Department : **CSE**

Year & Semester : **III & I**

Sub Code & Sub Name : **23CAI421 & Artificial Intelligence**

Unit-I

S.No	Part-A Questions
1.	**Define Artificial Intelligence.** AI is the science of building intelligent machines that can perform tasks requiring human-like intelligence such as learning, reasoning, problem-solving, and decision-making.
2.	**List the four approaches to AI.** Acting humanly, Thinking humanly, Thinking rationally, and Acting rationally.
3.	**What is cognitive science?** Cognitive science is the interdisciplinary study of thought, learning, and mental organization, combining psychology, linguistics, philosophy, and computer modelling.
4.	**What is Bayesian analysis?** It is a probabilistic reasoning method based on Bayes' theorem, used to update beliefs and handle uncertainty in AI.
5.	**Name two contributors of mathematics to AI.** (i) Gerolamo Cardano – probability theory. (ii) Thomas Bayes – Bayes' rule for updating probabilities.
6.	**Define an agent in AI.** An agent is an entity that perceives its environment through sensors and acts upon it using actuators to achieve goals.
7.	**Differentiate between agent function and agent program.** Agent function is a mathematical mapping from percepts to actions, while agent program is its implementation running on hardware.
8.	**What is a rational agent?** A rational agent is one that chooses actions expected to maximize goal achievement, based on available information.
9.	**Give one example of a simple reflex agent.** Vacuum cleaner agent: "If dirty → suck.
10.	**State a drawback of table-driven agents.** They require huge lookup tables, making them impractical for complex tasks (e.g., Chess has ~10^120 states).
11.	**What is a utility function in AI?** A function that assigns a numerical value (utility) to states, helping an agent choose actions that maximize expected happiness or success.
12.	**ist four components of a learning agent.** Learning element, Performance element, Critic, Problem generator.
13.	**Define percept and percept sequence.** Percept: Input received by an agent at a given instant. Percept sequence: Complete history of all percepts received by the agent.
14.	**What is the Turing Test?** Proposed by Alan Turing (1950), it checks whether a machine can imitate human behavior well enough to fool a human judge in a conversation.
15.	**Mention two state-of-the-art applications of AI.** (i) IBM's Deep Blue defeating Garry Kasparov in chess. (ii) Spam filtering using machine learning algorithms.

S.No	Part-B Questions
1.	Explain in detail the different approaches to Artificial Intelligence (Acting humanly, Thinking humanly, Thinking rationally, Acting rationally).
2.	Discuss the Turing Test. What are its advantages and limitations as a measure of machine intelligence?
3.	Trace the history of Artificial Intelligence highlighting the major milestones from 1950 to the present.
4.	Explain the contributions of Philosophy, Mathematics, Economics, and Psychology to the foundations of Artificial Intelligence.
5.	Write short notes on Bayesian analysis and its role in modern AI systems.
6.	Discuss the role of expert systems in the early development of AI with suitable examples (e.g., MYCIN, DENDRAL).
7.	Explain the successes and limitations of AI in the fields of robotics and natural language processing.
8.	Define an intelligent agent. Explain the structure of agents with the help of a neat diagram.
9.	Differentiate between agent function and agent program with suitable examples.
10.	What is PEAS? Describe how the PEAS framework can be applied to the self-driving car environment.
11.	Describe the properties of task environments in AI. Illustrate each property with an example.
12.	Explain the working of simple reflex agents and model-based reflex agents with neat diagrams.
13.	Discuss the need for goal-based and utility-based agents. How do they differ from reflex agents?
14.	With a neat diagram, explain the architecture and functioning of a learning agent.
15.	Write notes on state-of-the-art AI applications in different domains (games, robotics, natural language processing, computer vision).

Unit-II

S.No	Part-A Questions
1.	**What is a problem-solving agent?** A problem-solving agent searches for a sequence of actions that leads to the goal. It uses percepts, problem formulation, and search strategies to decide actions.
2.	**What are the four steps of problem-solving?** (i) Goal formulation, (ii) Problem formulation, (iii) Search, (iv) Execution. These steps guide an agent from problem definition to solution
3.	**What is a well-defined problem?** It clearly specifies the initial state, actions, transition model, goal test, and path cost. This makes the problem suitable for algorithmic solutions.
4.	**Give an example of a toy problem.** The **8-puzzle problem** is a toy problem used for studying search strategies. It requires sliding tiles into the correct goal configuration.
5.	**Give an example of a real-world problem.** The Travelling Salesperson Problem (TSP) is a real-world optimization problem. It requires finding the shortest route covering all cities.
6.	**Define state space.** The state space is the set of all possible states reachable from the initial state. It represents the environment the agent can explore.
7.	**Define search tree.** Search tree is a representation of sequences of actions applied to states. Nodes show states, and edges show actions between them.
8.	**Differentiate between search tree and state space.** State space is the abstract set of states, while a search tree is the exploration process. The tree may include repeated states, unlike the abstract space.
9.	**What is frontier in a search tree?** Frontier is the set of nodes generated but not yet expanded. It acts as the boundary between explored and unexplored nodes.
10.	**Define completeness of a search strategy.** A search is complete if it always finds a solution when one exists. For example, BFS is complete in finite state spaces.
11.	**Define optimality of a search strategy.** A search is optimal if it always finds the least-cost solution. A* is optimal when using admissible heuristics.
12.	**What is time complexity in AI search?** It measures the number of nodes generated during search. It reflects the computation time required by the algorithm.
13.	**What is a heuristic function?** A heuristic estimates the cost from the current state to the nearest goal. It guides informed search strategies like A* and Greedy.
14.	**State the difference between Greedy Best-First Search and A Search.** Greedy uses h(n) only, while A* uses f(n) = g(n) + h(n).
15.	**What is an admissible heuristic?** A heuristic is admissible if it never overestimates the true cost to the goal. This ensures that A* always finds the optimal solution.

S.No	Part-B Questions
1.	Explain the concept of problem-solving agents in Artificial Intelligence. Describe the assumptions made about the environment in which these agents operate.
2.	What is a well-defined problem**?** Explain the essential components of a well-defined problem with suitable examples.
3.	Differentiate between toy problems and real-world problems in AI. Illustrate with examples like Vacuum World, 8-Puzzle, and Route-Finding.
4.	Explain the 8-Queens problem formulation in AI. Discuss how state-space reduction is achieved in its improved formulation.
5.	Write a detailed note on Route Finding and Travelling Salesperson Problems (TSP) as real-world AI problems. How are they different from toy problems?
6.	With neat diagrams, explain the search tree and search process for problem-solving in AI. Discuss the role of frontier, explored set, and redundant paths.
7.	What are the criteria for evaluating problem-solving performance in AI search algorithms? Explain completeness, optimality, time complexity, and space complexity with examples.
8.	Explain Breadth-First Search (BFS) algorithm with an example. Discuss its properties in terms of completeness, optimality, time, and space complexity.
9.	Discuss Uniform-Cost Search (UCS) strategy. How is it different from BFS? Explain its algorithm and properties.
10.	What is Depth-First Search (DFS)**?** Explain with an example. Discuss its advantages, disadvantages, and conditions under which it fails.
11.	Explain Depth-Limited Search and Iterative Deepening Search (IDS)**.** Why is IDS preferred over BFS and DFS in large state spaces?
12.	Explain the concept of Bidirectional Search. With an example, show how it reduces the search effort compared to other uninformed strategies.
13.	What is an informed search strategy? Explain the role **of** heuristic functions. Differentiate between Greedy Best-First Search and A* Search with suitable examples.
14.	Discuss the optimality and conditions for A *algorithm. What are admissible and consistent heuristics? Illustrate with the 8-puzzle example.
15.	What are heuristic functions**?** Explain the generation of admissible heuristics using relaxed problems, subproblems, and pattern databases with examples.

Unit-III

S.No	Part-A Questions
1.	**What is reinforcement in reinforcement learning?** Reinforcement is feedback in the form of rewards or punishments. It helps the agent decide which actions are good or bad.
2.	**Give an example of reinforcement in games.** In chess, a reward may come only at the end (win/loss). In ping-pong, each scored point can be treated as a reward.
3.	**What is an optimal policy?** An optimal policy maximizes the expected total reward. It defines the best actions for every state.
4.	**What is passive reinforcement learning?** Passive learning uses a fixed policy π and evaluates its utility. The agent does not change the policy, only learns state utilities.
5.	**What is active reinforcement learning?** Active learning requires the agent to learn both what to do and where to go. It must balance exploration and exploitation.
6.	**Define exploration in RL.** Exploration is trying new actions to discover better rewards. It prevents the agent from getting stuck in suboptimal policies.
7.	**How is SARSA different from Q-learning?** SARSA updates Q-values using the actual action taken. Q-learning updates using the best possible action.
8.	**Define NLP.** Natural Language Processing (NLP): Natural Language Processing is a field of computer science and artificial intelligence that focuses on the interaction between computers and human (natural) languages. It enables machines to understand, interpret, generate, and respond to human language in a meaningful way. Example: Applications include language translation, speech recognition, chatbots, sentiment analysis, and text classification.
9.	**What is an n-gram?** An n-gram is a contiguous sequence of n items (characters, words, or syllables) from a text or speech.
10.	**What is the main difference between formal and natural languages?** Formal languages have precise syntax and semantics; natural languages are ambiguous and defined probabilistically.
11.	**What is a corpus?** A corpus is a large and structured collection of texts used to train language models.
12.	**Define perplexity in language models.** Perplexity measures how well a language model predicts a sample; lower perplexity indicates better prediction.
13.	**What is the "bag of words" model?** A representation where a text is treated as an unordered collection of words, ignoring grammar and word order.
14.	**Define precision in information retrieval.** Precision is the proportion of retrieved documents that are relevant: Precision=Relevant retrieved/Relevant retrieved
15.	**What is the difference between HITS and PageRank?** HITS is query-dependent and identifies hubs and authorities; PageRank is query-independent and ranks pages based on link structure.

S.No	Part-B Questions
1.	Explain the active reinforcement learning with examples.
2.	Describe the three methods of passive reinforcement learning: Direct Utility Estimation, Adaptive Dynamic Programming, and Temporal-Difference Learning
3.	Explain Q-learning and compare it with SARSA.
4.	Discuss applications of reinforcement learning in (i) Game playing and (ii) Robotics.
5.	Explain policy search in reinforcement learning. How does it differ from Q-learning?
6.	Explain the concept of language models in NLP and differentiate between formal and natural language models.
7.	Describe n-gram models and their use in language modeling.
8.	Explain smoothing techniques in n-gram language models. Why are they needed?
9.	What is text classification? Discuss methods used in NLP.
10.	Explain information retrieval (IR) and IR scoring functions.
11.	Discuss link analysis algorithms: PageRank and HITS.
12.	Describe information extraction and approaches: finite-state automata, HMMs, and CRFs.
13.	Explain cross-validation and perplexity as methods for evaluating language models.
14.	Compare rule-based and probabilistic approaches in information extraction.
15.	Discuss the role of information extraction in ontology construction and machine reading.

Unit-IV

S.No	Part-A Questions
1.	**What are Phrase Structure Grammars (PSGs)?** PSGs are grammars that define how words combine into phrases and sentences using hierarchical structures. Lexical categories (noun, verb) form syntactic categories (NP, VP), which are combined into phrase-structure trees.
2.	**Define Generative Capacity of grammars**. Generative capacity refers to the set of languages a grammar can describe. Chomsky's hierarchy classifies grammars into recursively enumerable, context-sensitive, context-free, and regular grammars, each increasingly restricted in expressive power.
3.	**What is a Probabilistic Context-Free Grammar (PCFG)?** A PCFG is a CFG where each rule has an associated probability. It assigns probabilities to sentences, helping to handle ambiguities in natural language parsing.
4.	**What is parsing in NLP?** Parsing is analyzing a sentence to uncover its phrase structure using grammar rules. Chart parsers, like the CYK algorithm, efficiently store sub-analyses to avoid repeated computation.
5.	**Explain lexicalized PCFG.** A lexicalized PCFG augments rules with the head word of a phrase, capturing dependencies between words (e.g., verb-object relations) to improve parsing accuracy
6.	**What is machine translation (MT)?** MT automatically translates text from a source language to a target language. Approaches include interlingua-based, transfer-based, and statistical MT using parallel corpora and probabilistic models.
7.	**What is speech recognition?** Answer: Speech recognition identifies a sequence of words from an acoustic signal. It uses acoustic models (sound-to-phoneme mapping) and language models (word probabilities) to find the most likely word sequence.
8.	**Define acoustic and language models in speech recognition.** Acoustic model: Represents the relationship between sounds and words/phonemes. Language model: Assigns probabilities to sequences of words, helping disambiguate similar-sounding phrases
9.	**Give types of ambiguity in language.** Lexical ambiguity: Word has multiple meanings. Syntactic ambiguity: Phrase has multiple parses. Semantic ambiguity: Syntactic ambiguity leads to multiple interpretations.
10.	**What is a definite clause grammar (DCG)?** A DCG translates grammar rules into logical clauses, allowing parsing as logical inference and supporting bottom-up or top-down reasoning.
11.	**What is diffuse reflection and why is it important in vision?** Diffuse reflection scatters light evenly in all directions from a surface. Its brightness does not depend on viewing angle. It helps in perceiving the shape and colour of objects accurately.
12.	**Define trichromacy in human vision.** Human eyes have three types of cone cells sensitive to red, green, and blue wavelengths. Any spectral color can be represented as a mixture of these three colors. This principle allows devices like cameras and displays to reproduce color.
13.	**Define optical flow and its application.** Optical flow is the apparent motion of pixels between video frames. It provides direction and speed of moving objects. It helps in depth estimation, motion analysis, and action recognition.

14.	**Define optical flow and its application.** Optical flow is the apparent motion of pixels between video frames. It provides direction and speed of moving objects. It helps in depth estimation, motion analysis, and action recognition.
15.	**What is a deformable template model in pose estimation?** Deformable template models represent body parts as connected segments with allowable positions and orientations. They account for human body articulation. This helps locate limbs, torso, and head in images or videos.

S.No	Part-B Questions
1.	Explain Phrase Structure Grammars (PSGs) and their importance in NLP.
2.	Describe Chomsky's hierarchy of grammars and give examples of each type.
3.	What is a Probabilistic Context-Free Grammar (PCFG)? How does it handle ambiguity in natural language?
4.	Explain the process of parsing in NLP and discuss the role of chart parsers like the CYK algorithm.
5.	Define lexicalized PCFGs and explain how they improve upon standard PCFGs.
6.	Discuss case agreement and subject–verb agreement in augmented grammars.
7.	Explain semantic interpretation in NLP with an example (e.g., arithmetic expressions).
8.	Describe the main challenges in machine translation and the differences between interlingua, transfer-based, and statistical approaches.
9.	Explain the statistical approach to machine translation using parallel corpora and EM algorithm.
10.	What are the key components of a speech recognition system? Explain the roles of acoustic and language models.
11.	Explain perception in AI and the challenges involved in interpreting sensor inputs.
12.	Describe image formation including pinhole cameras, perspective projection, lens systems, and scaled orthographic projection.
13.	Explain image segmentation, including low-level (brightness, colour, texture) approaches and the concept of superpixels.
14.	Describe object recognition by appearance, including sliding-window detection and the use of features robust to illumination changes.
15.	Explain pedestrian detection using HOG features, including feature construction, moving-window classification, and non-maximum suppression

Unit-V

S.No	Part-A Questions
1.	**What is a robot and what are its primary components?** A robot is a physical agent that performs tasks by interacting with the environment. Primary components include effectors (e.g., grippers, wheels) for action and sensors for perception. They work together to perceive and manipulate the world.
2.	**Name the three primary types of robots.** Robots are categorized as manipulators (anchored robot arms), mobile robots (wheeled, legged, or tracked), and mobile manipulators (combining mobility and manipulation). Each type is designed for specific tasks and environments.
3.	**Differentiate between active and passive sensors.** Passive sensors observe signals from the environment (e.g., cameras). Active sensors emit energy and measure reflections (e.g., sonar, lidar). Active sensors provide more info but consume more power and may interfere with each other.
4.	**What are proprioceptive sensors?** Proprioceptive sensors measure the robot's own motion or configuration, such as shaft encoders for joint angles or wheel revolutions, and gyroscopes for orientation. They help in odometry and state estimation.
5.	**Define degrees of freedom (DOF) in robotics.** DOF is the number of independent ways a robot or effector can move. For example, a mobile AUV has six DOFs (three positions + three orientations). More DOFs often mean better maneuverability and control.
6.	**What is the difference between holonomic and nonholonomic robots?** Holonomic robots have equal controllable and effective DOFs, making them easier to control. Nonholonomic robots have fewer controllable DOFs than effective DOFs (e.g., cars), making some movements restricted.
7.	**Explain Monte Carlo Localization (MCL).** MCL uses particle filtering to estimate a robot's position in a map. Particles represent hypotheses, which are updated using motion and sensor models. Over time, particles cluster around the robot's true location.
8.	**What is a configuration space in motion planning?** Configuration space represents all possible robot states (position, orientation, joint angles). Planning in this space accounts for linkage and motion constraints better than direct workspace planning.
9.	**How does a potential field help in path planning?** Potential fields assign higher costs to states near obstacles. Robots minimize both distance to goal and potential cost, producing safer trajectories while avoiding collisions.
10.	**What is the difference between MDP and POMDP in robotics?** A: MDP assumes full observability; optimal policies are based on exact states. POMDP handles partial observability; policies depend on a belief state (probability distribution), enabling information-gathering actions under uncertainty.
11.	**What is a dynamic state in robotics?** A dynamic state includes a robot's position and velocity, and sometimes acceleration. It captures how the robot's state changes under forces over time. Dynamic states are modeled using differential equations.
12.	**What is a PD controller?** A PD (Proportional-Derivative) controller corrects errors using current deviation and rate of change. It generates forces to keep the robot on a preplanned path. It can cause vibration if not properly tuned.
13.	**What is a reactive control in robots?** Reactive control uses sensor inputs to make real-time decisions.

	It handles low-level tasks without planning the global path. It is fast but may not achieve optimal long-term goals.
14.	**What is strong AI?** Strong AI or AGI aims to replicate human-level intelligence. It can learn, reason, and plan independently. It remains theoretical and is not yet realized.
15.	**Name two applications of deep learning.** Self-driving cars – for detecting objects and decision making. Image recognition – for classifying and tagging images accurately. Also used in speech recognition and pattern prediction.

S.No	Part-B Questions
1.	Explain the main types of robots and their applications, highlighting manipulators, mobile robots, and mobile manipulators.
2.	Describe the role of sensors in robotics. Distinguish between active and passive sensors with examples.
3.	Explain range sensors in robotics. Compare sonar, stereo vision, optical range finders, and lidar.
4.	Discuss the various types of location sensors used in robotics and explain the working of GPS and Differential GPS.
5.	What are proprioceptive sensors? Explain their importance in robot motion and state estimation.
6.	Define effectors and degrees of freedom (DOF). Explain their significance in manipulator and mobile robot design.
7.	Compare holonomic and nonholonomic robots. Give examples of each and explain their advantages and limitations.
8.	Describe the process of robotic perception. Explain how belief states are updated using sensor and motion models.
9.	Explain Monte Carlo Localization (MCL) and its working in mobile robot localization. Include a discussion on particle filtering.
10.	Describe configuration space in motion planning. How does it differ from workspace representation and why is it important?
11.	Explain cell decomposition and skeletonization methods in robot path planning. Discuss their advantages and limitations.
12.	Discuss planning under uncertainty in robotics. Explain MDPs, POMDPs, and approaches for handling uncertainty in motion planning
13.	Explain the difference between kinematic and dynamic states in robots, and discuss the role of controllers in robot motion.
14.	Describe the subsumption architecture and the three-layer hybrid architecture in robotic software systems
15.	Define strong AI and weak AI, and explain the Turing Test and Chinese Room Argument.