

1. Who coined the term “Artificial Intelligence”?

- A) Alan Turing
- B) John McCarthy
- C) Marvin Minsky
- D) Herbert Simon

Answer: B) John McCarthy

Explanation: John McCarthy introduced the term *Artificial Intelligence* in 1956 during the Dartmouth Conference.

2. The Turing Test evaluates a machine’s ability to ____.

- A) Learn from data
- B) Exhibit human-like conversation
- C) Solve optimization problems
- D) Recognize images

Answer: B) Exhibit human-like conversation

Explanation: Proposed by Alan Turing, the test checks if a machine’s replies can fool a human interrogator.

3. In the PEAS framework, the letter P stands for ____.

- A) Perception
- B) Performance measure

C) Planning

D) Policy

Answer: B) Performance measure

Explanation: PEAS = Performance measure, Environment, Actuators, Sensors – used to specify an AI agent's task.

4. Which of the following is *not* an AI environment property?

A) Observable vs Partially observable

B) Deterministic vs Stochastic

C) Static vs Dynamic

D) Linear vs Nonlinear

Answer: D) Linear vs Nonlinear

Explanation: Linearity is a mathematical concept, not a property used to classify AI environments.

5. A reflex agent selects actions based on ____.

A) Future states

B) Condition-action rules from current percepts

C) Utility values

D) Search trees

Answer: B) Condition-action rules from current percepts

Explanation: Reflex agents act immediately using rules without reasoning about future states.

6. A rational agent is one that ____.

- A) Always achieves success
- B) Chooses actions to maximize expected performance
- C) Uses symbolic logic only
- D) Mimics human emotion

Answer: B) Chooses actions to maximize expected performance

Explanation: Rationality means selecting actions expected to produce the best outcome given available information.

7. Which search strategy expands the shallowest node first?

- A) Depth-first search
- B) Best-first search
- C) Breadth-first search
- D) Hill-climbing

Answer: C) Breadth-first search

Explanation: BFS explores all nodes at a given depth before moving deeper.

8. An admissible heuristic ____.

- A) Overestimates the true cost
- B) Never overestimates the true cost to goal
- C) Underestimates randomly
- D) Is domain-specific only

Answer: B) Never overestimates the true cost to goal

Explanation: Admissible heuristics guarantee optimality for A* search because they are optimistic.

9. The A algorithm is optimal if the heuristic is ____.*

- A) Greedy
- B) Admissible and consistent
- C) Random
- D) Non-monotonic

Answer: B) Admissible and consistent

Explanation: These conditions ensure A* finds the lowest-cost path without re-expanding nodes.

10. The MRV (Minimum Remaining Values) heuristic is used in ____.

- A) Search trees
- B) Neural networks
- C) Constraint Satisfaction Problems
- D) Decision trees

Answer: C) Constraint Satisfaction Problems

Explanation: MRV chooses the variable with fewest legal values to reduce branching early.

11. Fuzzy logic allows truth values ____.

- A) Only 0 and 1
- B) Between 0 and 1
- C) Greater than 1
- D) Negative

Answer: B) Between 0 and 1

Explanation: Fuzzy logic represents partial truths rather than crisp true/false values.

12. Which algorithm is complete and optimal for equal step costs?

- A) Depth-first search

- B) Breadth-first search
- C) Greedy search
- D) Hill-climbing

Answer: B) Breadth-first search

Explanation: BFS finds the shallowest goal node first when each step has equal cost.

13. Knowledge representation deals with ____.

- A) Learning weights
- B) How information is stored and used for reasoning
- C) Data compression
- D) Graphics rendering

Answer: B) How information is stored and used for reasoning

Explanation: Knowledge representation defines structures (logic, frames, semantic nets) to support AI reasoning.

14. Which inference rule is complete for propositional logic?

- A) Modus Ponens
- B) Resolution
- C) Backward chaining
- D) Forward chaining

Answer: B) Resolution

Explanation: Resolution can derive any logical consequence from a set of clauses in propositional logic.

15. Naïve Bayes classifier assumes that features are ____.

- A) Dependent
- B) Independent given the class
- C) Correlated by weight
- D) Always binary

Answer: B) Independent given the class

Explanation: Conditional independence simplifies Bayes rule for efficient probability estimation.

16. Overfitting occurs when a model ____.

- A) Fits training data too well and fails on new data
- B) Performs poorly on training data
- C) Is too simple
- D) Uses regularization

Answer: A) Fits training data too well and fails on new data

Explanation: Overfitting means the model captures noise instead of general patterns.

17. Which technique helps reduce overfitting?

- A) Increasing epochs indefinitely
- B) Regularization or Dropout
- C) Ignoring validation data
- D) Adding noise to labels

Answer: B) Regularization or Dropout

Explanation: These methods limit model complexity and improve generalization.

18. The bias-variance trade-off balances ____.

- A) Training and testing speed
- B) Underfitting and overfitting
- C) Accuracy and precision
- D) Data and features

Answer: B) Underfitting and overfitting

Explanation: High bias → underfit; high variance → overfit; good models balance both.

19. k-Nearest Neighbors (k-NN) relies on ____.

- A) Training parameters

- B) Distance metric between samples
- C) Hidden layers
- D) Label probabilities only

Answer: B) Distance metric between samples

Explanation: k-NN classifies based on majority labels among closest neighbors in feature space.

20. The kernel trick is used in which algorithm?

- A) Naïve Bayes
- B) Support Vector Machines
- C) Decision Trees
- D) k-Means

Answer: B) Support Vector Machines

Explanation: The kernel trick implicitly maps data to higher dimensions for non-linear classification.

Perfect ☐ Continuing Topic 1 – General AI Concepts (Q21 – Q100) in the same readable format:

21. Principal Component Analysis (PCA) is mainly used for ____.

- A) Classification
- B) Clustering
- C) Dimensionality reduction
- D) Reinforcement learning

Answer: C) Dimensionality reduction

Explanation: PCA reduces data to a smaller number of orthogonal components capturing the greatest variance.

22. k-Means algorithm minimizes ____.

- A) Entropy
- B) Within-cluster squared distances
- C) Cross-entropy loss
- D) Information gain

Answer: B) Within-cluster squared distances

Explanation: k-Means updates cluster centroids to reduce the total intra-cluster variance.

23. The EM algorithm alternates between ____.

- A) Expectation and Maximization steps
- B) Exploration and exploitation
- C) Training and validation
- D) Expansion and pruning

Answer: A) Expectation and Maximization steps

Explanation: EM estimates hidden variables (E-step) and updates parameters (M-step) iteratively.

24. In a Hidden Markov Model, the Forward algorithm computes ____.

- A) Most likely sequence of states
- B) Probability of an observation sequence
- C) Maximum likelihood parameters
- D) Transition matrix

Answer: B) Probability of an observation sequence

Explanation: The forward procedure sums probabilities over all possible hidden state paths.

25. Q-Learning is a form of ____.

- A) Supervised learning
- B) Reinforcement learning
- C) Unsupervised learning
- D) Semi-supervised learning

Answer: B) Reinforcement learning

Explanation: Q-Learning learns optimal action values through reward feedback without requiring a model.

26. The Bellman equation expresses ____.

- A) Relationship between policy and reward
- B) Recursive definition of state values
- C) Gradient update in neural nets
- D) Statistical independence

Answer: B) Recursive definition of state values

Explanation: $\text{Value} = \text{Immediate reward} + \text{discount} \times \text{expected future value}$.

27. ReLU activation helps to reduce ____.

- A) Exploding gradient
- B) Vanishing gradient
- C) Dropout rate
- D) Overfitting directly

Answer: B) Vanishing gradient

Explanation: ReLU keeps positive gradients constant, avoiding vanishing seen in sigmoids.

28. Backpropagation is used to ____.

- A) Propagate inputs
- B) Compute weight gradients efficiently
- C) Store patterns
- D) Normalize data

Answer: B) Compute weight gradients efficiently

Explanation: It applies the chain rule to update network weights during training.

29. CNNs exploit ____.

- A) Temporal recurrence
- B) Local spatial correlations
- C) Symbolic reasoning
- D) Random noise

Answer: B) Local spatial correlations

Explanation: Convolutions capture nearby pixel dependencies using shared weights.

30. Pooling in CNNs provides ____.

- A) Temporal memory

- B) Translation invariance and down-sampling
- C) Higher resolution
- D) Weight normalization

Answer: B) Translation invariance and down-sampling

Explanation: Pooling reduces spatial size and highlights most relevant activations.

31. LSTM networks are designed to ____.

- A) Classify images
- B) Handle long-term dependencies in sequences
- C) Cluster features
- D) Reduce dimensions

Answer: B) Handle long-term dependencies in sequences

Explanation: LSTMs use gates to maintain and forget information across time steps.

32. Attention mechanisms allow a model to ____.

- A) Ignore irrelevant data
- B) Focus on relevant input parts adaptively
- C) Compress entire sequence
- D) Eliminate recurrence

Answer: B) Focus on relevant input parts adaptively

Explanation: Attention weights emphasize important features for prediction.

33. Transformers replace recurrence with ____.

- A) Convolutions
- B) Self-attention
- C) Pooling layers
- D) Residual connections only

Answer: B) Self-attention

Explanation: Transformers capture relationships between all tokens simultaneously via attention.

34. Word2Vec learns embeddings by ____.

- A) Matrix factorization
- B) Predicting context words or target words
- C) Random initialization
- D) Clustering

Answer: B) Predicting context words or target words

Explanation: Skip-Gram and CBOW train networks to predict surrounding words in a sentence.

35. TF-IDF down-weights words that are ____.

- A) Rare in a document
- B) Common across many documents
- C) Highly informative
- D) Stop-listed

Answer: B) Common across many documents

Explanation: TF-IDF emphasizes words that are frequent in one document but rare overall.

36. BERT is a ____.

- A) Decoder-only autoregressive model
- B) Bidirectional masked language model
- C) Recurrent sequence model
- D) Vision transformer

Answer: B) Bidirectional masked language model

Explanation: BERT learns deep bidirectional context by predicting masked tokens.

37. GPT models are trained with ____.

- A) Next-token prediction
- B) Masked token prediction
- C) Contrastive loss
- D) Autoencoder objective

Answer: A) Next-token prediction

Explanation: GPT uses an autoregressive objective to predict the next word sequentially.

38. A decision tree split aims to ____.

- A) Reduce dimensionality
- B) Increase node purity
- C) Increase entropy
- D) Reduce bias

Answer: B) Increase node purity

Explanation: Splits minimize impurity metrics like Gini index or entropy.

39. Random Forest combines ____.

- A) Neural layers
- B) Multiple decision trees on bootstrapped samples
- C) Linear regressors
- D) Clustering centroids

Answer: B) Multiple decision trees on bootstrapped samples

Explanation: Random Forest averages many diverse trees to reduce variance.

40. Boosting builds models ____.

- A) Independently in parallel
- B) Sequentially, correcting previous errors
- C) Randomly
- D) Using gradient descent only

Answer: B) Sequentially, correcting previous errors

Explanation: Boosting (e.g., AdaBoost, XGBoost) trains each weak learner to fix prior mistakes.

41. Dropout prevents ____.

- A) Gradient vanishing
- B) Overfitting by randomly disabling neurons
- C) Training slowdown
- D) Underfitting

Answer: B) Overfitting by randomly disabling neurons

Explanation: It forces the network not to rely on specific units during training.

42. The Softmax function outputs ____.

- A) Feature maps
- B) Normalized probabilities
- C) Binary decisions
- D) Loss values

Answer: B) Normalized probabilities

Explanation: Softmax converts logits into class probabilities that sum to 1.

43. Cross-entropy loss is used for ____.

- A) Regression tasks
- B) Classification tasks
- C) Clustering
- D) PCA

Answer: B) Classification tasks

Explanation: It measures divergence between predicted probabilities and true one-hot labels.

44. Precision is defined as ____.

- A) $TP / (TP + FN)$
- B) $TP / (TP + FP)$
- C) $TN / (TN + FP)$
- D) $TP / (TP + TN)$

Answer: B) $TP / (TP + FP)$

Explanation: Precision measures correctness of positive predictions.

45. Recall (Sensitivity) is ____.

- A) $TP / (TP + FP)$
- B) $TP / (TP + FN)$
- C) $TN / (TN + FP)$
- D) $FP / (FP + TN)$

Answer: B) $TP / (TP + FN)$

Explanation: Recall quantifies how many actual positives were correctly identified.

46. ROC-AUC evaluates ____.

- A) Regression fit
- B) Trade-off between true-positive and false-positive rates
- C) Data imbalance
- D) Overfitting

Answer: B) Trade-off between TPR and FPR

Explanation: The area under the ROC curve measures overall classifier discrimination.

47. In supervised learning, labeled data contain ____.

- A) Inputs only
- B) Inputs and known outputs
- C) Outputs only
- D) Random noise

Answer: B) Inputs and known outputs

Explanation: Models learn mappings from examples of input-output pairs.

48. Unsupervised learning uses ____.

- A) Labeled data
- B) Unlabeled data
- C) Partially labeled data
- D) Reward signals

Answer: B) Unlabeled data

Explanation: It discovers hidden structures (clusters, features) without predefined labels.

49. Semi-supervised learning uses ____.

- A) Only labeled data

- B) Only unlabeled data
- C) Both labeled and unlabeled data
- D) Synthetic data only

Answer: C) Both labeled and unlabeled data

Explanation: It leverages small labeled sets and large unlabeled sets for better generalization.

50. Reinforcement learning differs because it ____.

- A) Learns from labeled examples
- B) Uses rewards and punishments
- C) Is unsupervised
- D) Is rule-based

Answer: B) Uses rewards and punishments

Explanation: Agents learn optimal actions through trial-and-error interactions with an environment.

51. The agent's goal in reinforcement learning is to maximize ____.

- A) Immediate reward only
- B) Cumulative discounted reward
- C) Prediction accuracy
- D) Entropy

Answer: B) Cumulative discounted reward

Explanation: Agents aim for long-term benefit, balancing short- and long-term returns.

52. The value function estimates ____.

- A) State transition probabilities
- B) Expected return from a state
- C) Reward variance
- D) Learning rate

Answer: B) Expected return from a state

Explanation: It measures how good a state is in terms of expected future rewards.

53. Policy gradient methods optimize ____.

- A) Action-value table
- B) Parameters of the policy directly
- C) Q-values indirectly
- D) Reward function

Answer: B) Parameters of the policy directly

Explanation: They adjust policy parameters using gradients of expected reward.

54. Game-playing AI like AlphaGo combines ____.

- A) Search and deep learning
- B) Symbolic reasoning only
- C) Pure reinforcement rules
- D) Statistical sampling only

Answer: A) Search and deep learning

Explanation: AlphaGo used deep neural nets with Monte Carlo Tree Search.

55. The first AI winter occurred due to ____.

- A) Funding cuts after unmet expectations
- B) Lack of neural networks
- C) Quantum computing rise
- D) Hardware explosion

Answer: A) Funding cuts after unmet expectations

Explanation: Early symbolic AI failed to deliver promised results, causing reduced research funding.

56. Expert systems are built on ____.

- A) Neural network weights
- B) If-then rule bases
- C) Reinforcement agents
- D) Decision forests

Answer: B) If-then rule bases

Explanation: They emulate human experts using rules and inference engines.

57. The inference engine in an expert system ____.

- A) Stores facts
- B) Applies rules to derive conclusions
- C) Collects user data
- D) Displays outputs

Answer: B) Applies rules to derive conclusions

Explanation: It matches conditions in the knowledge base to infer new facts.

58. A knowledge base contains ____.

- A) User preferences
- B) Facts and rules
- C) Graph edges
- D) Sensor data

Answer: B) Facts and rules

Explanation: It's the core repository of domain knowledge used for reasoning.

59. Forward chaining starts from ____.

- A) Goals and moves backward
- B) Known facts and moves forward
- C) Heuristic guesses
- D) Random sampling

Answer: B) Known facts and moves forward

Explanation: It applies rules to known data to derive new information.

60. Backward chaining starts from ____.

- A) Facts
- B) Hypotheses or goals
- C) Sensors
- D) Rules only

Answer: B) Hypotheses or goals

Explanation: It works backward to check if facts support a desired conclusion.

Excellent ☐ — here's the continuation of Topic 1 - General AI Concepts (Q61-Q100) in the same clear "copy-paste" text format:

61. Knowledge representation can use which method?

- A) Frames and semantic networks
- B) Convolutional filters
- C) Recurrent loops
- D) Gradient descent

Answer: A) Frames and semantic networks

Explanation: Frames and semantic networks organize concepts and their relations for reasoning.

62. Ontologies in AI describe ____.

- A) Low-level data features
- B) The formal relationships between concepts
- C) Visual objects
- D) Neural connections

Answer: B) The formal relationships between concepts

Explanation: Ontologies define shared vocabularies and class hierarchies for knowledge exchange.

63. A production rule in expert systems has the form ____.

- A) Input → Output
- B) IF condition THEN action
- C) Data → Model
- D) Cause → Effect only

Answer: B) IF condition THEN action

Explanation: Production rules encode reasoning steps used by the inference engine.

64. A utility function in AI expresses ____.

- A) Execution time
- B) The desirability of states or outcomes
- C) Memory usage
- D) Probability of failure

Answer: B) The desirability of states or outcomes

Explanation: Utility quantifies preferences, enabling rational decision-making under uncertainty.

65. In game trees, the Minimax algorithm assumes ____.

- A) Random opponents
- B) Optimal play by both sides
- C) Greedy player only
- D) Infinite depth always

Answer: B) Optimal play by both sides

Explanation: Minimax computes moves that minimize the possible loss for a worst-case opponent.

66. Alpha-Beta pruning improves Minimax by ____.

- A) Changing the evaluation function
- B) Eliminating branches that cannot affect the final decision
- C) Doubling tree depth
- D) Averaging utilities

Answer: B) Eliminating branches that cannot affect the final decision

Explanation: It skips evaluation of moves that are provably irrelevant, improving efficiency.

67. A heuristic function provides ____.

- A) Exact cost from start to goal
- B) An estimate of cost to reach goal
- C) Random guidance
- D) Total path length

Answer: B) An estimate of cost to reach goal

Explanation: Heuristics guide informed search algorithms like A* toward promising paths.

68. Admissible heuristics guarantee ____.

- A) Fastest runtime
- B) Optimal solutions
- C) Minimum branching factor
- D) Minimal memory

Answer: B) Optimal solutions

Explanation: Because admissible heuristics never overestimate, they ensure A* finds the best path.

69. Constraint satisfaction problems (CSPs) can be solved by ____.

- A) Backtracking with heuristics
- B) Reinforcement learning
- C) Decision trees
- D) Gradient descent

Answer: A) Backtracking with heuristics

Explanation: CSP solvers use variable-ordering and consistency checks to prune invalid choices.

70. Arc consistency ensures that ____.

- A) Every variable has at least one consistent value
- B) Graph is connected
- C) Search tree is complete
- D) Path cost is minimal

Answer: A) Every variable has at least one consistent value

Explanation: Arc consistency removes inconsistent domain values in binary constraints.

71. Which data structure is used in BFS?

- A) Stack
- B) Queue
- C) Priority queue
- D) Tree

Answer: B) Queue

Explanation: BFS expands nodes in FIFO order for level-wise exploration.

72. Which data structure is used in DFS?

- A) Queue
- B) Stack
- C) Heap
- D) List

Answer: B) Stack

Explanation: DFS explores by pushing and popping nodes from a stack (LIFO order).

73. Greedy best-first search expands nodes based on ____.

- A) Depth level
- B) Path cost only
- C) Heuristic estimate $h(n)$
- D) Random choice

Answer: C) Heuristic estimate $h(n)$

Explanation: It selects nodes appearing closest to the goal according to the heuristic.

74. Uniform-cost search expands nodes by ____.

- A) Smallest path cost $g(n)$

- B) Highest heuristic
- C) Random order
- D) Breadth-level

Answer: A) Smallest path cost $g(n)$

Explanation: UCS guarantees optimality when step costs are positive.

75. The evaluation function $f(n)=g(n)+h(n)$ is used in ____.

- A) A* search
- B) Hill climbing
- C) Beam search
- D) Genetic algorithms

Answer: A) A* search

Explanation: A* combines actual cost (g) and heuristic estimate (h) to rank nodes.

76. Hill-climbing search may get stuck in ____.

- A) Global optima
- B) Local maxima or plateaus
- C) Infinite loops
- D) Heuristic updates

Answer: B) Local maxima or plateaus

Explanation: It only moves to better neighbors, so it can stop at suboptimal points.

77. Simulated annealing avoids local maxima by ____.

- A) Adding random noise to weights
- B) Accepting worse moves with decreasing probability
- C) Increasing temperature continuously
- D) Restarting search each time

Answer: B) Accepting worse moves with decreasing probability

Explanation: The “temperature” controls exploration, reducing over time.

78. Beam search differs from BFS because it ____.

- A) Expands all nodes
- B) Keeps only a fixed number of best candidates
- C) Works without heuristics
- D) Always finds optimal solution

Answer: B) Keeps only a fixed number of best candidates

Explanation: Beam width limits memory and computation using heuristic scores.

79. Genetic algorithms use ____.

- A) Backtracking

- B) Mutation, crossover, and selection
- C) Neural activations
- D) Logical inference

Answer: B) Mutation, crossover, and selection

Explanation: They simulate evolution to find optimal or near-optimal solutions.

80. Fitness function in genetic algorithms measures ____.

- A) Randomness
- B) Solution quality
- C) Diversity
- D) Reproduction count

Answer: B) Solution quality

Explanation: It determines how well each candidate meets problem objectives.

81. In fuzzy logic, a membership function defines ____.

- A) Exact truth values
- B) Degree of belonging of an element to a set
- C) Probability of outcome
- D) Weight initialization

Answer: B) Degree of belonging of an element to a set

Explanation: It maps input values to membership grades between 0 and 1.

82. Defuzzification converts ____.

- A) Crisp values to fuzzy
- B) Fuzzy results back to crisp outputs
- C) Probabilities to weights
- D) Boolean logic

Answer: B) Fuzzy results back to crisp outputs

Explanation: It produces actionable numeric values from fuzzy inference.

83. A knowledge graph consists of ____.

- A) Images and pixels
- B) Entities and relationships
- C) Programs and files
- D) Layers and nodes

Answer: B) Entities and relationships

Explanation: Knowledge graphs link data via nodes (entities) and edges (relations).

84. RDF and OWL are standards for ____.

- A) Deep learning

- B) Knowledge representation on the Semantic Web
- C) Image processing
- D) Data compression

Answer: B) Knowledge representation on the Semantic Web

Explanation: RDF and OWL define structured, machine-understandable data formats.

85. Machine perception refers to ____.

- A) High-level reasoning
- B) Sensing and interpreting environment information
- C) Logical deduction
- D) Text translation

Answer: B) Sensing and interpreting environment information

Explanation: It covers computer vision, speech, and other sensory modalities.

86. Natural Language Processing (NLP) involves ____.

- A) Image enhancement
- B) Understanding and generating human language
- C) Reinforcement signals
- D) Logical resolution

Answer: B) Understanding and generating human language

Explanation: NLP enables AI to process text and speech meaningfully.

87. Tokenization in NLP means ____.

- A) Assigning parts of speech
- B) Splitting text into words or subwords
- C) Translating languages
- D) Counting frequencies

Answer: B) Splitting text into words or subwords

Explanation: Tokenization prepares raw text for further linguistic analysis.

88. Named Entity Recognition (NER) identifies ____.

- A) Sentiment polarity
- B) People, places, organizations in text
- C) Part-of-speech tags
- D) Coreference chains

Answer: B) People, places, organizations in text

Explanation: NER locates and classifies proper nouns into predefined categories.

89. Sentiment analysis aims to ____.

- A) Detect emotions or opinions in text

- B) Parse syntax trees
- C) Generate embeddings
- D) Recognize entities

Answer: A) Detect emotions or opinions in text

Explanation: It classifies text as positive, negative, or neutral.

90. Word embeddings capture ____.

- A) Syntactic grammar only
- B) Semantic similarity between words
- C) Random distributions
- D) Logical rules

Answer: B) Semantic similarity between words

Explanation: Similar words have closer vector representations in embedding space.

91. Speech recognition converts ____.

- A) Text to speech
- B) Speech to text
- C) Text to numbers
- D) Images to text

Answer: B) Speech to text

Explanation: ASR (Automatic Speech Recognition) transcribes spoken input into text.

92. Computer vision tasks include ____.

- A) Image classification and object detection
- B) Text summarization
- C) Data encryption
- D) Audio synthesis

Answer: A) Image classification and object detection

Explanation: Vision models interpret visual data for recognition or segmentation.

93. Object detection outputs ____.

- A) Single labels
- B) Bounding boxes with class labels
- C) Sentiment scores
- D) Hidden states

Answer: B) Bounding boxes with class labels

Explanation: Detection identifies what and where objects appear in an image.

94. Semantic segmentation assigns ____.

- A) One label per image

- B) A class label to each pixel
- C) A color histogram
- D) An embedding vector

Answer: B) A class label to each pixel

Explanation: It provides dense per-pixel classification for visual understanding.

95. Reinforcement learning problems are modeled as ____.

- A) Markov Decision Processes (MDP)
- B) Hidden Markov Models
- C) Linear systems
- D) Regression tasks

Answer: A) Markov Decision Processes (MDP)

Explanation: MDPs define states, actions, rewards, and transitions for agent learning.

96. The exploration-exploitation dilemma concerns ____.

- A) Choosing between learning new actions and using known good ones
- B) Overfitting
- C) Regularization
- D) Random sampling only

Answer: A) Choosing between learning new actions and using known good ones

Explanation: Balancing discovery of new strategies and maximizing reward is key in RL.

97. The discount factor γ in RL determines ____.

- A) Learning rate
- B) Importance of future rewards
- C) Reward variance
- D) Exploration rate

Answer: B) Importance of future rewards

Explanation: Lower γ emphasizes immediate rewards; higher γ values long-term returns.

98. Ethics in AI mainly focus on ____.

- A) Computation speed
- B) Fairness, transparency, and accountability
- C) Optimization accuracy
- D) Dataset size

Answer: B) Fairness, transparency, and accountability

Explanation: Ethical AI seeks to avoid bias and ensure responsible deployment.

99. Explainable AI (XAI) aims to ____.

- A) Hide model decisions

- B) Make AI decisions understandable to humans
- C) Increase model depth
- D) Replace humans entirely

Answer: B) Make AI decisions understandable to humans

Explanation: XAI techniques reveal reasoning behind predictions for trust and compliance.

100. The ultimate goal of Artificial Intelligence is ____.

- A) Automating mathematical operations
- B) Creating systems that exhibit intelligent behavior
- C) Building faster computers
- D) Eliminating all human jobs

Answer: B) Creating systems that exhibit intelligent behavior

Explanation: AI seeks to design machines that can perceive, reason, learn, and act autonomously.

Related posts:

1. November 2023 Current affairs
2. On November 3, 1957, the Soviet Union launched Sputnik 2
3. Neil Wagner Retirement: A Tribute to New Zealand's Cricket Legend
4. Daily Current Affairs 18 October 2024

5. Current Affairs
6. Daily Current Affairs 19 October 2024
7. Daily Current Affairs 20 October 2024
8. Daily Current Affairs 21 October 2024
9. Daily Current Affairs 22 October 2024
10. Daily Current Affairs 22 October 2024
11. Daily Current Affairs 23 October 2024
12. Daily Current Affairs 24 October 2024
13. Daily Current Affairs 25 October 2024
14. Daily Current Affairs 26 October 2024
15. Fix ERR_TOO_MANY_REDIRECTS: Solve Redirect Loop in WordPress & Cloudflare
16. What is learned weight matrix ?