Features of Amazon MLS-C01 Desktop Practice Exam Software

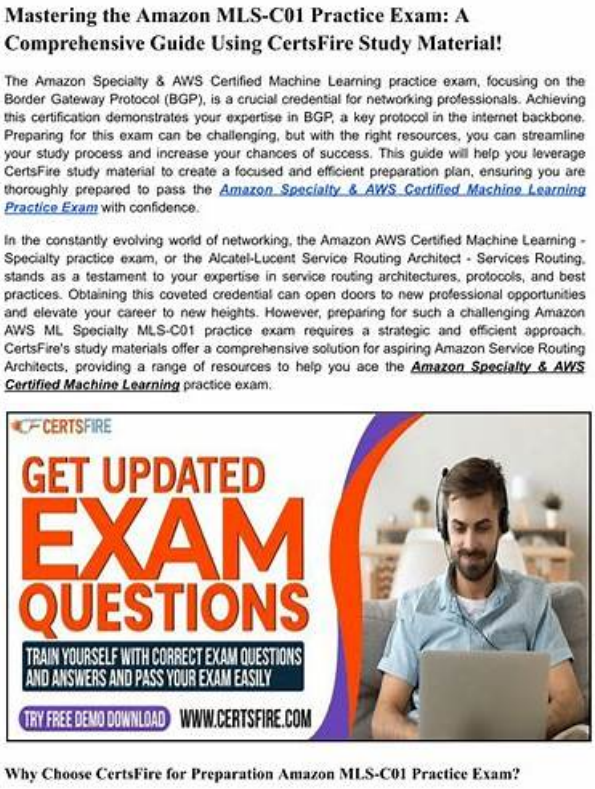

BTW, DOWNLOAD part of Prep4cram MLS-C01 dumps from Cloud Storage: https://drive.google.com/open?id=1BKqR0A9ESs0vovdiAyZdAquN4ePBtyDW

It helps you to pass the Amazon MLS-C01 test with excellent results. Amazon MLS-C01 imitates the actual MLS-C01 exam environment. You can take the MLS-C01 practice exam many times to evaluate and enhance your Amazon MLS-C01 Exam Preparation level. Desktop MLS-C01 practice test software is compatible with windows and the web-based software will work on these operating systems: Android, IOS, Windows, and Linux.

The AWS Certified Machine Learning - Specialty certification exam is ideal for professionals who are already working in the field of machine learning and want to demonstrate their expertise in the AWS environment. It is also suitable for individuals who are interested in pursuing a career in machine learning and want to gain a competitive advantage in the job market. The AWS Certified Machine Learning - Specialty certification is recognized by top companies in the tech industry, making it a valuable asset for professionals looking to advance their careers.

>> Detailed MLS-C01 Answers <<

MLS-C01 Dump Collection, MLS-C01 Vce Files

Some of our customers are white-collar workers with no time to waste, and need a Amazon certification urgently to get their promotions, meanwhile the other customers might aim at improving their skills. So we try to meet different requirements by setting different versions of our MLS-C01 question dumps. The first one is online MLS-C01 engine version. As an online tool, it is convenient and easy to study, supports all Web Browsers and system including Windows, Mac, Android, iOS and so on. You can

practice online anytime and check your test history and performance review, which will do help to your study. The second is MLS-C01 Desktop Test Engine. As an installable MLS-C01 software application, it simulated the real MLS-C01 exam environment, and builds 200-125 exam confidence. The third one is Practice PDF version. PDF Version is easy to read and print. So you can study anywhere, anytime.

Career Path

In case you want to specialize in more specific AWS services, you can opt for other Amazon specialty certifications like the AWS Certified Advanced Networking Specialty, the AWS Certified Alexa Skill Builder Specialty, or the AWS Certified Database Specialty if to name a few.

Amazon AWS Certified Machine Learning - Specialty Sample Questions (Q126-Q131):

NEW QUESTION # 126
An agricultural company is interested in using machine learning to detect specific types of weeds in a 100-acre grassland field. Currently, the company uses tractor-mounted cameras to capture multiple images of the field as 10 * 10 grids. The company also has a large training dataset that consists of annotated images of popular weed classes like broadleaf and non-broadleaf docks.
The company wants to build a weed detection model that will detect specific types of weeds and the location of each type within the field. Once the model is ready, it will be hosted on Amazon SageMaker endpoints. The model will perform real-time inferencing using the images captured by the cameras.
Which approach should a Machine Learning Specialist take to obtain accurate predictions?

- A. Prepare the images in Apache Parquet format and upload them to Amazon S3. Use Amazon SageMaker to train, test, and validate the model using an image classification algorithm to categorize images into various weed classes.
- B. Prepare the images in RecordIO format and upload them to Amazon S3. Use Amazon SageMaker to train, test, and validate the model using an object-detection single-shot multibox detector (SSD) algorithm.
- C. Prepare the images in Apache Parquet format and upload them to Amazon S3. Use Amazon SageMaker to train, test, and validate the model using an object-detection single-shot multibox detector (SSD) algorithm.
- D. Prepare the images in RecordIO format and upload them to Amazon S3. Use Amazon SageMaker to train, test, and validate the model using an image classification algorithm to categorize images into various weed classes.

Answer: B

NEW QUESTION # 127
A company wants to classify user behavior as either fraudulent or normal. Based on internal research, a Machine Learning Specialist would like to build a binary classifier based on two features: age of account and transaction month. The class distribution for these features is illustrated in the figure provided.

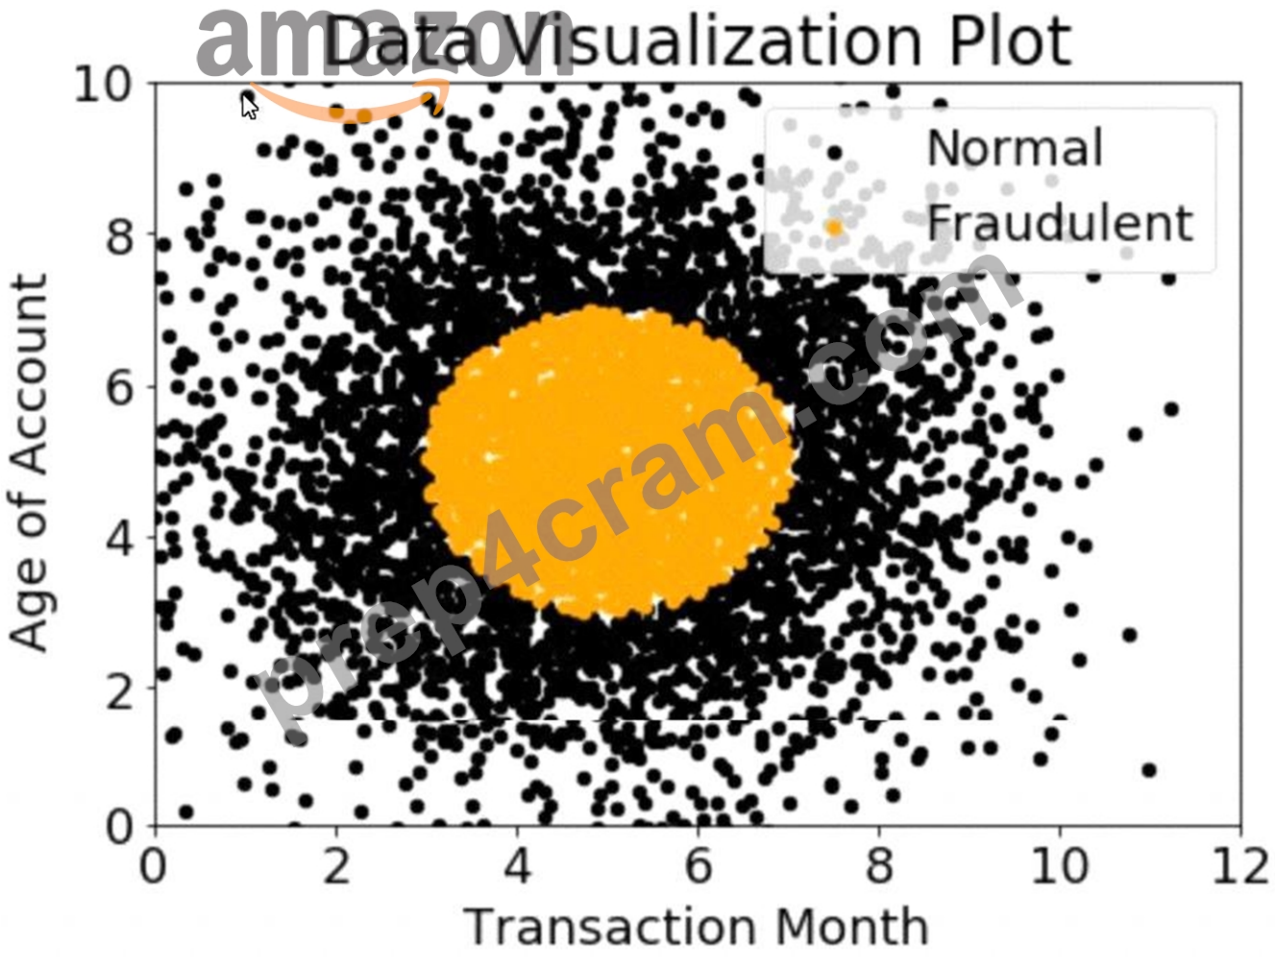

Based on this information, which model would have the HIGHEST recall with respect to the fraudulent class?

- A. Single Perceptron with sigmoidal activation function
- B. Decision tree
- C. Naive Bayesian classifier
- D. Linear support vector machine (SVM)

Answer: B

Explanation:
Explanation
Based on the figure provided, a decision tree would have the highest recall with respect to the fraudulent class. Recall is a model evaluation metric that measures the proportion of actual positive instances that are correctly classified by the model. Recall is calculated as follows:
Recall = True Positives / (True Positives + False Negatives)
A decision tree is a type of machine learning model that can perform classification tasks by splitting the data into smaller and purer subsets based on a series of rules or conditions. A decision tree can handle both linear and non-linear data, and can capture complex patterns and interactions among the features. A decision tree can also be easily visualized and interpreted1 In this case, the data is not linearly separable, and has a clear pattern of seasonality. The fraudulent class forms a large circle in the center of the plot, while the normal class is scattered around the edges. A decision tree can use the transaction month and the age of account as the splitting criteria, and create a circular boundary that separates the fraudulent class from the normal class. A decision tree can achieve a high recall for the fraudulent class, as it can correctly identify most of the black dots as positive instances, and minimize the number of false negatives. A decision tree can also adjust the depth and complexity of the tree to balance the trade-off between recall and precision23 The other options are not valid or suitable for achieving a high recall for the fraudulent class. A linear support vector machine (SVM) is a type of machine learning model that can perform classification tasks by finding a linear hyperplane that maximizes the margin between the classes. A linear SVM can handle linearly separable data, but not non-linear data. A linear SVM cannot capture the circular pattern of the fraudulent class, and may misclassify many of the black dots as negative instances, resulting in a low recall4 A naive Bayesian classifier is a type of machine learning model that can perform classification tasks by applying the Bayes' theorem and assuming conditional independence among the features. A naive Bayesian classifier can handle both linear and non-linear data, and can incorporate prior knowledge and probabilities into the model. However, a naive Bayesian classifier may not perform well when the features are correlated or dependent, as in this case. A naive Bayesian classifier may not capture the circular

pattern of the fraudulent class, and may misclassify many of the black dots as negative instances, resulting in a low recall5 A single perceptron with sigmoidal activation function is a type of machine learning model that can perform classification tasks by applying a weighted linear combination of the features and a non-linear activation function. A single perceptron with sigmoidal activation function can handle linearly separable data, but not non-linear data. A single perceptron with sigmoidal activation function cannot capture the circular pattern of the fraudulent class, and may misclassify many of the black dots as negative instances, resulting in a low recall.

NEW QUESTION # 128
A retail company uses a machine learning (ML) model for daily sales forecasting. The company's brand manager reports that the model has provided inaccurate results for the past 3 weeks.
At the end of each day, an AWS Glue job consolidates the input data that is used for the forecasting with the actual daily sales data and the predictions of the model. The AWS Glue job stores the data in Amazon S3. The company's ML team is using an Amazon SageMaker Studio notebook to gain an understanding about the source of the model's inaccuracies.
What should the ML team do on the SageMaker Studio notebook to visualize the model's degradation MOST accurately?

- A. Create a histogram of the daily sales over the last 3 weeks. In addition, create a histogram of the daily sales from before that period.
- B. Create a line chart with the weekly mean absolute error (MAE) of the model.
- C. Create a histogram of the model errors over the last 3 weeks. In addition, create a histogram of the model errors from before that period.
- D. Create a scatter plot of daily sales versus model error for the last 3 weeks. In addition, create a scatter plot of daily sales versus model error from before that period.

Answer: C

Explanation:
The best way to visualize the model's degradation is to create a histogram of the model errors over the last 3 weeks and compare it with a histogram of the model errors from before that period. A histogram is a graphical representation of the distribution of numerical data. It shows how often each value or range of values occurs in the data. A model error is the difference between the actual value and the predicted value. A high model error indicates a poor fit of the model to the data. By comparing the histograms of the model errors, the ML team can see if there is a significant change in the shape, spread, or center of the distribution. This can indicate if the model is underfitting, overfitting, or drifting from the data. A line chart or a scatter plot would not be as effective as a histogram for this purpose, because they do not show the distribution of the errors. A line chart would only show the trend of the errors over time, which may not capture the variability or outliers. A scatter plot would only show the relationship between the errors and another variable, such as daily sales, which may not be relevant or informative for the model's performance. References:
Histogram - Wikipedia
Model error - Wikipedia
SageMaker Model Monitor - visualizing monitoring results

NEW QUESTION # 129
A Machine Learning Specialist deployed a model that provides product recommendations on a company's website. Initially, the model was performing very well and resulted in customers buying more products on average. However, within the past few months, the Specialist has noticed that the effect of product recommendations has diminished and customers are starting to return to their original habits of spending less. The Specialist is unsure of what happened, as the model has not changed from its initial deployment over a year ago.
Which method should the Specialist try to improve model performance?

- A. The model needs to be completely re-engineered because it is unable to handle product inventory changes.
- B. The model's hyperparameters should be periodically updated to prevent drift.
- C. The model should be periodically retrained using the original training data plus new data as product inventory changes.
- D. The model should be periodically retrained from scratch using the original data while adding a regularization term to handle product inventory changes

Answer: C

NEW QUESTION # 130
Acybersecurity company is collecting on-premises server logs, mobile app logs, and loT sensor data. The company backs up the

ingested data in an Amazon S3 bucket and sends the ingested data to Amazon OpenSearch Service for further analysis. Currently, the company has a custom ingestion pipeline that is running on Amazon EC2 instances. The company needs to implement a new serverless ingestion pipeline that can automatically scale to handle sudden changes in the data flow.
Which solution will meet these requirements MOST cost-effectively?

- A. Create one Amazon Kinesis data stream. Create two Amazon Data Firehose delivery streams to send data to the S3 bucket and OpenSearch Service. Connect the delivery streams to the data stream.
  Configure the data sources to send data to the data stream.
- B. Create two Amazon Data Firehose delivery streams to send data to the S3 bucket and OpenSearch Service. Configure the data sources to send data to the delivery streams.
- C. Create one Amazon Kinesis data stream. Create one Amazon Data Firehose delivery stream to send data to OpenSearch Service. Configure the delivery stream to back up the data to the S3 bucket.Connect the delivery stream to the data stream. Configure the data sources to send data to the data stream.
- D. Create one Amazon Data Firehose delivery stream to send data to OpenSearch Service. Configure the delivery stream to back up the raw data to the S3 bucket. Configure the data sources to send data to the delivery stream.

Answer: A

Explanation:
To build a scalable, serverless, and cost-effective data ingestion pipeline, this solution uses a Kinesis data stream to handle fluctuations in data flow, buffering and distributing incoming data in real time. By connecting two Amazon Kinesis Data Firehose delivery streams to the Kinesis data stream, the company can simultaneously route data to Amazon S3 for backup and Amazon OpenSearch Service for analysis.
This approach meets all requirements by providing automatic scaling, reducing operational overhead, and ensuring data storage and analysis without duplicating efforts or needing additional infrastructure.

NEW QUESTION # 131
.....

myportal.utt.edu.tt, myportal.utt.edu.tt, www.stes.tyc.edu.tw, motionentrance.edu.np, Disposable vapes

BONUS!!! Download part of Prep4cram MLS-C01 dumps for free: https://drive.google.com/open?id=1BKqR0A9ESs0vovdiAyZdAquN4ePBtyDW