

Advanced Concepts in AI and ML

Artificial Intelligence And Machine Learning · Practice Test · 18 Questions

1. Which mathematical technique is primarily used in neural networks to calculate the gradient of the loss function with respect to the weights?

- A) Backpropagation
- B) Linear Regression
- C) K-Nearest Neighbors
- D) Random Forest

2. In the context of machine learning, what is the 'Vanishing Gradient Problem'?

- A) When the model deletes training data
- B) When gradients become too small to effectively update early layers
- C) When the learning rate is set to zero
- D) When a model overfits the test set

3. What is the primary purpose of the 'Softmax' activation function in an output layer of a multi-class classifier?

- A) To normalize inputs to a range of 0 to 1
- B) To convert raw scores into a probability distribution
- C) To prevent overfitting in deep layers
- D) To increase the speed of backpropagation

4. Which architectural component allows Transformers to weigh the importance of different words in a sequence regardless of their distance?

- A) Recurrent connections
- B) Attention mechanism
- C) Convolutional filters
- D) Pooling layers

5. In reinforcement learning, what does the 'Exploration vs. Exploitation' trade-off refer to?

- A) Training time vs. Inference time
- B) Memory usage vs. Processing power
- C) Trying new actions vs. choosing known high-reward actions
- D) Data labeling vs. Data cleaning

6. Which type of machine learning model is specifically designed to handle sequential data by maintaining an internal state?

- A) Recurrent Neural Networks (RNN)
- B) Support Vector Machines
- C) Decision Trees
- D) Clustering algorithms

7. What is the fundamental difference between supervised and unsupervised learning?

- A) Unsupervised learning uses labeled data
- B) Supervised learning requires ground truth labels
- C) Supervised learning uses reinforcement
- D) Unsupervised learning is only for images

8. In Python's Scikit-learn library, what is the primary function of a 'Pipeline' object?

- A) To visualize data distributions
- B) To automate the sequence of data transformations and model estimation
- C) To debug code syntax
- D) To compress large datasets

9. What does the term 'Hyperparameter' refer to in a machine learning model?

- A) The weights learned during training
- B) The parameters set by the developer before training begins
- C) The output values of the model
- D) The input noise in the dataset

10. Which algorithm is commonly used for dimensionality reduction by transforming data into principal components?

- A) Principal Component Analysis (PCA)
- B) K-Means
- C) Logistic Regression
- D) Linear Discriminant Analysis

11. What is the goal of 'Regularization' techniques like L1 and L2 in machine learning?

- A) To increase model complexity
- B) To prevent overfitting by penalizing large weights
- C) To speed up model training
- D) To eliminate the need for training data

12. What is a 'Generative Adversarial Network' (GAN) composed of?

- A) A generator and a discriminator
- B) An encoder and a decoder
- C) A transformer and an attention head
- D) Two supervised models

13. Which metric is calculated as the harmonic mean of precision and recall?

- A) F1-Score
- B) Accuracy
- C) Mean Squared Error
- D) R-squared

14. What is the purpose of a 'Validation Set' during the model training process?

- A) To train the model parameters
- B) To adjust hyperparameters and monitor for overfitting
- C) To calculate the final test accuracy
- D) To store the raw training data

15. In deep learning, what does 'ReLU' stand for?

- A) Rectified Linear Unit
- B) Randomized Efficient Logic Unit
- C) Relative Expected Loss Update
- D) Recurrent Linear Utility

16. Which of the following is an example of an unsupervised learning task?

- A) Image classification
- B) Clustering
- C) Stock price prediction
- D) Credit scoring

17. What occurs during 'Overfitting' in a machine learning model?

- A) The model fails to learn the training data
- B) The model performs well on training data but poorly on unseen data
- C) The model training takes too long
- D) The model requires too much memory

18. Who is often referred to as the 'father of AI' who proposed a test to determine if a machine can exhibit human-like intelligence?

- A) Alan Turing
- B) Ada Lovelace
- C) John McCarthy
- D) Claude Shannon