1. What is Supervised Learning?

Ans. Supervised learning is a type of machine learning where the algorithm learns from labeled data, making predictions or decisions based on input features.

2. What is the difference between Regression and Classification?

Ans. Regression predicts continuous numerical values, while classification predicts categorical labels or classes.

3. What is Overfitting?

Ans. Overfitting occurs when a model learns the training data too well, capturing noise and random fluctuations. It performs poorly on new, unseen data.

4. What is a Hyperparameter?

Ans. A hyperparameter is a configuration setting of a model that is set before training and remains constant during training.

5. Give an example of a Classification Algorithm.

Ans. Logistic Regression is an example of a classification algorithm, used to predict binary outcomes.

6. Explain Bias-Variance Tradeoff.

Ans. The bias-variance tradeoff is the balance between a model's simplicity (bias) and its

ability to fit diverse data patterns (variance). Finding the right balance minimizes prediction errors.

7. What is Cross-Validation?

Ans. Cross-validation is a technique used to assess the performance of a model. The dataset is divided into subsets for training and testing, allowing for more robust evaluation.

8. Why do we use a Test Set?

Ans. The test set is used to evaluate the model's performance on data it has never seen before, providing an unbiased estimate of its predictive power.

9. What is Feature Engineering?

Ans. Feature engineering involves creating or transforming features from raw data to improve a model's performance.

10. Explain Precision in Classification.

Ans. Precision is the ratio of true positive predictions to the total predicted positives. It measures the accuracy of positive predictions.

11. Define Recall in Classification.

Ans. Recall (Sensitivity) is the ratio of true positive predictions to the total actual positives. It measures the ability of the model to identify all relevant cases.

12. What is the F1-Score?

Ans. The F1-Score is the harmonic mean of precision and recall. It provides a balanced measure of a model's performance.

13. Explain Regularization.

Ans. Regularization is a technique used to prevent overfitting by adding a penalty term to the model's loss function.

14. Name an Ensemble Learning Technique.

Ans. Bagging (Bootstrap Aggregating) is an ensemble learning technique that combines the predictions of multiple base learners to improve overall performance.

15. What is Gradient Descent?

Ans. Gradient Descent is an iterative optimization algorithm used to minimize the loss function of a model by adjusting the model's parameters in the direction of steepest descent.

16. Define Confusion Matrix.

Ans. A confusion matrix is a table that visualizes the performance of a classification algorithm, showing the true positive, true negative, false positive, and false negative counts.

17. Explain One-Hot Encoding.

Ans. One-hot encoding is a technique used to represent categorical variables as binary

vectors, where only one bit is 'hot' (1) indicating the category.

18. What is a Decision Tree?

Ans. A decision tree is a tree-like model used for both classification and regression. It makes decisions based on the values of input features.

19. Define K-Nearest Neighbors (KNN).

Ans. K-Nearest Neighbors is a simple, instance-based learning algorithm where predictions are made by averaging the target values of 'K' nearest neighbors in the training set.

20. What is a Neural Network?

Ans. A neural network is a computational model inspired by the structure of the human brain. It consists of layers of interconnected nodes (neurons) and is used for various machine learning tasks.

Related Posts:

- 1. What is Machine Learning ?
- 2. Types of Machine Learning ?
- 3. Applications of Machine Learning
- 4. Data Preprocessing
- 5. Data Cleaning
- 6. Handling Missing Data
- 7. Feature Scaling
- 8. Labeled data in Machine learning
- 9. Difference between Supervised vs Unsupervised vs Reinforcement learning

- 10. Machine learning algorithms for Big data
- 11. Difference between Supervised vs Unsupervised vs Reinforcement learning
- 12. What is training data in Machine learning
- 13. What is Ordinary Least Squares (OLS) estimation
- 14. Scalar in Machine Learning
- 15. Scalars in Loss Functions | Machine Learning
- 16. Linear Algebra for Machine Learning Practitioners
- 17. Supervised Learning
- 18. Define machine learning and explain its importance in real-world applications.
- 19. Differences Between Machine Learning and Artificial Intelligence
- 20. Machine Learning works on which type of data ?
- 21. What is target variable and independent variable in machine learning
- 22. Machine Learning Scope and Limitations
- 23. What is Regression in Machine learning
- 24. Statistics and linear algebra for machine learning
- 25. Finding Machine Learning Datasets
- 26. What is hypothesis function and testing
- 27. Explain computer vision with an appropriate example
- 28. Explain Reinformcement learning with an appropriate exaple
- 29. Reinforcement Learning Framework
- 30. Data augmentation
- 31. Normalizing Data Sets in Machine Learning
- 32. Machine learning models
- 33. Unsupervised machine learning
- 34. Neural Network in Machine Learning
- 35. Recurrent neural network
- 36. Support Vector Machines

- 37. Long short-term memory (LSTM) networks
- 38. Convolutional neural network
- 39. How to implement Convolutional neural network in Python
- 40. What is MNIST ?
- 41. What does it mean to train a model on a dataset ?
- 42. Can a textual dataset be used with an openCV?
- 43. Name some popular machine learning libraries.
- 44. Introduction to Machine Learning
- 45. Some real time examples of machine learning