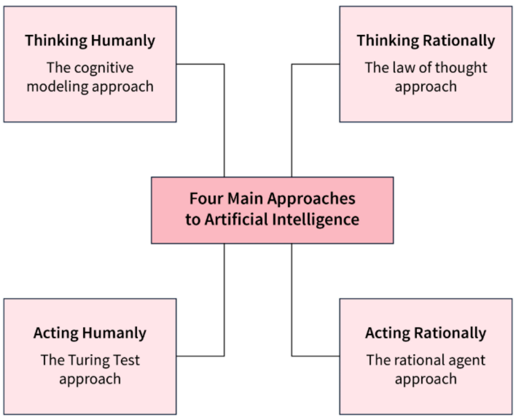

Thinking Humanly (The Cognitive Approach)

Thinking humanly, or a cognitive approach is an approach to artificial intelligence (AI) and machine learning that is inspired by the way humans think and learn. The cognitive approach aims to develop AI systems that can mimic human thought processes and behaviours, such as perception, reasoning, and problem-solving.

This approach emphasizes the importance of understanding human cognition and how it can be replicated in machines, rather than focusing solely on statistical or mathematical models.

One example of the cognitive approach is the development of expert systems, which are computer programs that can solve complex problems in a particular domain, such as medical diagnosis or financial planning.

Acting Humanly (The Turing Test Approach)

Acting humanly, also known as the **Turing Test approach**, is an approach to artificial intelligence (AI) and machine learning that focuses on creating machines that can simulate human-like behaviour and thought processes to the point where they are indistinguishable from humans.

The Turing Test approach is based on the idea that a machine can be considered intelligent if it can convincingly pass a test that was proposed by British mathematician and computer scientist Alan Turing.

The Turing Test involves a human evaluator engaging in a natural language conversation with a machine and a human, without knowing which is which. If the machine can successfully convince the evaluator that it is the human, then it is considered to have passed the Turing Test.

The Turing Test approach has led to the development of a wide range of AI technologies, including chatbots, virtual assistants, and recommendation engines.

Thinking Rationally (The Laws of Thought Approach)

Thinking rationally, or the laws of thought approach is an approach to artificial intelligence (AI) and machine learning that is based on the principles of formal logic and reasoning. The laws of thought approach aims to develop AI systems that can reason logically and make decisions based on a set of predefined rules.

In the laws of thought approach, AI systems are designed to reason deductively, by starting with a set of premises and using logical rules to conclude. This approach is often used in expert systems, where a knowledge base of facts and rules is used to solve complex problems in a particular domain.

Acting Rationally (The Rational Agent Approach)

Acting rationally, also known as the rational agent approach, is an approach to artificial intelligence (AI) and machine learning that focuses on creating intelligent agents that can act in the world to achieve their goals. The rational agent approach is based on the idea of rationality, which involves making decisions that maximize the chances of achieving one's goals, given the available information and resources.

The rational agent approach emphasizes the importance of designing agents that can reason under uncertainty and adapt to changing environments, rather than simply following a set of predefined rules.

One example of the rational agent approach is reinforcement learning, which involves training an agent to make decisions in an environment based on rewards and punishments. The agent learns to maximize its rewards by trying different actions and observing the outcomes.

Thinking Humanly	**Thinking Rationally**
"The exciting new effort to make computers think . . . *machines with minds*, in the full and literal sense." (Haugeland, 1985) "[The 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 and McDermott, 1985) "The study of the computations that make it possible to perceive, reason, and act." (Winston, 1992)
Acting Humanly	**Acting Rationally**
"The art of creating machines that perform functions that require intelligence when performed by people." (Kurzweil, 1990) "The study of how to make computers do things at which, at the moment, people are better." (Rich and Knight, 1991)	"Computational Intelligence is the study of the design of intelligent agents." (Poole *et al.*, 1998) "AI . . . is concerned with intelligent behavior in artifacts." (Nilsson, 1998)

Figure 1.1 Some definitions of artificial intelligence, organized into four categories.

Note:

Please refer : Understanding the Four Types of Artificial Intelligence (AI) | by Fahmisajid | Medium

For ease of understanding. But for VTU exam you must write as given in textbook.

2 Explain the evolution of AI. What are the key milestones which contributed towards the development of AI and have shaped the AI field to the current state. [10] CO1 L1

Artificial Intelligence (AI) is fast becoming an indispensable part of our modern world, revolutionizing numerous industries and transforming the way we live and work. The path to this remarkable technological advancement was paved with significant milestones and breakthroughs that shaped the development of AI systems as we know them today.

- In 1931, Goedel layed the foundation of Theoretical Computer Science1920-30s: He published the first universal formal language and showed that math itself is either flawed or allows for unprovable but true statements.

- In 1936, Turing reformulated Goedel's result and church's extension thereof.

The gestation of artificial intelligence (1943–1955)

- The first work that is now generally recognized as AI was done by Warren McCulloch and Walter Pitts (1943). They drew on three sources: knowledge of the basic physiology and function of neurons in the brain; a formal analysis of propositional logic due to Russell and Whitehead; and Turing's theory of computation. They proposed a model of artificial neurons in which each neuron is characterized as being "on" or "off," with a switch to "on" occurring in response to stimulation by a sufficient number of neighboring neurons.
- Donald Hebb (1949) demonstrated a simple updating rule for modifying the connection strengths between neurons. His rule, now called **Hebbian learning**, remains an influential model to this day.
- Two undergraduate students at Harvard, Marvin Minsky and Dean Edmonds, built the first neural network computer in 1950.
- 1950 Alan Turing in his article "Computing Machinery and Intelligence." Therein, he introduced the Turing Test, machine learning, genetic algorithms, and reinforcement learning.

The birth of artificial intelligence (1956)

- In 1956, John McCarthy coined the term "Artificial Intelligence" as the topic of the Dartmouth Conference, the first conference devoted to the subject.

Early enthusiasm, great expectations (1952–1969)

- In 1957, The General Problem Solver (GPS) demonstrated by Newell, Shaw & Simon
- In 1958, John McCarthy (MIT) invented the Lisp language.
- In 1959, Arthur Samuel (IBM) wrote the first game-playing program, for checkers, to achieve sufficient skill to challenge a world champion.
- In 1963, Ivan Sutherland's MIT dissertation on Sketchpad introduced the idea of interactive graphics into computing.
- In 1966, Ross Quillian (PhD dissertation, Carnegie Inst. of Technology; now CMU) demonstrated semantic nets
- In 1967, Dendral program (Edward Feigenbaum, Joshua Lederberg, Bruce Buchanan, Georgia Sutherland at Stanford) demonstrated to interpret mass spectra on organic chemical compounds. First successful knowledge-based program for scientific reasoning.
- In 1967, Doug Engelbart invented the mouse at SRI
- In 1968, Marvin Minsky & Seymour Papert publish Perceptrons, demonstrating limits of simple neural nets

A dose of reality (1966–1973)

The first AI winter (1974-1980)

The initial AI winter, occurring from 1974 to 1980, is known as a tough period for artificial intelligence (AI). During this time, there was a substantial decrease in research funding, and AI faced a sense of letdown.

- The duration between years 1974 to 1980 was the first AI winter duration. AI winter refers to the time period where computer scientist dealt with a severe shortage of funding from government for AI researches.
- During AI winters, an interest of publicity on artificial intelligence was decreased.

A boom of AI (1980-1987)

Between 1980 and 1987, AI underwent a renaissance and newfound vitality after the challenging era of the First AI Winter. Here are notable occurrences from this timeframe:

- In 1980, the first national conference of the American Association of Artificial Intelligence was held at Stanford University.
- **Year 1980:** After AI's winter duration, AI came back with an "Expert System". Expert systems were programmed to emulate the decision-making ability of a human expert. Additionally, Symbolics Lisp machines were brought into commercial use, marking the onset of an AI resurgence. However, in subsequent years, the Lisp machine market experienced a significant downturn.
- **Year 1981:** Danny Hillis created parallel computers tailored for AI and various computational functions, featuring an architecture akin to contemporary GPUs.
- **Year 1984:** Marvin Minsky and Roger Schank introduced the phrase "AI winter" during a gathering of the Association for the Advancement of Artificial Intelligence. They cautioned the business world that exaggerated expectations about AI would result in disillusionment and the eventual downfall of the industry, which indeed occurred three years later.
- **Year 1985:** Judea Pearl introduced Bayesian network causal analysis, presenting statistical methods for encoding uncertainty in computer systems.

The second AI winter (1987-1993)

- The duration between the years 1987 to 1993 was the second AI Winter duration.
- Again Investors and government stopped in funding for AI research as due to high cost but not efficient result. The expert system such as XCON was very cost effective.

The emergence of intelligent agents (1993-2011)

Between 1993 and 2011, there were significant leaps forward in artificial intelligence (AI), particularly in the development of intelligent computer programs. During this era, AI professionals shifted their emphasis from attempting to match human intelligence to crafting pragmatic, ingenious software tailored to specific tasks. Here are some noteworthy occurrences from this timeframe:

- **Year 1997:** In 1997, IBM's Deep Blue achieved a historic milestone by defeating world chess champion Gary Kasparov, marking the first time a computer triumphed over a reigning world chess champion. Moreover, Sepp Hochreiter and Jürgen Schmidhuber introduced the Long Short-Term Memory recurrent neural network, revolutionizing the capability to process entire sequences of data such as speech or video.
- **Year 2002:** for the first time, AI entered the home in the form of Roomba, a vacuum cleaner.
- **Year 2006:** AI came into the Business world till the year 2006. Companies like Facebook, Twitter, and Netflix also started using AI.
- **Year 2009:** Rajat Raina, Anand Madhavan, and Andrew Ng released the paper titled "Utilizing Graphics Processors for Extensive Deep Unsupervised Learning," introducing the concept of employing GPUs for the training of expansive neural networks.
- **Year 2011:** Jürgen Schmidhuber, Dan Claudiu Cire?an, Ueli Meier, and Jonathan Masci created the initial CNN that attained "superhuman" performance by emerging as the victor in the German Traffic Sign Recognition competition. Furthermore, Apple launched Siri, a voice-activated personal assistant capable of generating responses and executing actions in response to voice commands.
- Deep learning, big data and artificial general intelligence (2011-present)

3	A Water Jug Problem: You are given two jugs, a 4-gallon one and a 3-gallon one, a pump which has unlimited water which you can use to fill the jug, and the ground on which water may be poured. Neither jug has any measuring markings on it. How can you get exactly 2 gallons of water in the 4-gallon jug.	[10]	CO1	L2

a) Write down the production rules for the above problem

State Representation and Initial State : We will represent a state of the problem as a tuple (x, y) where x represents the amount of water in the 4-gallon jug and y represents the amount of water in the 3-gallon jug.

Note $0 \leq x \leq 4$, and $0 \leq y \leq 3$.

Our initial state: (0, 0)

Goal Predicate - state = (2, y) where $0 \leq y \leq 3$.

Operators -we must defi ne a set of operators that will take us from one state to another:

1. Fill 4-gal jug (x,y) → (4,y) x < 4
2. Fill 3-gal jug (x,y) → (x,3) y < 3
3. Empty 4-gal jug on ground (x,y) → (0,y) x > 0
4. Empty 3-gal jug on ground (x,y) → (x,0) y > 0
5. Pour water from 3-gal jug (x,y) →! (4, y - (4 - x)) to ll 4-gal jug 0 < x+y 4 and y > 0
6. Pour water from 4-gal jug (x,y) → (x - (3-y), 3) to ll 3-gal-jug 0 < x+y 3 and x > 0
7. Pour all of water from 3-gal jug (x,y) → (x+y, 0) into 4-gal jug 0 < x+y 4 and y 0
8. Pour all of water from 4-gal jug (x,y) → (0, x+y) into 3-gal jug 0 < x+y 3 and x 0

b) Write any one solution to the above problem

Through Graph Search, the following solution is found :

Gals in 4-gal jug	Gals in 3-gal jug	Rule Applied
0	0	
		1. Fill 4
4	0	
		6. Pour 4 into 3 to ll
1	3	
		4. Empty 3
1	0	
		8. Pour all of 4 into 3
0	1	
		1. Fill 4
4	1	
		6. Pour into 3
2	3	

Second Solution:

Number of Steps	*Rules applied*	*4-g jug*	*3-g jug*
1	**Initial State**	0	0
2	R2 {Fill 3-g jug}	0	3
3	R7{Pour all water from 3 to 4-g jug }	3	0
4	R2 {Fill 3-g jug}	3	3
5	R5 {Pour from 3 to 4-g jug until it is full}	4	2
6	R3 {Empty 4-gallon jug}	0	2
7	R7 {Pour all water from 3 to 4-g jug}	2	0

Goal State

4	What is an Intelligent agent? Explain in detail the structure of an intelligent agent and the interactions of the different components of an intelligent agent with neat diagram. Agent: An Agent is anything that can be viewed as perceiving its environment through	[10]	CO1	L1

	sensors and acting upon that environment through actuators. For example o A human agent has eyes, ears, and other organs for sensors and hands, legs, mouth, and other body parts foractuators. o A robotic agent might have cameras and infrared range finders for sensors and various motors foractuators. o A software agent receives keystrokes, file contents, and network packets as sensory inputs and acts on the environment by displaying on the screen, writing files, and sending network packets. 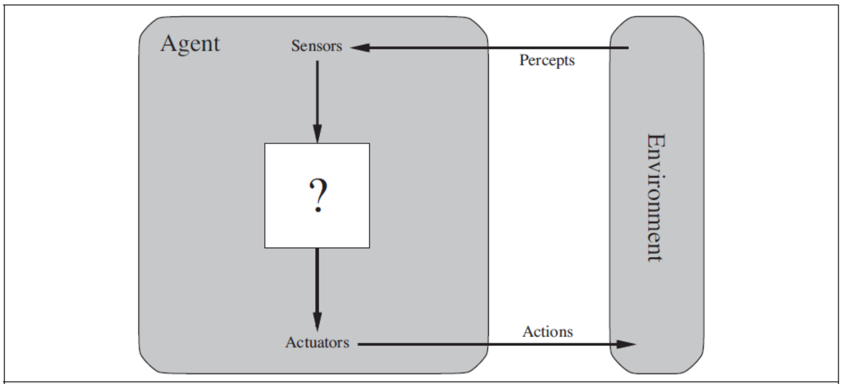 **Figure 2.1** Agents interact with environments through sensors and actuators. We use the term **percept** to refer to the agent's perceptual inputs at any given instant. An agent's **percept sequence** is the complete history of everything the agent has ever perceived. In general, *an agent's choice of action at any given instant can depend on the entire percept sequence observed to date, but not on anything it hasn't perceived.* By specifying the agent's choice of action for every possible percept sequence, we have said more or less everything there is to say about the agent. Mathematically speaking, we say that an agent's behavior is described by the **agent function** AGENT FUNCTION that maps any given percept sequence to an action. This particular world has just two locations: squares A and B. The vacuum agent perceives which square it is in and whether there is dirt in the square. It can choose to move left, move right, suck up the dirt, or do nothing. One very simple agent function is the following: if the current square is dirty, then suck, otherwise move to the other square. A partial tabulation of this agent function is shown in Figure			

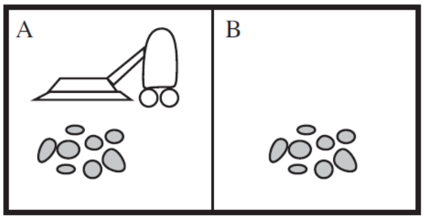

Figure 2.2 A vacuum-cleaner world with just two locations.

Percept sequence	Action
[*A*, *Clean*]	*Right*
[*A*, *Dirty*]	*Suck*
[*B*, *Clean*]	*Left*
[*B*, *Dirty*]	*Suck*
[*A*, *Clean*], [*A*, *Clean*]	*Right*
[*A*, *Clean*], [*A*, *Dirty*]	*Suck*
⋮	⋮
[*A*, *Clean*], [*A*, *Clean*], [*A*, *Clean*]	*Right*
[*A*, *Clean*], [*A*, *Clean*], [*A*, *Dirty*]	*Suck*
⋮	⋮

Figure 2.3 Partial tabulation of a simple agent function for the vacuum-cleaner world shown in Figure 2.2.

5 Explain different types of Environments, the intelligent agent may be exposed to and how do they impact the design of the agent? [10] CO1 L1

An environment in artificial intelligence is the surrounding of the agent. The agent takes input from the environment through sensors and delivers the output to the environment through actuators. There are several types of environments:

- Fully Observable vs Partially Observable
- Deterministic vs Stochastic
- Competitive vs Collaborative
- Single-agent vs Multi-agent
- Static vs Dynamic
- Discrete vs Continuous
- Episodic vs Sequential
- Known vs Unknown

Fully Observable vs Partially Observable

- When an agent sensor is capable to sense or access the complete state of an agent at each point in time, it is said to be a fully observable environment else it is partially observable.
- Maintaining a fully observable environment is easy as there is no need to keep track of the history of the surrounding.
- An environment is called **unobservable** when the agent has no sensors in all environments.
- **Examples:**
    - **Chess –** the board is fully observable, and so are the opponent's moves.
    - **Driving –** the environment is partially observable because what's around the corner is not known.

Deterministic vs Stochastic

- When a uniqueness in the agent's current state completely determines the next state of the agent, the environment is said to be deterministic.

- The stochastic environment is random in nature which is not unique and cannot be completely determined by the agent.
- **Examples:**
    - **Chess –** there would be only a few possible moves for a chess piece at the current state and these moves can be determined.
    - **Self-Driving Cars-** the actions of a self-driving car are not unique, it varies time to time.

3. Competitive vs Collaborative

- An agent is said to be in a competitive environment when it competes against another agent to optimize the output.
- The game of chess is competitive as the agents compete with each other to win the game which is the output.
- An agent is said to be in a collaborative environment when multiple agents cooperate to produce the desired output.
- When multiple self-driving cars are found on the roads, they cooperate with each other to avoid collisions and reach their destination which is the output desired.

4. Single-agent vs Multi-agent

- An environment consisting of only one agent is said to be a single-agent environment.
- A person left alone in a maze is an example of the single-agent system.
- An environment involving more than one agent is a multi-agent environment.
- The game of football is multi-agent as it involves 11 players in each team.

5. Dynamic vs Static

- An environment that keeps constantly changing itself when the agent is up with some action is said to be dynamic.
- A roller coaster ride is dynamic as it is set in motion and the environment keeps changing every instant.
- An idle environment with no change in its state is called a static environment.
- An empty house is static as there's no change in the surroundings when an agent enters.

6. Discrete vs Continuous

- If an environment consists of a finite number of actions that can be deliberated in the environment to obtain the output, it is said to be a discrete environment.
- The game of chess is discrete as it has only a finite number of moves. The number of moves might vary with every game, but still, it's finite.
- The environment in which the actions are performed cannot be numbered i.e. is not discrete, is said to be continuous.
- Self-driving cars are an example of continuous environments as their actions are driving, parking, etc. which cannot be numbered.

7.Episodic vs Sequential

- In **an Episodic task environment**, each of the agent's actions is divided into atomic incidents or episodes. There is no dependency between current and previous incidents. In each incident, an agent receives input from the environment and then performs the corresponding action.
- **Example:** Consider an example of **Pick and Place robot**, which is used to detect defective parts from the conveyor belts. Here, every time robot(agent) will make the decision on the current part i.e. there is no dependency between current and previous decisions.

- In a **Sequential environment**, the previous decisions can affect all future decisions. The next action of the agent depends on what action he has taken previously and what action he is supposed to take in the future.
- **Example:**
    - **Checkers-** Where the previous move can affect all the following moves.

8. Known vs Unknown

- In a known environment, the output for all probable actions is given. Obviously, in case of unknown environment, for an agent to make a decision, it has to gain knowledge about how the environment works.

6	What is the PEAS framework, and how is it used to define the components of an agent? Provide PEAS descriptions of the following agents: a) Medical diagnosis system b) Part-picking Robot	[10]	CO1	L2

In Artificial Intelligence (AI), various types of agents operate to achieve specific goals. The PEAS system is a critical framework used to categorize these agents based on their performance, environment, actuators, and sensors. Understanding the PEAS system is essential for grasping how different AI agents function effectively in diverse environments. Among these agents, Rational Agents are considered the most efficient, consistently choosing the optimal path for maximum efficiency.

PEAS stands for Performance measure, Environment, Actuator, Sensor.

- **Performance Measure**: Performance measure is a quantitative measure that evaluates the outcomes of an agent's actions against a predefined goal. The performance measure is crucial because it guides the agent's decision-making process, ensuring that it acts in a way that maximizes its success.

  For example, in a self-driving car, the performance measure could include criteria such as safety (avoiding accidents), efficiency (minimizing travel time), and comfort (ensuring a smooth ride). The car's AI will aim to optimize these factors through its actions.
- **Environment**: The environment includes all external factors and conditions that the agent must consider when making decisions. The environment can vary significantly depending on the type of agent and its task.

  For instance, in the case of a robotic vacuum cleaner, the environment includes the layout of the room, the presence of obstacles, and the type of floor surface. The robot must adapt to these conditions to clean effectively.

  Understanding the environment is critical for designing AI systems because it affects how the agent perceives its surroundings and interacts with them.
- **Actuators**: They are responsible for executing the actions decided by the agent based on its perceptions and decisions. In essence, actuators are the "hands and feet" of the agent, enabling it to carry out tasks.

  In a robotic arm, actuators would be the motors and joints that move the arm to pick up objects. In a software-based AI, actuators might be commands sent to other software components or systems to perform specific tasks.

  The design and choice of actuators are crucial because they directly affect the agent's ability to perform its functions in the environment.
- **Sensors**: Sensors collect data from the environment, which is then processed by the agent to make informed decisions. Sensors are the "eyes and ears" of the agent, providing it with the necessary information to act intelligently.

  In the context of an autonomous drone, sensors might include cameras, GPS,

	and altimeters, which provide the drone with information about its location, altitude, and surroundings. This data is essential for the drone to navigate and perform tasks like aerial photography or package delivery. The quality and variety of sensors used in an AI system greatly influence its ability to perceive and understand its environment. This framework not only helps in structuring intelligent agents but also ensures that they are well-equipped to achieve their goals in various environments. By carefully considering the performance measure, environment, actuators, and sensors, developers can create AI systems that are more capable, adaptable, and successful in their tasks.			

Agent Type	Performance Measure	Environment	Actuators	Sensors
Medical diagnosis system	Healthy patient, reduced costs	Patient, hospital, staff	Display of questions, tests, diagnoses, treatments, referrals	Keyboard entry of symptoms, findings, patient's answers
Part-picking robot	Percentage of parts in correct bins	Conveyor belt with parts; bins	Jointed arm and hand	Camera, joint angle sensors

7

a) What are the four main performance parameters (completeness, optimality, time complexity, and space complexity) used to evaluate search algorithms. [5] CO2 L2

Following are the four essential performance parameters of search algorithms to compare the efficiency of these algorithms:

Completeness: A search algorithm is said to be complete if it guarantees to return a solution if at least any solution exists for any random input.

Optimality: If a solution found for an algorithm is guaranteed to be the best solution (lowest path cost) among all other solutions, then such a solution for is said to be an optimal solution.

Time Complexity: Time complexity is a measure of time for an algorithm to complete its task.

Space Complexity: It is the maximum storage space required at any point during the search, as the complexity of the problem.

b) Explain Depth first search Strategy with functional description and performance measures. [5] CO2 L2

Depth-first search is a traversing algorithm used in tree and graph-like data structures. It generally starts by exploring the deepest node in the frontier. Starting at the root node, the algorithm proceeds to search to the deepest level of the search tree until nodes with no successors are reached. Suppose the node with unexpanded successors is encountered then the search backtracks to the next deepest node to explore alternative paths.

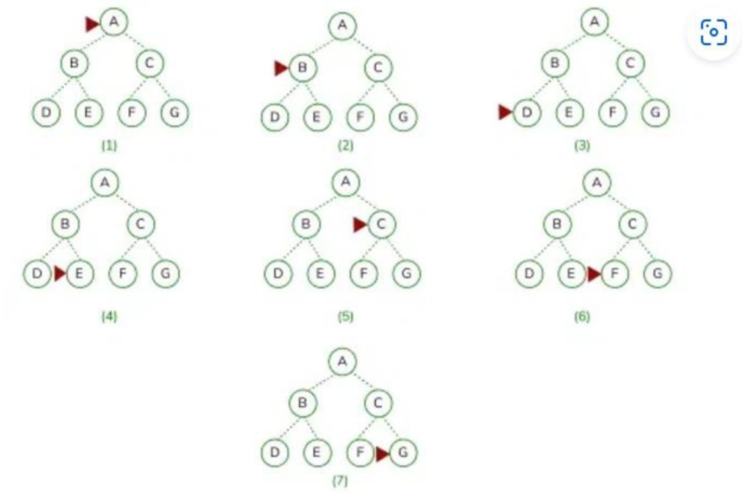

Search operation of the depth-first search

Advantage:

- DFS requires very less memory as it only needs to store a stack of the nodes on the path from root node to the current node.
- It takes less time to reach to the goal node than BFS algorithm (if it traverses in the right path).
- With the help of this we can stores the route which is being tracked in memory to save time as it only needs to keep one at a particular time.

Disadvantage:

- There is the possibility that many states keep re-occurring, and there is no guarantee of finding the solution.
- DFS algorithm goes for deep down searching and sometime it may go to the infinite loop.
- The depth-first search (DFS) algorithm does not always find the shortest path to a solution.

Completeness: DFS search algorithm is complete within finite state space as it will expand every node within a limited search tree.

Time Complexity: Time complexity of DFS will be equivalent to the node traversed by the algorithm. It is given by:

$T(n)= 1+ n^2+ n^3 +.........+ n^m=O(n^m)$

Where, m= maximum depth of any node and this can be much larger than d (Shallowest solutio

Space Complexity: DFS algorithm needs to store only single path from the root node, hence space of DFS is equivalent to the size of the fringe set, which is **O(bm)**.

CCI HOD