

1. What is the technique used to measure the similarity between data points in K-Means Algorithm?

- a) Euclidean Distance
- b) Manhattan Distance
- c) Mahalanobis Distance
- d) Cosine Similarity

Answer: A. Euclidean Distance

2. Which method is used to combine the outputs of multiple weak learners in Boosting?

- a) Weighted Sum
- b) Max Voting
- c) Average
- d) Majority Voting

Answer: A. Weighted Sum

3. What is the primary purpose of Bagging in ensemble learning?

- a) Reduce Variance
- b) Increase Bias
- c) Improve Interpretability
- d) Enhance Scalability

Answer: A. Reduce Variance

4. Which technique involves training multiple models on different subsets of the training data?

- a) Bootstrap Aggregating
- b) Dropout
- c) Regularization

d) Batch Normalization

Answer: A. Bootstrap Aggregating

5. What is the key characteristic of Ensemble Learning?

- a) Combination of Weak Learners
- b) Use of Single Strong Learner
- c) Use of Sequential Training
- d) Application of Reinforcement Learning

Answer: A. Combination of Weak Learners

6. Which ensemble learning technique places equal weights on the predictions of all classifiers?

- a) Majority Voting
- b) Weighted Voting
- c) Stacking
- d) Boosting

Answer: A. Majority Voting

7. What is the primary advantage of using ensemble methods in machine learning?

- a) Improved Performance
- b) Reduced Complexity
- c) Reduced Variance
- d) Increased Bias

Answer: A. Improved Performance

8. Which technique involves iteratively adjusting the weights of misclassified samples in Boosting?

- a) AdaBoost
- b) Gradient Boosting
- c) XGBoost
- d) LightGBM

Answer: A. AdaBoost

9. What is the primary objective of boosting algorithms?

- a) Sequentially Improve Model
- b) Reduce Model Complexity
- c) Increase Interpretability
- d) Decrease Bias

Answer: A. Sequentially Improve Model

10. Which algorithm constructs the decision boundary by fitting a sequence of weak learners?

- a) Gradient Boosting
- b) AdaBoost
- c) Random Forest
- d) Bagging

Answer: A. Gradient Boosting

11. What is the primary limitation of AdaBoost algorithm?

- a) Sensitive to Noisy Data
- b) Overfitting
- c) Requires Large Dataset
- d) Prone to Underfitting

Answer: A. Sensitive to Noisy Data

12. Which ensemble learning method trains multiple models independently and then combines their predictions?

- a) Bagging
- b) Boosting
- c) Stacking
- d) Random Forest

Answer: A. Bagging

13. What is the main advantage of using Random Forest over a single Decision Tree?

- a) Reduced Overfitting
- b) Increased Interpretability
- c) Faster Training Time
- d) Higher Complexity

Answer: A. Reduced Overfitting

14. Which technique is used to introduce randomness in Random Forests?

- a) Bootstrap Sampling
- b) Gradient Descent
- c) Back Propagation
- d) Feature Scaling

Answer: A. Bootstrap Sampling

15. What is the purpose of feature subsampling in Random Forests?

- a) Decorrelate Trees
- b) Increase Variance
- c) Decrease Bias
- d) Improve Interpretability

Answer: A. Decorrelate Trees

16. Which algorithm constructs multiple decision trees and combines their predictions through voting?

- a) Random Forest
- b) AdaBoost
- c) Gradient Boosting
- d) Bagging

Answer: A. Random Forest

17. What is the concept of feature bagging in Random Forests?

- a) Randomly selecting subsets of features
- b) Randomly selecting subsets of samples
- c) Randomly selecting subsets of trees
- d) Randomly selecting subsets of splits

Answer: A. Randomly selecting subsets of features

18. What is the key difference between Bagging and Boosting techniques?

- a) Weighting of Samples
- b) Weighting of Trees
- c) Weighting of Features
- d) Weighting of Classifiers

Answer: A. Weighting of Samples

19. Which algorithm constructs an ensemble of decision trees with different subsets of training data?

- a) Bagging
- b) Random Forest

- c) AdaBoost
- d) Gradient Boosting

Answer: A. Bagging

20. What is the main advantage of ensemble learning methods such as Bagging and Boosting?

- a) Improved Generalization
- b) Reduced Variance
- c) Decreased Complexity
- d) Increased Bias

Answer: A. Improved Generalization

21. Which ensemble method constructs decision trees sequentially with each tree correcting the errors of the previous one?

- a) AdaBoost
- b) Gradient Boosting
- c) Random Forest
- d) Bagging

Answer: A. AdaBoost

22. What is the primary disadvantage of ensemble learning methods like Bagging and Boosting?

- a) Increased Computational Complexity
- b) Decreased Model Performance
- c) Reduced Variance
- d) Reduced Interpretability

Answer: A. Increased Computational Complexity

23. Which technique is used to reduce the variance of individual decision trees in Random Forests?

- a) Feature Bagging
- b) Bootstrap Sampling
- c) Gradient Boosting
- d) AdaBoost

Answer: A. Feature Bagging

24. What is the main characteristic of ensemble learning methods like Bagging and Boosting?

- a) Diversity of Models
- b) Homogeneity of Models
- c) Single Model Approach
- d) Noisy Predictions

Answer: A. Diversity of Models

25. Which ensemble learning technique combines multiple weak learners into a single strong learner through a weighted sum of their predictions?

- a) Stacking
- b) AdaBoost
- c) Gradient Boosting
- d) Random Forest

Answer: A. Stacking

26. What is the main advantage of using Random Forest over a single Decision Tree?

- a) Reduced Overfitting
- b) Increased Interpretability
- c) Faster Training Time

d) Higher Complexity

Answer: A. Reduced Overfitting

27. Which technique is used to introduce randomness in Random Forests?

- a) Bootstrap Sampling
- b) Gradient Descent
- c) Back Propagation
- d) Feature Scaling

Answer: A. Bootstrap Sampling

28. What is the purpose of feature subsampling in Random Forests?

- a) Decorrelate Trees
- b) Increase Variance
- c) Decrease Bias
- d) Improve Interpretability

Answer: A. Decorrelate Trees

29. Which algorithm constructs multiple decision trees and combines their predictions through voting?

- a) Random Forest
- b) AdaBoost
- c) Gradient Boosting
- d) Bagging

Answer: A. Random Forest

30. What is the concept of feature bagging in Random Forests?

- a) Randomly selecting subsets of features
- b) Randomly selecting subsets of samples



- c) Randomly selecting subsets of trees
- d) Randomly selecting subsets of splits

Answer: A. Randomly selecting subsets of features

31. What is the key difference between Bagging and Boosting techniques?

- a) Weighting of Samples
- b) Weighting of Trees
- c) Weighting of Features
- d) Weighting of Classifiers

Answer: A. Weighting of Samples

32. Which algorithm constructs an ensemble of decision trees with different subsets of training data?

- a) Bagging
- b) Random Forest
- c) AdaBoost
- d) Gradient Boosting

Answer: A. Bagging

33. What is the main advantage of ensemble learning methods such as Bagging and Boosting?

- a) Improved Generalization
- b) Reduced Variance
- c) Decreased Complexity
- d) Increased Bias

Answer: A. Improved Generalization

34. Which ensemble method constructs decision trees sequentially with each tree correcting the errors of the previous one?

- a) AdaBoost
- b) Gradient Boosting
- c) Random Forest
- d) Bagging

Answer: A. AdaBoost

35. What is the primary disadvantage of ensemble learning methods like Bagging and Boosting?

- a) Increased Computational Complexity
- b) Decreased Model Performance
- c) Reduced Variance
- d) Reduced Interpretability

Answer: A. Increased Computational Complexity

36. Which technique is used to reduce the variance of individual decision trees in Random Forests?

- a) Feature Bagging
- b) Bootstrap Sampling
- c) Gradient Boosting
- d) AdaBoost

Answer: A. Feature Bagging

37. What is the main characteristic of ensemble learning methods like Bagging and Boosting?

- a) Diversity of Models
- b) Homogeneity of Models
- c) Single Model Approach

d) Noisy Predictions

Answer: A. Diversity of Models

38. Which ensemble learning technique combines multiple weak learners into a single strong learner through a weighted sum of their predictions?

- a) Stacking
- b) AdaBoost
- c) Gradient Boosting
- d) Random Forest

Answer: A. Stacking

39. What is the primary objective of Dimensionality Reduction techniques?

- a) Preserve Data Structure
- b) Reduce Model Complexity
- c) Increase Dimensionality
- d) Enhance Computational Speed

Answer: A. Preserve Data Structure

40. Which method is used to find the optimal projection to discriminate between classes in Linear Discriminant Analysis?

- a) Fisher Criterion
- b) PCA
- c) Factor Analysis
- d) ICA

Answer: A. Fisher Criterion

41. What is the main advantage of Principal Component Analysis (PCA)?

- a) Capture Maximum Variance

- b) Handles Non-Linearity
- c) Suitable for Categorical Data
- d) Requires Feature Engineering

Answer: A. Capture Maximum Variance

42. Which technique assumes that the observed data are generated by a linear combination of some underlying factors?

- a) Factor Analysis
- b) PCA
- c) LDA
- d) ICA

Answer: A. Factor Analysis

43. What is the purpose of Independent Component Analysis (ICA)?

- a) Extract Independent Sources
- b) Maximize Variance
- c) Minimize Dimensionality
- d) Preserve Data Structure

Answer: A. Extract Independent Sources

44. Which method focuses on preserving the local structure of the data during dimensionality reduction?

- a) Locally Linear Embedding
- b) PCA
- c) Factor Analysis
- d) Independent Component Analysis

Answer: A. Locally Linear Embedding

45. What is the primary objective of Isomap algorithm?

- a) Preserve Geodesic Distances
- b) Maximize Variance
- c) Minimize Reconstruction Error
- d) Reduce Computational Complexity

Answer: A. Preserve Geodesic Distances

46. Which optimization technique is commonly used in Locally Linear Embedding (LLE)?

- a) Least Squares Optimization
- b) Gradient Descent
- c) Random Search
- d) Evolutionary Algorithms

Answer: A. Least Squares Optimization

47. What is the key concept behind Evolutionary Learning?

- a) Mimicking Natural Selection
- b) Gradient Descent
- c) Random Search
- d) Pattern Recognition

Answer: A. Mimicking Natural Selection

48. Which technique in Evolutionary Learning selects the best individuals for reproduction based on their fitness?

- a) Genetic Algorithms
- b) Neural Networks
- c) Decision Trees
- d) Support Vector Machines

Answer: A. Genetic Algorithms

49. What is the term for the offspring produced by combining genetic material from two parent individuals?

- a) Genetic Offspring
- b) Generational Offspring
- c) Evolutionary Offspring
- d) Parental Offspring

Answer: A. Genetic Offspring

50. Which operators are used to modify the genetic material during the reproduction process in Genetic Algorithms?

- a) Genetic Operators
- b) Logical Operators
- c) Arithmetic Operators
- d) Boolean Operators

Answer: A. Genetic Operators

51. Which operator in Genetic Algorithms is responsible for exploring new regions of the search space?

- a) Mutation
- b) Crossover
- c) Selection
- d) Elitism

Answer: A. Mutation

52. What is the primary purpose of Genetic Algorithms?

- a) Optimization

- b) Classification
- c) Clustering
- d) Regression

Answer: A. Optimization

53. Which technique in Genetic Algorithms selects individuals for reproduction based on their fitness and probability of selection?

- a) Selection
- b) Mutation
- c) Crossover
- d) Elitism

Answer: A. Selection

54. What is the role of crossover operator in Genetic Algorithms?

- a) Exchange Genetic Material
- b) Introduce Diversity
- c) Improve Fitness
- d) Explore New Regions

Answer: A. Exchange Genetic Material

55. Which phase of Genetic Algorithms evaluates the fitness of individuals in the population?

- a) Evaluation
- b) Initialization
- c) Selection
- d) Mutation

Answer: A. Evaluation

56. What is the main advantage of Genetic Algorithms over traditional optimization techniques?

- a) Global Optimization
- b) Local Optimization
- c) Fast Convergence
- d) Less Computational Complexity

Answer: A. Global Optimization

57. Which technique in Genetic Algorithms retains the best-performing individuals in the population for the next generation?

- a) Elitism
- b) Selection
- c) Mutation
- d) Crossover

Answer: A. Elitism

58. What is the primary disadvantage of Genetic Algorithms?

- a) Computational Complexity
- b) Slow Convergence
- c) Lack of Parallelism
- d) Limited Application

Answer: A. Computational Complexity

59. Which phase of Genetic Algorithms generates an initial population of individuals for optimization?

- a) Initialization
- b) Evaluation
- c) Selection



d) Crossover

Answer: A. Initialization

60. What is the primary objective of crossover operator in Genetic Algorithms?

- a) Explore New Solutions
- b) Exploit Existing Solutions
- c) Improve Diversity
- d) Improve Fitness

Answer: A. Explore New Solutions

61. Which type of problem is Genetic Algorithms best suited for?

- a) Optimization
- b) Classification
- c) Clustering
- d) Regression

Answer: A. Optimization

62. What is the main challenge of using Genetic Algorithms?

- a) Computational Complexity
- b) Slow Convergence
- c) Lack of Diversity
- d) Limited Search Space

Answer: A. Computational Complexity

63. Which operator in Genetic Algorithms is responsible for combining genetic material from parent individuals?

- a) Crossover
- b) Mutation

- c) Selection
- d) Elitism

Answer: A. Crossover

64. What is the primary advantage of using Evolutionary Learning techniques?

- a) Adaptability
- b) Scalability
- c) Low Complexity
- d) High Interpretability

Answer: A. Adaptability

65. Which technique in Genetic Algorithms introduces random changes in the genetic material of individuals?

- a) Mutation
- b) Crossover
- c) Selection
- d) Elitism

Answer: A. Mutation

66. What is the role of fitness function in Genetic Algorithms?

- a) Measure Individual Performance
- b) Select Parents
- c) Create Offspring
- d) Evaluate Solutions

Answer: A. Measure Individual Performance

67. Which phase of Genetic Algorithms selects individuals from the current population for reproduction?

- a) Selection
- b) Initialization
- c) Evaluation
- d) Crossover

Answer: A. Selection

68. What is the main objective of Genetic Algorithms in optimization tasks?

- a) Find Optimal Solutions
- b) Minimize Error
- c) Maximize Accuracy
- d) Optimize Parameters

Answer: A. Find Optimal Solutions

69. Which technique in Genetic Algorithms allows the best-performing individuals to survive from one generation to the next?

- a) Elitism
- b) Selection
- c) Mutation
- d) Crossover

Answer: A. Elitism

70. What is the main challenge of using Genetic Algorithms in optimization tasks?

- a) Computational Complexity
- b) Lack of Guarantee
- c) Limited Convergence
- d) Sensitivity to Noise

Answer: A. Computational Complexity

71. What is the primary role of crossover operator in Genetic Algorithms?

- a) Exchange Genetic Material
- b) Introduce Diversity
- c) Improve Fitness
- d) Explore New Regions

Answer: A. Exchange Genetic Material

72. Which phase of Genetic Algorithms generates an initial population of individuals for optimization?

- a) Initialization
- b) Evaluation
- c) Selection
- d) Crossover

Answer: A. Initialization

73. What is the primary objective of crossover operator in Genetic Algorithms?

- a) Explore New Solutions
- b) Exploit Existing Solutions
- c) Improve Diversity
- d) Improve Fitness

Answer: A. Explore New Solutions

74. Which technique in Genetic Algorithms allows the best-performing individuals to survive from one generation to the next?

- a) Elitism
- b) Selection
- c) Mutation

d) Crossover

Answer: A. Elitism

75. What is the primary disadvantage of using Genetic Algorithms?

- a) Computational Complexity
- b) Slow Convergence
- c) Lack of Diversity
- d) Limited Search Space

Answer: A. Computational Complexity

76. What is the primary objective of Reinforcement Learning?

- a) Learn Optimal Actions
- b) Minimize Prediction Error
- c) Maximize Likelihood
- d) Reduce Dimensionality

Answer: A. Learn Optimal Actions

77. Which technique in Reinforcement Learning involves learning from trial and error by interacting with an environment?

- a) Q-Learning
- b) Supervised Learning
- c) Unsupervised Learning
- d) Semi-Supervised Learning

Answer: A. Q-Learning

78. What is the role of the reward signal in Reinforcement Learning?

- a) Provides Feedback
- b) Defines the Policy

- c) Shapes the Environment
- d) Determines the Model

Answer: A. Provides Feedback

79. Which algorithm in Reinforcement Learning aims to estimate the value of being in a particular state and taking a specific action?

- a) Q-Learning
- b) SARSA
- c) Policy Gradient Methods
- d) Deep Q-Networks

Answer: A. Q-Learning

80. What is the term for the function that maps states to the probability distribution over actions in Reinforcement Learning?

- a) Policy
- b) Value Function
- c) Reward Function
- d) Transition Probability

Answer: A. Policy

81. What is the main challenge of Reinforcement Learning in complex environments?

- a) Exploration vs Exploitation
- b) Model Complexity
- c) Data Availability
- d) Computational Cost

Answer: A. Exploration vs Exploitation

82. Which Reinforcement Learning algorithm updates the Q-values based on the observed rewards and the maximum Q-value of the next state?

- a) SARSA
- b) Q-Learning
- c) Policy Gradient Methods
- d) Deep Q-Networks

Answer: A. SARSA

83. What is the primary difference between Q-Learning and SARSA algorithms?

- a) Off-Policy vs On-Policy
- b) Value vs Policy
- c) Deterministic vs Stochastic
- d) Exploration vs Exploitation

Answer: A. Off-Policy vs On-Policy

84. What is the primary objective of Markov Chain Monte Carlo (MCMC) methods?

- a) Sample from Complex Distributions
- b) Optimize Objective Function
- c) Reduce Dimensionality
- d) Estimate Uncertainty

Answer: A. Sample from Complex Distributions

85. Which technique is used in Markov Chain Monte Carlo (MCMC) to generate proposals for the next state?

- a) Proposal Distribution
- b) Reward Function
- c) Policy

d) Transition Probability

Answer: A. Proposal Distribution

86. What is the primary advantage of using Markov Chain Monte Carlo (MCMC) methods for sampling?

- a) Suitable for Complex Distributions
- b) Fast Convergence
- c) Global Optimization
- d) Low Computational Cost

Answer: A. Suitable for Complex Distributions

87. Which method in Markov Chain Monte Carlo (MCMC) iteratively constructs a Markov Chain to sample from a target distribution?

- a) Metropolis-Hastings
- b) Gradient Descent
- c) Back Propagation
- d) Evolutionary Algorithms

Answer: A. Metropolis-Hastings

88. What is the role of graphical models in representing probabilistic relationships between variables?

- a) Capture Dependencies
- b) Model Non-Linearity
- c) Reduce Dimensionality
- d) Optimize Objective

Answer: A. Capture Dependencies

89. Which type of graphical model represents conditional dependencies between variables through a directed acyclic graph?



- a) Bayesian Networks
- b) Markov Random Fields
- c) Hidden Markov Models
- d) Isomap

Answer: A. Bayesian Networks

90. What is the primary objective of Bayesian Networks?

- a) Represent Conditional Dependencies
- b) Model Temporal Data
- c) Capture Latent Variables
- d) Estimate Uncertainty

Answer: A. Represent Conditional Dependencies

91. Which type of graphical model represents pairwise relationships between variables through an undirected graph?

- a) Markov Random Fields
- b) Bayesian Networks
- c) Hidden Markov Models
- d) Locally Linear Embedding

Answer: A. Markov Random Fields

92. Which type of graphical model is commonly used for image segmentation tasks?

- a) Markov Random Fields
- b) Bayesian Networks
- c) Hidden Markov Models
- d) Reinforcement Learning

Answer: A. Markov Random Fields

93. What is the primary characteristic of Hidden Markov Models (HMMs)?

- a) Temporal Dynamics
- b) Spatial Relationships
- c) Non-Stationarity
- d) Directed Acyclic Graph

Answer: A. Temporal Dynamics

94. Which method is used to estimate the hidden states in Hidden Markov Models (HMMs) given observed data?

- a) Forward-Backward Algorithm
- b) Gradient Descent
- c) Back Propagation
- d) Evolutionary Algorithms

Answer: A. Forward-Backward Algorithm

95. What is the primary application of Hidden Markov Models (HMMs) in machine learning?

- a) Sequential Data Modeling
- b) Image Classification
- c) Regression Analysis
- d) Anomaly Detection

Answer: A. Sequential Data Modeling

96. Which method is used to update the probability distribution over states in Hidden Markov Models (HMMs) given observed data?

- a) Forward Algorithm
- b) Backward Algorithm

- c) Baum-Welch Algorithm
- d) Viterbi Algorithm

Answer: A. Forward Algorithm

97. What is the main challenge of using Hidden Markov Models (HMMs) in practice?

- a) Choosing Appropriate States
- b) Estimating Transition Probabilities
- c) Handling Missing Data
- d) Interpreting Model Outputs

Answer: A. Choosing Appropriate States

98. Which method in Reinforcement Learning deals with the problem of finding the optimal policy in Markov Decision Processes?

- a) Dynamic Programming
- b) Q-Learning
- c) SARSA
- d) Policy Gradient Methods

Answer: A. Dynamic Programming

99. What is the primary objective of Markov Chain Monte Carlo (MCMC) methods in machine learning?

- a) Sampling from Complex Distributions
- b) Optimization
- c) Model Interpretability
- d) Clustering

Answer: A. Sampling from Complex Distributions

100. Which method is commonly used in Reinforcement Learning to handle sequential decision-making problems with uncertainty?

- a) Markov Decision Processes
- b) Support Vector Machines
- c) K-Nearest Neighbors
- d) Random Forest

Answer: A. Markov Decision Processes

101. Which technique in Reinforcement Learning involves modeling the environment as a Markov Decision Process (MDP)?

- a) Markov Decision Processes
- b) Q-Learning
- c) SARSA
- d) Policy Gradient Methods

Answer: A. Markov Decision Processes

102. What is the primary benefit of using Markov Chain Monte Carlo (MCMC) methods in Bayesian inference?

- a) Approximate Posterior Sampling
- b) Exact Posterior Sampling
- c) Variational Inference
- d) Maximum Likelihood Estimation

Answer: A. Approximate Posterior Sampling

103. Which method in Reinforcement Learning is used for estimating the value function of each state-action pair?

- a) Q-Learning
- b) SARSA
- c) Dynamic Programming

d) Policy Gradient Methods

Answer: A. Q-Learning

104. What is the primary limitation of Markov Chain Monte Carlo (MCMC) methods?

- a) Computational Cost
- b) Limited Convergence
- c) Lack of Parallelism
- d) Difficulty in Implementation

Answer: A. Computational Cost

105. Which method in Reinforcement Learning focuses on learning a policy directly without explicitly estimating value functions?

- a) Policy Gradient Methods
- b) Q-Learning
- c) SARSA
- d) Dynamic Programming

Answer: A. Policy Gradient Methods

106. What is the key characteristic of Markov Chain Monte Carlo (MCMC) methods in sampling from complex distributions?

- a) Markov Chain Property
- b) Stochastic Gradient Descent
- c) Model Flexibility
- d) Reduced Variance

Answer: A. Markov Chain Property

107. Which method in Reinforcement Learning learns the value function by bootstrapping from the estimates of subsequent states?

- a) SARSA
- b) Q-Learning
- c) Dynamic Programming
- d) Policy Gradient Methods

Answer: A. SARSA

108. What is the primary advantage of Policy Gradient Methods in Reinforcement Learning?

- a) Suitable for Continuous Action Spaces
- b) Guaranteed Convergence
- c) Reduced Variance
- d) Low Computational Cost

Answer: A. Suitable for Continuous Action Spaces

109. Which technique in Reinforcement Learning focuses on directly optimizing the policy with respect to the expected return?

- a) Policy Gradient Methods
- b) Q-Learning
- c) SARSA
- d) Dynamic Programming

Answer: A. Policy Gradient Methods

110. What is the main advantage of Markov Chain Monte Carlo (MCMC) methods over traditional Monte Carlo methods?

- a) Handles High-Dimensional Spaces
- b) Faster Convergence
- c) Lower Computational Cost
- d) Exact Sampling

Answer: A. Handles High-Dimensional Spaces

111. Which method in Reinforcement Learning is suitable for environments with continuous state and action spaces?

- a) Policy Gradient Methods
- b) Q-Learning
- c) SARSA
- d) Dynamic Programming

Answer: A. Policy Gradient Methods

112. What is the primary challenge of Policy Gradient Methods in Reinforcement Learning?

- a) High Variance
- b) Model Complexity
- c) Lack of Exploration
- d) Slow Convergence

Answer: A. High Variance

113. Which technique in Reinforcement Learning is prone to the issue of "credit assignment" in long-term rewards?

- a) Q-Learning
- b) SARSA
- c) Dynamic Programming
- d) Policy Gradient Methods

Answer: A. Q-Learning

114. What is the primary benefit of using Policy Gradient Methods in Reinforcement Learning?

- a) Handles Continuous Action Spaces

- b) Exact Learning
- c) Model Interpretability
- d) Low Computational Cost

Answer: A. Handles Continuous Action Spaces

115. Which method in Reinforcement Learning learns the value function by bootstrapping from the estimate of the subsequent state with the highest value?

- a) Q-Learning
- b) SARSA
- c) Dynamic Programming
- d) Policy Gradient Methods

Answer: A. Q-Learning

116. What is the primary limitation of Policy Gradient Methods in Reinforcement Learning?

- a) High Variance
- b) Slow Convergence
- c) Model Complexity
- d) Lack of Exploration

Answer: A. High Variance

117. Which Reinforcement Learning algorithm updates the policy parameters in the direction of the gradient of the expected return?

- a) Policy Gradient Methods
- b) Q-Learning
- c) SARSA
- d) Dynamic Programming

Answer: A. Policy Gradient Methods



118. What is the primary characteristic of Q-Learning in Reinforcement Learning?

- a) Off-Policy Learning
- b) On-Policy Learning
- c) Direct Policy Optimization
- d) Temporal-Difference Learning

Answer: A. Off-Policy Learning

119. Which method in Reinforcement Learning learns the value function by iteratively updating Q-values based on observed rewards and estimated future rewards?

- a) Q-Learning
- b) SARSA
- c) Dynamic Programming
- d) Policy Gradient Methods

Answer: A. Q-Learning

120. What is the main advantage of SARSA algorithm in Reinforcement Learning?

- a) On-Policy Learning
- b) Off-Policy Learning
- c) Model Flexibility
- d) Low Computational Cost

Answer: A. On-Policy Learning

121. Which Reinforcement Learning algorithm updates the policy based on the action-value function learned from the current policy?

- a) SARSA

- b) Q-Learning
- c) Dynamic Programming
- d) Policy Gradient Methods

Answer: A. SARSA

122. What is the primary characteristic of SARSA algorithm in Reinforcement Learning?

- a) On-Policy Learning
- b) Off-Policy Learning
- c) Direct Policy Optimization
- d) Temporal-Difference Learning

Answer: A. On-Policy Learning

123. Which method in Reinforcement Learning learns the value function by considering the value of the subsequent state chosen by the current policy?

- a) SARSA
- b) Q-Learning
- c) Dynamic Programming
- d) Policy Gradient Methods

Answer: A. SARSA

124. What is the main advantage of Dynamic Programming in Reinforcement Learning?

- a) Finds Optimal Policy
- b) Suitable for Continuous State Spaces
- c) Low Variance
- d) Handles Non-Stationarity

Answer: A. Finds Optimal Policy

125. Which method in Reinforcement Learning is commonly used for planning in deterministic environments with known transition dynamics?

- a) Dynamic Programming
- b) Q-Learning
- c) SARSA
- d) Policy Gradient Methods

Answer: A. Dynamic Programming

