CertNexus AIP-210 CAIP Practice Questions

Shared by Crosby on 21-10-2026

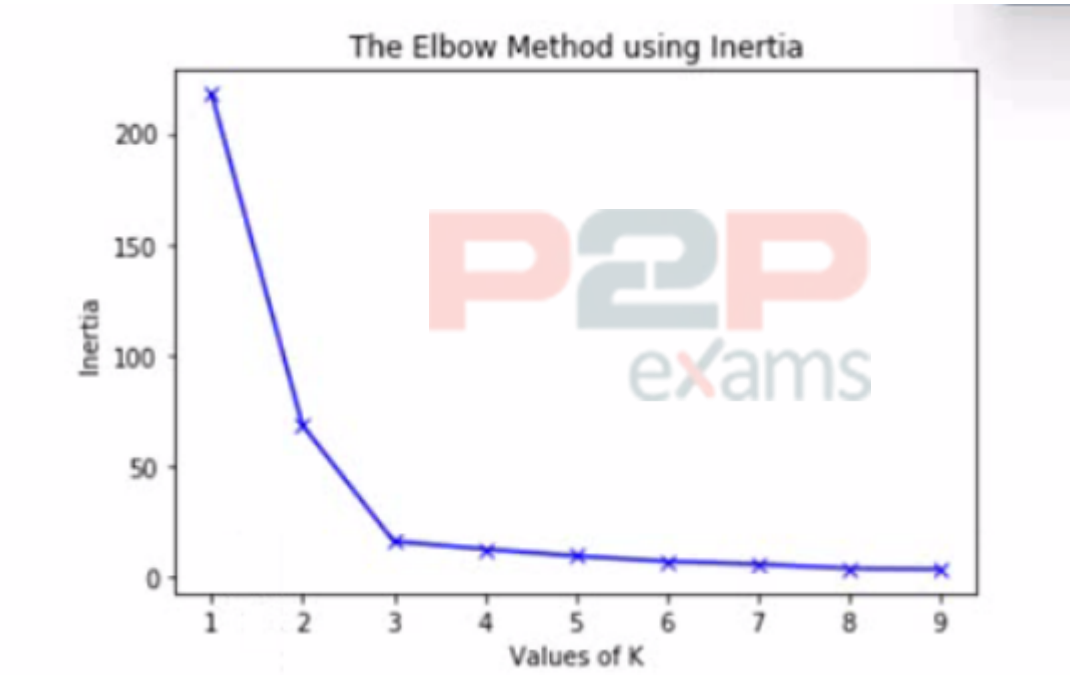
For More Free Questions and Preparation Resources

Check the Links on Last Page

Question 1

Question Type: MultipleChoice

Refer to the exhibit.



The graph is an elbow plot showing the inertia or within-cluster sum of squares on the y-axis and number of clusters (also called K) on the x-axis, denoting the change in inertia as the clusters change using k-means algorithm.

What would be an optimal value of K to ensure a good number of clusters?

Options:

A- 2
B- 3
C- 5
D- 9

Answer:

B

Explanation:

The optimal value of K is the one that minimizes the inertia or within-cluster sum of squares, while avoiding too many clusters that may overfit the data. The elbow plot shows a sharp decrease in inertia from K = 1 to K = 2, and then a more gradual decrease from K = 2 to K = 3.

After K = 3, the inertia does not change much as K increases. Therefore, the elbow point is at K = 3, which is the optimal value of K for this data. Reference: How to Run K-Means Clustering in Python, K-means clustering - Wikipedia

Question 2

Question Type: MultipleChoice

What is the primary benefit of the Federated Learning approach to machine learning?

Options:

A- It does not require a labeled dataset to solve supervised learning problems.
B- It protects the privacy of the user's data while providing well-trained models.
C- It requires less computation to train the same model using a traditional approach.
D- It uses large, centralized data stores to train complex machine learning models.

Answer:

B

Explanation:

Federated learning is a distributed approach to machine learning that allows multiple parties to collaboratively train a model without sharing their data with each other or a central server. This protects the privacy of the user's data while still enabling well-trained models that can benefit from diverse and large-scale datasets. Reference: [Federated Learning - Wikipedia], [Federated Learning for Mobile Keyboard Prediction - Google AI Blog]

Question 3

Question Type: MultipleChoice

Which three security measures could be applied in different ML workflow stages to defend them against malicious activities? (Select three.)

Options:

A- Disable logging for model access.
B- Launch ML Instances In a virtual private cloud (VPC).
C- Monitor model degradation.
D- Use data encryption.
E- Use max privilege to control access to ML artifacts.
F- Use Secrets Manager to protect credentials.

Answer:

B, D, F

Explanation:

Security measures can be applied in different ML workflow stages to defend them against malicious activities, such as data theft, model tampering, or adversarial attacks. Some of the security measures are:

Launch ML Instances In a virtual private cloud (VPC): A VPC is a logically isolated section of a cloud provider's network that allows users to launch and control their own resources. By launching ML instances in a VPC, users can enhance the security and privacy of their data and models, as well as restrict the access and traffic to and from the instances.

Use data encryption: Data encryption is the process of transforming data into an unreadable format using a secret key or algorithm. Data encryption can protect the confidentiality, integrity, and availability of data at rest (stored in databases or files) or in transit (transferred over networks). Data encryption can prevent unauthorized access, modification, or leakage of sensitive data.

Use Secrets Manager to protect credentials: Secrets Manager is a service that helps users securely store, manage, and retrieve secrets, such as passwords, API keys, tokens, or certificates. Secrets Manager can help users protect their credentials from unauthorized access or exposure, as well as rotate them automatically to comply with security policies.

Question 4

Question Type: MultipleChoice

Which option best describes a typical use case of video tracking?

Options:

A- Augmented dreaming
B- Medical diagnosis
C- Traffic monitoring
D- Video composition

Answer:

C

Explanation:

Video tracking is a technique that involves detecting and following moving objects in a video sequence. Video tracking can be used for various applications, such as surveillance, security, sports analysis, and human-computer interaction. One typical use case of video tracking is traffic monitoring, where video tracking can help measure traffic flow, detect congestion, identify violations, and optimize traffic signals.

Question 5

Question Type: MultipleChoice

Which option best scenarios is an example of entanglement in ML pipelines?

Options:

A- Add a new method for drift detection in the model evaluation step.
B- Add a new pipeline for retraining the model in the model training step.
C- Change in normalization function in the feature engineering step.
D- Change the way output is visualized in the monitoring step.

Answer:

C

Explanation:

Entanglement in ML pipelines occurs when a change in one step affects other steps that depend on it. Changing the normalization function in the feature engineering step would affect the model training and evaluation steps, as they rely on the features generated by the feature engineering step. Therefore, this scenario is an example of entanglement in ML pipelines. The other scenarios are not examples of entanglement, as they do not affect other steps in the

pipeline.

Question 6

Question Type: MultipleChoice

Which of the following unsupervised learning models can a bank use for fraud detection?

Options:

A- Anomaly detection
B- DB5CAN
C- Hierarchical clustering
D- k-means

Answer:

A

Explanation:

Anomaly detection is an unsupervised learning technique that identifies outliers or abnormal patterns in data, which can be useful for fraud detection. Anomaly detection algorithms can learn the normal behavior of transactions and flag the ones that deviate significantly from the norm, indicating possible fraud.

Question 7

Question Type: MultipleChoice

Your dependent variable Y is a count, ranging from 0 to infinity. Because Y is approximately log-normally distributed, you decide to log-transform the data prior to performing a linear regression.

What should you do before log-transforming Y?

Options:

A- Add 1 to all of the Y values.
B- Divide all the Y values by the standard deviation of Y.
C- Explore the data for outliers.
D- Subtract the mean of Y from all the Y values.

Answer:

A

Explanation:

Before log-transforming Y, we should add 1 to all of the Y values. This is because log transformation is undefined for zero or negative values, and some of the Y values may be zero. Adding 1 to all of the Y values can avoid this problem and ensure that the log transformation is valid and meaningful. Adding 1 to all of the Y values is also known as a log-plus-one transformation.

Question 8

Question Type: MultipleChoice

Which two encoders can be used to transform categorical data into numerical features? (Select two.)

Options:

A- Count Encoder
B- Log Encoder
C- Mean Encoder
D- Median Encoder
E- One-Hot Encoder

Answer:

A, E

Question 9

Question Type: MultipleChoice

Personal data should not be disclosed, made available, or otherwise used for purposes other than specified with Which option best exceptions? (Select two.)

Options:

A- If it is for a good cause.
B- If it was collected accidentally.
C- If it was requested by the authority of law.
D- If it was with consent of the person it is collected from.
E- If the data is only collected once.

P2P exams

Answer:

C, D

Explanation:

Personal data is any information that relates to an identified or identifiable individual, such as name, address, email, phone number, or biometric data. Personal data should not be disclosed, made available, or otherwise used for purposes other than specified, except with:

The consent of the person it is collected from: Consent is a clear and voluntary indication of agreement by the person to the processing of their personal data for a specific purpose. Consent can be given by a statement or a clear affirmative action, such as ticking a box or clicking a button.

The authority of law: The authority of law is a legal basis or obligation that requires or permits the processing of personal data for a legitimate purpose. For example, the authority of law could be a court order, a subpoena, a warrant, or a statute.

P2P exams

To Get Premium Files for AIP-210 Visit

https://www.p2pexams.com/products/aip-210

For More Free Questions Visit

https://www.p2pexams.com/certnexus/pdf/aip-210

