Paper XII (A): Introduction to Artificial Intelligence

Comprehensive Question Bank (Units I - IV)
Strictly mapped to Semester Syllabus Requirements

Question Type	Marks per Qn.	Qns. per Unit	Total Qns.
Short Answer Questions	1 Mark	20	80
Brief Explanation Questions	2 Marks	15	60
Analytical/Detailed Questions	5 Marks	10	40
Long Essay Questions	8 Marks	5	20

UNIT-I: INTRODUCTION TO AI & UNINFORMED SEARCH

1 Mark Questions

Q1. Define Artificial Intelligence according to John McCarthy.
Q2. What is the core objective of the Turing Test?
Q3. Define an Intelligent Agent in AI.
Q4. What constitutes the environment of an AI agent?
Q5. What is meant by 'State Space Search'?
Q6. List the components of a production system.
Q7. What is a goal test in state space search?
Q8. Define an un-informed search strategy.
Q9. State the completeness property of Breadth-First Search (BFS).
Q10. What is the space complexity of Depth-First Search (DFS) for depth d and branching factor b?
Q11. Define Iterative Deepening Search (IDS).
Q12. What is the primary advantage of bidirectional search?
Q13. Define Beam Search.
Q14. What determines the width of a Beam Search?
Q15. What is a production rule?
Q16. Give one example of a characteristic of an AI problem.
Q17. What is the role of an actuator in an intelligent agent?
Q18. Define a rational agent.
Q19. What is a closed-world assumption in state space?
Q20. What is the branching factor of a search tree?

2 Marks Questions

Q1. Differentiate between a strong AI and a weak AI approach.

Q2. Explain the concept of acting rationally vs. thinking humanly.

Q3. Describe the significance of the Total Turing Test over the classic Turing Test.

Q4. What is the difference between a simple reflex agent and a model-based reflex agent?

Q5. Define the PEAS specification with an example for an automated taxi driver.

Q6. What are the major characteristics of an AI production system?

Q7. Explain why Depth-First Search can be non-optimal and incomplete in infinite spaces.

Q8. Compare the time and space complexity trade-offs of Breadth-First Search.

Q9. Why is Iterative Deepening Search preferred for large state spaces with unknown depth?

Q10. Explain how a bidirectional search reduces the search time complexity.

Q11. Describe the evaluation mechanism used in Beam Search.

Q12. What is the difference between a monotonic and non-monotonic production system?

Q13. Define the concept of path cost in state space representation.

Q14. Distinguish between accessible and inaccessible environments.

Q15. What is a deterministic environment? Give an example.

5 Marks Questions

Q1. Discuss the primary scope and diverse modern applications of Artificial Intelligence.

Q2. Analyze the characteristics of AI problems that distinguish them from conventional computing problems.

Q3. Explain the architecture and execution cycle of a typical Production System with a neat block diagram.

Q4. Elaborate on the concept of Intelligent Agents. Describe the structure of a Goal-based agent.

Q5. Provide a complete formulation of the 8-Puzzle problem as a State Space Search model.

Q6. Compare and contrast Breadth-First Search and Depth-First Search across completeness, optimality, time, and space complexities.

Q7. Explain the working mechanism of Iterative Deepening Depth-First Search (IDDFS) and prove why it is highly memory efficient.

Q8. Describe the operational framework of Bidirectional Search. What are the practical difficulties in implementing it?

Q9. Discuss Beam Search in detail, explaining how the beam width parameter affects performance and risk of missing the optimal path.

Q10. Explain the classification of environments (Discrete vs. Continuous, Static vs. Dynamic, Episodic vs. Sequential).

8 Marks Questions

Q1. Define the Turing Test in AI. Critically evaluate its limitations, the cognitive arguments against it (such as Searle's Chinese Room), and its modern relevance.

Q2. Provide an in-depth classification of Intelligent Agents based on their structural designs. Supplement your answer with architectural diagrams for Utility-based and Learning agents.

Q3. Formulate the State Space Search framework. Detail how an abstract real-world problem is systematically mapped into states, initial state, actions, transition models, path costs, and goal criteria.

Q4. Write detailed algorithms for both Breadth-First Search (BFS) and Depth-First Search (DFS). Trace their execution on an arbitrary graph containing loops and dead-ends.

Q5. Explain the Uninformed Search landscape. Compare BFS, DFS, Depth-Limited Search, Iterative Deepening, and Bidirectional Search in a tabular format using standard complexity notations.

UNIT-II: INFORMED SEARCH & CONSTRAINT SATISFACTION

1 Mark Questions

Q1. What is a heuristic function?

Q2. State the basic principle of Generate-and-Test method.

Q3. What is the main limitation of pure Hill Climbing search?

Q4. Define Best-First Search.

Q5. Write the evaluation formula f(n) for the A* algorithm.

Q6. What does g(n) represent in the A* algorithm?

Q7. What does h(n) represent in the A* algorithm?

Q8. When is a heuristic function said to be admissible?

Q9. What is problem reduction?

Q10. What kind of graph is used to represent problem reduction?

Q11. Name the algorithm used to search an AND-OR graph.

Q12. What is a Constraint Satisfaction Problem (CSP)?

Q13. Define the components of a CSP.

Q14. What is backtracking search in CSP?

Q15. Give an example of a classic CSP problem.

Q16. Define Means-End Analysis.

Q17. What is the role of an operator in Means-End Analysis?

Q18. Define the plateau problem in hill climbing.

Q19. What is a local maximum?

Q20. What is a ridge in the context of heuristic search?

2 Marks Questions

Q1. Differentiate between informed (heuristic) search and uninformed search.

Q2. Explain the difference between a local maximum and a global maximum in Hill Climbing.

Q3. What is the significance of the heuristic function being consistent or monotonic?

Q4. Describe the conditions under which the A* algorithm is guaranteed to be optimal.

Q5. How does the AO* algorithm handle interaction between subgoals in an AND-OR graph?

Q6. Explain the concept of forward checking in Constraint Satisfaction Problems.

Q7. What is the difference between an AND node and an OR node in problem reduction?

Q8. Describe how Means-End Analysis utilizes a difference table.

Q9. Explain the term 'Constraint Propagation' with a simple example.

Q10. What is the random-restart hill climbing method and why is it used?

Q11. Compare Greedy Best-First Search with A* search regarding path accuracy.

Q12. Explain the state representation of the Map Coloring problem as a CSP.

Q13. How does node consistency differ from arc consistency in a CSP grid?

Q14. What is the main purpose of the General Problem Solver (GPS)?

Q15. Explain the significance of the heuristic evaluation function h(n) = 0 in A*.

5 Marks Questions

Q1. Describe the Generate-and-Test paradigm and discuss its limitations when applied to large search spaces.

Q2. Explain the simple Hill Climbing algorithm and analyze its key drawbacks: Local Maxima, Plateaus, and Ridges.

Q3. Elaborate on the working principle of Greedy Best-First Search. Illustrate its behavior with a routing example.

Q4. Explain the mathematical framework of the A* algorithm. Show how admissibility guarantees finding the shortest path.

Q5. Discuss Problem Reduction using AND-OR graphs. Explain why standard state space trees are insufficient.

Q6. Outline the steps of the AO* algorithm for searching an AND-OR graph.

Q7. Define Constraint Satisfaction Problems (CSPs). Explain how backtracking search is systematically used to solve CSPs.

Q8. Explain the concept of Arc Consistency (AC-3 algorithm) and how it eliminates invalid domains in CSPs.

Q9. Elaborate on Means-End Analysis (MEA) and explain how operators are selected based on difference reduction.

Q10. Illustrate the application of A* search to solve the 8-puzzle problem using an appropriate heuristic like Manhattan Distance.

8 Marks Questions

Q1. Provide a rigorous, step-by-step description of the A* algorithm. Prove mathematically that if a heuristic function h(n) is admissible, A* search using a tree-search framework is optimal.

Q2. Compare and contrast the Hill Climbing search strategy with the Simulated Annealing approach. Explain how Simulated Annealing avoids getting trapped in local optima using thermodynamic probabilities.

Q3. Detail the AO* algorithm for searching problem-reduction graphs. Explain its execution cycle, node-marking strategy, and how cost estimates are dynamically updated backward through the graph.

Q4. Formulate the Cryptarithmetic problem (e.g., SEND + MORE = MONEY) as a Constraint Satisfaction Problem. Walk through the backtracking search with constraint propagation steps to find a solution.

Q5. Provide an in-depth review of Means-End Analysis. Explain its core data structures, its implementation in early AI systems like GPS, and trace a complete blocks-world problem using a difference table.

UNIT-III: KNOWLEDGE REPRESENTATION & REASONING

1 Mark Questions

Q1. What is Propositional Logic?

Q2. Define Predicate Logic (First-Order Logic).

Q3. What is a literal in logic?

Q4. Define Horn Clause.

Q5. What is the purpose of Resolution in logic?

Q6. What does the Unification algorithm do?

Q7. Define a substitution in First-Order Logic.

Q8. What is forward chaining?

Q9. What is backward chaining?

Q10. Define conflict resolution in production rule systems.

Q11. What is a Semantic Network?

Q12. Define a Frame in knowledge representation.

Q13. What are slots and fillers in a frame?

Q14. What is a Conceptual Dependency (CD)?

Q15. Who introduced Conceptual Dependency theory?

Q16. What is a Script in knowledge representation?

Q17. Define a Bayesian Belief Network.

Q18. What does a node represent in a Bayesian Network?

Q19. State the core purpose of the Dempster-Shafer theory.

Q20. What is a conditional probability table (CPT)?

2 Marks Questions

Q1. Distinguish between Propositional Logic and First-Order Predicate Logic.

Q2. Explain the concept of Skolemization in converting to Conjunctive Normal Form (CNF).

Q3. What is the difference between a ground term and a variable in the Unification process?

Q4. Explain the Modus Ponens inference rule with a logical expression.

Q5. Compare forward chaining and backward chaining based on whether they are data-driven or goal-driven.

Q6. What is the inheritance mechanism in Semantic Networks?

Q7. Describe the purpose of a Script tracker or component parts like 'Props' and 'Roles'.

Q8. How does Dempster-Shafer theory handle ignorance compared to classical Bayesian probability?

Q9. Explain the concept of Joint Probability Distribution in Bayesian networks.

Q10. What is a conflict set in an inference engine?

Q11. Define the primitive acts 'ATRANS' and 'PTRANS' in Conceptual Dependency.

Q12. Explain the concept of default values in Frame-based systems.

Q13. What is the role of quantifier switching or scoping in Predicate Logic?

Q14. Define monotonic reasoning vs. non-monotonic reasoning.

Q15. What is a resolution refutation proof?

5 Marks Questions

Q1. Explain the systematic process of converting an arbitrary First-Order Logic formula into Conjunctive Normal Form (CNF).

Q2. Describe the Unification Algorithm in First-Order Logic. Work out an example with expressions P(x, f(y)) and P(g(z), w).

Q3. Compare and contrast Forward Chaining and Backward Chaining mechanisms in rule-based expert systems.

Q4. Explain the strategy of Resolution Refutation. Show how to prove a conclusion using resolution given a set of simple premises.

Q5. Elaborate on Semantic Nets as a knowledge representation tool. Discuss how inheritance works and point out its ambiguities.

Q6. Explain the structure of a Frame system. Create an appropriate Frame network representing a 'University Faculty Member'.

Q7. Discuss the core primitive actions and conceptual dependency structures proposed by Schank for representing human language semantics.

Q8. What is a Script? Design a complete Script detailing the common sequences of events in a 'Restaurant Visit'.

Q9. Explain the mathematical framework of Bayesian Belief Networks and how they represent conditional independence among variables.

Q10. Discuss the Dempster-Shafer Theory of Evidence. Explain how belief functions and plausibility levels are determined.

8 Marks Questions

Q1. Elaborate on the resolution principle for Predicate Logic. Given the premises: 'All dogs are animals', 'All animals eat food', and 'Fido is a dog', prove using Resolution Refutation that 'Fido eats food'.

Q2. Provide a definitive comparison of structural knowledge representation schemas: Semantic Networks, Frames, Conceptual Dependencies, and Scripts. Use diagrams to highlight their differences and best use cases.

Q3. Explain the foundational challenges of Reasoning Under Uncertainty. Provide a comprehensive description of Bayesian Networks, detailing how conditional probabilities are formulated and chained.

Q4. Detail the Dempster-Shafer theory. Explain the mathematical formula for Dempster's Rule of Combination to combine two independent bodies of evidence, illustrating with a numeric diagnostic scenario.

Q5. Discuss Conflict Resolution in Rule-Based Systems. Analyze at least five different conflict resolution strategies (e.g., Recency, Specificity, Refractory, Salience) and explain their impact on inference execution.

UNIT-IV: NATURAL LANGUAGE PROCESSING & EXPERT SYSTEMS

1 Mark Questions

Q1. What is Natural Language Processing (NLP)?

Q2. Name the five primary levels of knowledge in language understanding.

Q3. What is phonetic analysis?

Q4. Define lexical analysis in NLP.

Q5. What is syntactic parsing?

Q6. Define semantic analysis.

Q7. What is pragmatic analysis?

Q8. What is a Transition Network in syntactic parsing?

Q9. What does a Recursive Transition Network (RTN) handle?

Q10. Define top-down parsing.

Q11. Define bottom-up parsing.

Q12. What is an Expert System?

Q13. Name the two main components of an Expert System core.

Q14. What is the role of an Inference Engine?

Q15. What is a Knowledge Base?

Q16. Name the domain for which MYCIN was developed.

Q17. What was the primary function of the DENDRAL expert system?

Q18. Define the role of a Knowledge Engineer.

Q19. What is user interface shell in expert systems?

Q20. What are certainty factors in expert systems?

2 Marks Questions

Q1. Explain the difference between syntax and semantics in NLP.

Q2. What is syntactic ambiguity? Provide a classic sentence showing ambiguity.

Q3. Describe the basic concept of a Recursive Transition Network (RTN).

Q4. Compare top-down parsing and bottom-up parsing strategies briefly.

Q5. What is the distinction between an Expert System and a conventional program?

Q6. Explain the significance of the MYCIN expert system in AI history.

Q7. What contribution did DENDRAL make to the development of rule-based systems?

Q8. Describe the Expert System Development Cycle (ESDC) phases in short.

Q9. What is a explanation facility in an expert system?

Q10. Define the concept of discourse integration in NLP.

Q11. What is morphological analysis? Give an example.

Q12. What are the limitations of rule-based expert systems?

Q13. Explain what a shell is in the context of building expert systems.

Q14. What is an augmented transition network (ATN)?

Q15. Explain how knowledge acquisition acts as a bottleneck in expert system creation.

5 Marks Questions

Q1. Elaborate on the distinct Levels of Knowledge required for comprehensive Natural Language Understanding.

Q2. Explain the key phases of Natural Language Understanding, detailing the transition from tokenization to pragmatic interpretation.

Q3. Describe the Top-Down Parsing approach with a suitable grammar example and trace its parsing tree creation.

Q4. Discuss Bottom-Up Parsing, highlight its shift-reduce actions, and compare its efficiency with top-down approaches.

Q5. Explain Transition Networks and show how Recursive Transition Networks (RTNs) model complex grammatical structures.

Q6. Describe the architectural components of a standard Expert System with a comprehensive block diagram.

Q7. Detail the Expert System Development Cycle (ESDC), highlighting the responsibilities of domain experts and knowledge engineers.

Q8. Provide an in-depth case study of the MYCIN expert system, including its rule structure and certainty factor handling.

Q9. Provide an in-depth case study of the DENDRAL expert system, detailing how it used heuristic rules to deduce molecular structures.

Q10. Discuss the primary advantages and major bottlenecks faced during the development and real-world deployment of Expert Systems.

8 Marks Questions

Q1. Provide a comprehensive, structural breakdown of the Phases of Natural Language Understanding. Discuss the precise engineering challenges encountered at each level (Lexical, Syntactic, Semantic, Pragmatic).

Q2. Elaborate on Context-Free Grammars (CFGs) in NLP. Write a sample grammar rule set and demonstrate how a specific input sentence is systematically processed via both top-down and bottom-up parsing algorithms.

Q3. Design the complete system architecture of an industrial-grade Expert System. Detail the mechanics of the Knowledge Base, Inference Engine (Forward/Backward reasoning), Working Memory, and Explanation Facility.

Q4. Trace the evolution of early expert systems by performing an extensive comparative analysis of MYCIN and DENDRAL. Focus on their knowledge representation techniques, inference mechanics, and domain impacts.

Q5. Discuss the Knowledge Acquisition Bottleneck in Expert Systems. Detail the methodologies, automated tools, and engineering strategies employed to extract domain expertise and formulate a robust production rule base.