

#1. In reinforcement learning, what is the role of the exploration-exploitation tradeoff?

- ☐ Balancing between trying out new actions and exploiting known high-reward actions
- ☐ Balancing between the number of features and the number of samples
- ☐ Balancing between training time and testing time
- ☐ Balancing between underfitting and overfitting
- ☐ Balancing between model complexity and the amount of training data

#2. Which algorithm is commonly used for sentiment analysis in natural language processing (NLP)?

- ☐ Long Short-Term Memory (LSTM)
- ☐ Naive Bayes
- ☐ Logistic Regression
- ☐ Support Vector Machines (SVM)
- ☐ Random Forest

#3. What is the purpose of early stopping in neural network training?

- ☐ To prevent overfitting by stopping training when validation performance plateaus
- ☐

To increase model complexity

☐

To speed up the training process

☐

To add noise to the data

☐

To improve the interpretability of the model

#4. Which technique is used for reducing the dimensionality of high-dimensional data while preserving as much information as possible?

☐

t-SNE (t-Distributed Stochastic Neighbor Embedding)

☐

Random Projection

☐

Singular Value Decomposition (SVD)

☐

Apriori algorithm

☐

K-Means Clustering

#5. What is the main purpose of a Recurrent Neural Network (RNN) in deep learning?

☐

Handling sequential or time-series data

☐

Performing image classification

☐

Reducing model complexity

☐

Increasing the learning rate

☐

Adding non-linearity to the model

#6. Which method is commonly used for hyperparameter tuning in machine learning?

☐

Grid Search

☐

Random Search

☐

Gradient Descent

☐

K-Nearest Neighbors (KNN)

☐

Apriori algorithm

#7. What is the purpose of a one-hot encoding in preprocessing categorical data?

☐

Representing categorical variables as binary vectors

☐

Reducing the number of features

☐

Speeding up the training process

☐

Increasing the learning rate

☐

Adding noise to the data

#8. Which technique is used for outlier detection in a dataset?

☐

Isolation Forest

☐

K-Means Clustering

☐

PCA

☐

Linear Regression

☐

Logistic Regression

#9. What is the main objective of the Mean-Shift clustering algorithm?

☐

Identifying dense regions of data points in feature space

☐

Reducing model complexity

☐

Performing image segmentation

☐

Classifying data into predefined categories

☐

Increasing the learning rate

#10. What is the purpose of a loss function in machine learning?

☐

Quantifying the error between predicted and actual values

☐

Determining the number of hidden layers in a neural network

☐

Controlling the number of features in a dataset

☐

Adding non-linearity to the model

☐

Speeding up the training process

#11. In reinforcement learning, what is the role of the discount factor (gamma)?

☐

Balancing immediate rewards against future rewards

☐

Controlling the exploration rate

☐

Reducing model complexity

☐

Controlling the learning rate

☐

Controlling the number of episodes in training

#12. Which method is commonly used for imbalanced classification tasks?

☐

Resampling (e.g., oversampling or undersampling)

☐

Bagging

☐

Feature selection

☐

Principal Component Analysis (PCA)

☐

Random Projection

#13. What is the purpose of the Kullback-Leibler (KL) divergence in information theory?

☐

Measuring the difference between two probability distributions

☐

Classifying data into predefined categories

☐

Speeding up the training process

☐

Reducing model complexity

☐

Adding noise to the data

#14. Which algorithm is commonly used for face recognition in computer vision applications?

☐

Eigenfaces (PCA-based)

☐

Random Forest

☐

Logistic Regression

☐

K-Nearest Neighbors (KNN)

☐

Naive Bayes

#15. What is the role of the latent space in generative models like Variational Autoencoders (VAEs)?

☐

Encoding and decoding high-dimensional data into a lower-dimensional representation

☐

Controlling the learning rate

☐

Reducing model complexity

☐

Adding non-linearity to the model

☐

Performing image classification

#16. Which algorithm is commonly used for outlier detection based on distance measures?

☐

Mahalanobis Distance

☐

Random Forest

☐

Logistic Regression

☐

K-Means Clustering

☐

Naive Bayes

#17. What is the purpose of the term “momentum” in gradient descent optimization?

☐

To help accelerate convergence and escape local minima

☐

To reduce the learning rate

☐

To add noise to the data

☐

To increase the number of features

☐

To regularize the model

#18. Which technique is used for feature extraction in natural language processing (NLP)?

☐

Word Embeddings (e.g., Word2Vec, GloVe)

☐

Principal Component Analysis (PCA)

☐

Support Vector Machines (SVM)

☐

K-Means Clustering

☐

Decision Trees

#19. What is the main purpose of a Recurrent Neural Network (RNN) in deep learning?

☐

Handling sequential or time-series data

☐

Performing image classification

☐

Reducing model complexity

☐

Increasing the learning rate

☐

Adding non-linearity to the model

#20. In reinforcement learning, what is the role of the exploration-exploitation tradeoff?

☐

Balancing between trying out new actions and exploiting known high-reward actions

☐

Balancing between the number of features and the number of samples

☐

Balancing between training time and testing time

☐

Balancing between underfitting and overfitting

☐

Balancing between model complexity and the amount of training data

Next

Results



