CPMAI AI Project Exam Prep

Free Practice Test — 30 Real Exam-Style Questions

with full answer key & explanations

Unlock the full bank of 862 questions
+ unlimited timed mock exams + mistake book

Practice on the web: https://certs.theorypractice.app/cpmai
$2.99 / week · $6.99 / month · cancel anytime

What you unlock: all 862 questions • unlimited timed mock exams • mistake book • instant explanations

Study offline on the free app — search your exam on the App Store or Google Play

Practice Questions

Try all 30 first, then check the answer key at the back.

1. A project manager is working with a data science team on a customer churn prediction model. Which type of machine learning approach is most appropriate for this business problem?

A. Unsupervised learning with clustering algorithms
B. Reinforcement learning with reward functions
C. Semi-supervised learning with generative models
D. Supervised learning with classification algorithms

2. Which ensemble method combines multiple weak learners trained sequentially, with each new model attempting to correct errors made by previous models?

A. Stacking
B. Random Forest
C. Boosting
D. Bagging

3. An AI project manager is evaluating different approaches for a credit scoring application. What is the primary advantage of using a random forest over a single decision tree?

A. Lower computational resource requirements
B. Reduced risk of overfitting to training data
C. Faster training and prediction times
D. Perfect interpretability of model decisions

4. In a deep learning project, what is the primary function of an activation function in a neural network?

A. To introduce non-linearity into the network
B. To normalize input data
C. To prevent network parameters from changing
D. To reduce the computational complexity

5. A project team is developing a recommendation system that suggests products to users based on the purchasing patterns of similar customers. Which machine learning approach best describes this scenario?

A. Reinforcement learning
B. Anomaly detection
C. Supervised regression
D. Collaborative filtering

6. Which deep learning architecture is specifically designed for processing sequential data such as time series or natural language?

A. Multilayer Perceptrons (MLPs)
B. Radial Basis Function Networks (RBFNs)
C. Recurrent Neural Networks (RNNs)
D. Convolutional Neural Networks (CNNs)

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

A. Supervised learning requires human intervention during model training, while unsupervised learning is fully automated
B. Supervised learning requires labeled training data, while unsupervised learning works with unlabeled data
C. Supervised learning is computationally less expensive than unsupervised learning
D. Supervised learning is used for clustering, while unsupervised learning is used for classification

8. An AI project team is working with a dataset containing thousands of features. Which technique should they consider to reduce the dimensionality of the data while preserving its important characteristics?

A. Principal Component Analysis (PCA)
B. Gradient Descent Optimization
C. K-means Clustering
D. Cross-validation

9. In a reinforcement learning system, what is the purpose of the reward function?

A. To calculate the computational complexity of the algorithm
B. To determine the optimal number of training iterations
C. To classify input data into predefined categories
D. To provide feedback that guides the learning agent toward desired behaviors

10. When preparing text data for machine learning, what is the purpose of tokenization?

A. Removing outliers from numeric data
B. Converting categorical variables into numeric representations
C. Breaking text into smaller units like words or phrases that can be processed by algorithms
D. Encrypting sensitive information in the text

11. What is the primary advantage of using a probabilistic approach over a deterministic approach in machine learning?

A. It eliminates the need for data preprocessing
B. It provides a measure of uncertainty in predictions
C. It always produces more accurate results
D. It requires less computational resources

12. Which feature engineering technique is most appropriate when dealing with categorical variables that have a large number of possible values?

A. Target encoding
B. One-hot encoding
C. Min-max scaling
D. Principal component analysis

13. A company is implementing a machine learning model to predict equipment failures in a manufacturing plant. Which type of machine learning problem is this?

A. Multi-class classification
B. Regression
C. Clustering
D. Binary classification

14. What is the primary purpose of vectorization in natural language processing?

A. Identifying the most important sentences in a document
B. Translating text between different languages
C. Converting text data into numerical representations that algorithms can process
D. Reducing the number of unique words in a corpus

15. Which of the following is a key characteristic of Large Language Models (LLMs) like GPT?

A. They generate outputs based on explicit programming rules
B. They use transformer architectures with self-attention mechanisms
C. They require minimal computational resources for training
D. They can only be trained on specialized domain-specific data

16. In the context of machine learning, what is the primary difference between prediction and inference?

A. Prediction focuses on forecasting specific outcomes, while inference aims to understand underlying patterns and relationships
B. Prediction requires supervised learning, while inference uses unsupervised learning
C. Prediction works with numeric data, while inference only processes text data
D. Prediction is performed during model training, while inference happens after deployment

17. What is a key advantage of generative AI compared to traditional discriminative models?

A. Lower computational requirements for training
B. Perfect accuracy in classification tasks
C. Independence from large training datasets
D. The ability to create new content that resembles the training data

18. A project team is developing a system to categorize customer support tickets automatically. Which type of neural network architecture would be most appropriate for this text classification task?

A. Hopfield Networks
B. Self-Organizing Maps
C. Transformer-based models
D. Convolutional Neural Networks

19. In a heuristic search algorithm, what is the primary function of the heuristic?

A. To validate the correctness of the final solution
B. To estimate how close a given state is to the goal state
C. To guarantee the optimal solution is always found
D. To eliminate the need for a search algorithm

20. What distinguishes k-means from hierarchical clustering algorithms?

A. K-means requires the number of clusters to be specified in advance
B. K-means can only work with categorical data
C. K-means always produces better quality clusters
D. K-means is computationally less expensive in all cases

21. A data scientist has developed a machine learning model with 99% accuracy on the training data but only 72% accuracy on the test data. What is the most likely issue?

A. The test data is corrupted
B. The model is underfitting
C. The accuracy metric is inappropriate for this problem
D. The model is overfitting to the training data

22. When deploying a machine learning model to production, which validation technique is most appropriate for estimating how the model will perform on unseen data?

A. Manual review of model outputs
B. Comparing against a random baseline
C. K-fold cross-validation
D. Using the same data for training and validation

23. A company has deployed an AI recommendation system that initially showed good results but has seen declining performance over the past three months. What is the most likely cause?

A. Increased system load affecting performance
B. Concept drift in user preferences
C. Hardware degradation
D. Software bugs introduced during updates

24. When aligning AI model results with business KPIs, which approach is most effective?

A. Define business metrics before model development and design the model evaluation to reflect these metrics
B. Focus on technical metrics like accuracy and precision, as business value will naturally follow
C. Develop the model first, then find business KPIs that match the model's strengths
D. Rely on qualitative feedback rather than quantitative metrics to assess business impact

25. What is the primary benefit of implementing a blue/green deployment strategy for AI models?

A. Reduced computational resources required for deployment
B. Elimination of the need for model testing
C. Increased model accuracy through parallel processing
D. Ability to quickly roll back to the previous model version if issues are detected

26. A company is deciding between on-premise and cloud deployment for their new machine learning model. Which factor is MOST critical in this decision?

A. The company's existing cloud service provider
B. The personal preferences of the data science team
C. Data privacy requirements and regulatory constraints
D. The programming language used to develop the model

27. What is the primary purpose of implementing a canary deployment for an AI model?

A. To automatically optimize model hyperparameters in production
B. To test the model with a small percentage of users before full deployment
C. To increase the computational efficiency of the deployment process
D. To eliminate the need for model validation

28. In the context of MLOps, what is the primary purpose of a feature store?

A. To provide a centralized repository of features that can be reused across different models and applications
B. To store model artifacts and versions
C. To manage cloud computing resources
D. To automate the deployment pipeline

29. A company notices that their deployed fraud detection model has a high false positive rate. What is the most appropriate immediate action?

A. Immediately replace the model with a different algorithm
B. Add more features to the model without retraining
C. Ignore the issue as false positives are preferable to false negatives in fraud detection
D. Adjust the classification threshold to better balance precision and recall

30. What is the most effective strategy for monitoring an AI model in production?

A. Rely on periodic manual reviews of model outputs
B. Monitor only for system crashes and errors
C. Implement monitoring for input data distributions, model predictions, and business metrics
D. Focus solely on technical metrics like response time and throughput

Answer Key & Explanations

1. D — Supervised learning with classification algorithms
Customer churn prediction is a classification problem where we predict whether a customer will leave (churn) or stay. This requires supervised learning, as we have historical data with labeled examples of customers who have churned or not churned in the past.

2. C — Boosting
Boosting is an ensemble technique that builds models sequentially, with each new model focusing on correcting the errors made by previous models. AdaBoost and Gradient Boosting are examples of boosting algorithms.

3. B — Reduced risk of overfitting to training data
Random forests reduce overfitting by averaging predictions from multiple trees trained on different subsets of data and features. This ensemble approach provides more robust predictions than a single decision tree, which is prone to overfitting to training data.

4. A — To introduce non-linearity into the network
Activation functions introduce non-linearity into neural networks, allowing them to learn complex patterns. Without activation functions, neural networks would be limited to learning linear relationships regardless of depth.

5. D — Collaborative filtering
Collaborative filtering is a technique used in recommendation systems that identifies patterns in user behavior and preferences by finding similarities between users (user-based) or items (item-based) to make recommendations.

6. C — Recurrent Neural Networks (RNNs)
Recurrent Neural Networks (RNNs) are specifically designed to handle sequential data by maintaining an internal state (memory) that captures information about previous inputs in the sequence, making them well-suited for time series analysis and natural language processing.

7. B — Supervised learning requires labeled training data, while unsupervised learning works with unlabeled data
The fundamental difference between supervised and unsupervised learning is that supervised learning requires labeled training data (with input-output pairs), while unsupervised learning works with unlabeled data and aims to discover patterns or structures within the data without explicit guidance.

8. A — Principal Component Analysis (PCA)
Principal Component Analysis (PCA) is a dimensionality reduction technique that transforms the original features into a new set of uncorrelated features (principal components) that capture the maximum variance in the data, allowing for effective dimensionality reduction while preserving important information.

9. D — To provide feedback that guides the learning agent toward desired behaviors
The reward function in reinforcement learning provides feedback to the agent about the desirability of its actions in different states, guiding the learning process by signaling which actions lead to favorable outcomes and which should be avoided.

10. C — Breaking text into smaller units like words or phrases that can be processed by algorithms
Tokenization is the process of breaking down text into smaller units (tokens) such as words, phrases, or characters. This is a fundamental preprocessing step for text data, as it converts raw text into discrete elements that can be processed by machine learning algorithms.

11. B — It provides a measure of uncertainty in predictions
Probabilistic approaches in machine learning provide a measure of uncertainty or confidence in predictions, which is crucial for risk assessment and decision-making, especially in applications where understanding prediction confidence is important.

12. A — Target encoding
One-hot encoding creates binary features for each category, but with many categories, this can lead to excessive dimensionality. Target encoding replaces categories with their mean target value, effectively handling high-cardinality categorical variables without creating too many new features.

13. D — Binary classification
Predicting equipment failures is a binary classification problem where the model predicts whether a failure will occur (positive class) or not (negative class) based on historical data of equipment conditions and past failures.

14. C — Converting text data into numerical representations that algorithms can process
Vectorization in NLP converts text data into numerical representations (vectors) that machine learning algorithms can process. These vector representations capture semantic relationships between words or documents, enabling algorithms to work with text data.

15. B — They use transformer architectures with self-attention mechanisms
Large Language Models like GPT use transformer architectures with self-attention mechanisms that allow them to consider the context of words in relation to all other words in a sequence, leading to better understanding of language context and meaning.

16. A — Prediction focuses on forecasting specific outcomes, while inference aims to understand underlying patterns and relationships
Prediction focuses on using a trained model to make specific forecasts or classifications on new data, while inference involves drawing broader conclusions or insights about patterns, relationships, and underlying mechanisms in the data.

17. D — The ability to create new content that resembles the training data
Generative AI models learn the underlying data distribution and can create new content that resembles the training data, while discriminative models focus only on making predictions or classifications based on input features.

18. C — Transformer-based models
Transformers are particularly well-suited for text classification tasks because they can capture long-range dependencies and contextual information in text through their self-attention mechanism, leading to

state-of-the-art performance in NLP tasks like ticket categorization.

19. B — To estimate how close a given state is to the goal state
A heuristic provides an estimated cost or distance to the goal state, helping the search algorithm prioritize more promising paths and improve efficiency by avoiding exhaustive exploration of all possible paths.

20. A — K-means requires the number of clusters to be specified in advance
K-means requires the number of clusters (k) to be specified in advance, while hierarchical clustering builds a tree of clusters without requiring a pre-specified number, allowing users to choose the appropriate number of clusters after examining the cluster hierarchy.

21. D — The model is overfitting to the training data
This significant drop in performance between training and test data is a classic sign of overfitting, where the model has memorized the training data rather than learning generalizable patterns.

22. C — K-fold cross-validation
K-fold cross-validation provides a robust estimate of model performance by testing on multiple held-out data subsets, which better simulates real-world performance than a single train-test split.

23. B — Concept drift in user preferences
Model drift occurs when the statistical properties of the target variable change over time, causing the model's predictions to become less accurate. This is common in recommendation systems as user preferences evolve.

24. A — Define business metrics before model development and design the model evaluation to reflect these metrics
Defining clear business metrics before model development ensures the AI solution is designed to address specific business goals rather than just technical performance metrics.

25. D — Ability to quickly roll back to the previous model version if issues are detected
Blue/green deployment allows for immediate rollback to the previous version if issues are detected with the new model, minimizing downtime and user impact.

26. C — Data privacy requirements and regulatory constraints
Data privacy and regulatory requirements often dictate where data can be stored and processed, making this a critical factor in deployment decisions, especially for sensitive information.

27. B — To test the model with a small percentage of users before full deployment
Canary deployments limit risk by exposing the new model to a small percentage of users first, allowing for monitoring of performance and issues before full deployment.

28. A — To provide a centralized repository of features that can be reused across different models and applications
Feature stores serve as centralized repositories for preprocessed features, ensuring consistency across training and inference while reducing redundant computation.

29. D — Adjust the classification threshold to better balance precision and recall
Adjusting the classification threshold can help balance the trade-off between false positives and false negatives without requiring model retraining.

30. C — Implement monitoring for input data distributions, model predictions, and business metrics
A comprehensive monitoring strategy includes both input data and output predictions, along with business metrics, to detect issues like data drift, model drift, and misalignment with business goals.

Ready to pass?

Unlock the full CPMAI AI Project Exam Prep bank, every explanation, and unlimited timed mock exams.

Scan to start practising
https://certs.theorypractice.app/cpmai

Also on iOS & Android — search your exam name on the App Store or Google Play