Latest Amazon MLS-C01: AWS Certified Machine Learning - Specialty Reliable Test Topics - Authoritative DumpsFree MLS-C01 New Exam Bootcamp

BTW, DOWNLOAD part of DumpsFree MLS-C01 dumps from Cloud Storage: https://drive.google.com/open?id=19v7rajmZxqEOLfQWJHLydNmVlS25EuSS

Competition appear everywhere in modern society. There are many way to improve ourselves and learning methods of MLS-C01 exams come in different forms. Economy rejuvenation and social development carry out the blossom of technology; some MLS-C01 practice materials are announced which have a good quality. Certification qualification MLS-C01 Exam Materials are a big industry and many companies are set up for furnish a variety of services for it. And our MLS-C01 study guide has three different versions: PDF, Soft and APP versions to let you study in varied and comfortable ways.

To earn the AWS Certified Machine Learning - Specialty certification, candidates must pass the MLS-C01 exam, which consists of 65 multiple-choice and multiple-response questions. MLS-C01 exam covers a broad range of topics, including data engineering, feature engineering, model selection and training, and deployment and monitoring of ML models. MLS-C01 Exam also tests the candidate's knowledge of AWS's various ML services, such as Amazon SageMaker, Amazon Rekognition, and Amazon Comprehend.

>> MLS-C01 Reliable Test Topics <<

Helpful Features of MLS-C01 PDF Questions

Many candidates find the Amazon MLS-C01 exam preparation difficult. They often buy expensive study courses to start their Amazon MLS-C01 certification exam preparation. However, spending a huge amount on such resources is difficult for many AWS Certified Machine Learning - Specialty exam applicants. The latest Amazon MLS-C01 Exam Dumps are the right option for you to prepare for the Amazon MLS-C01 certification test at home.

The AWS Certified Machine Learning - Specialty certification exam is ideal for professionals who are looking to advance their careers in the field of machine learning and artificial intelligence. It is a great way to showcase your skills and expertise to potential employers and clients, and to demonstrate your commitment to staying up-to-date with the latest developments in this rapidly evolving field. Additionally, AWS certification exams are recognized globally, which means that earning this certification can help you

land new job opportunities in different countries and regions.

Amazon AWS Certified Machine Learning - Specialty Sample Questions (Q165-Q170):

NEW QUESTION # 165
A chemical company has developed several machine learning (ML) solutions to identify chemical process abnormalities. The time series values of independent variables and the labels are available for the past 2 years and are sufficient to accurately model the problem.
The regular operation label is marked as 0. The abnormal operation label is marked as 1 . Process abnormalities have a significant negative effect on the companys profits. The company must avoid these abnormalities.
Which metrics will indicate an ML solution that will provide the GREATEST probability of detecting an abnormality?

- A. Precision = 0.7
  Recall = 0.9
- B. Precision = 0.91
  Recall = 0.6
- C. Precision = 0.61
  Recall = 0.98
- D. Precision = 0.98
  Recall = 0.8

Answer: C

Explanation:
The metrics that will indicate an ML solution that will provide the greatest probability of detecting an abnormality are precision and recall. Precision is the ratio of true positives (TP) to the total number of predicted positives (TP + FP), where FP is false positives. Recall is the ratio of true positives (TP) to the total number of actual positives (TP + FN), where FN is false negatives. A high precision means that the ML solution has a low rate of false alarms, while a high recall means that the ML solution has a high rate of true detections. For the chemical company, the goal is to avoid process abnormalities, which are marked as 1 in the labels. Therefore, the company needs an ML solution that has a high recall for the positive class, meaning that it can detect most of the abnormalities and minimize the false negatives. Among the four options, option B has the highest recall for the positive class, which is 0.98. This means that the ML solution can detect 98% of the abnormalities and miss only 2%. Option B also has a reasonable precision for the positive class, which is 0.61. This means that the ML solution has a false alarm rate of 39%, which may be acceptable for the company, depending on the cost and benefit analysis. The other options have lower recall for the positive class, which means that they have higher false negative rates, which can be more detrimental for the company than false positive rates.
References:
1: AWS Certified Machine Learning - Specialty Exam Guide
2: AWS Training - Machine Learning on AWS
3: AWS Whitepaper - An Overview of Machine Learning on AWS
4: Precision and recall

NEW QUESTION # 166
A data scientist is developing a pipeline to ingest streaming web traffic data. The data scientist needs to implement a process to identify unusual web traffic patterns as part of the pipeline. The patterns will be used downstream for alerting and incident response. The data scientist has access to unlabeled historic data to use, if needed.
The solution needs to do the following:
* Calculate an anomaly score for each web traffic entry.
* Adapt unusual event identification to changing web patterns over time.
Which approach should the data scientist implement to meet these requirements?

- A. Collect the streaming data using Amazon Kinesis Data Firehose. Map the delivery stream as an input source for Amazon Kinesis Data Analytics. Write a SQL query to run in real time against the streaming data with the Amazon Random Cut Forest (RCF) SQL extension to calculate anomaly scores for each record using a sliding window.
- B. Use historic web traffic data to train an anomaly detection model using the Amazon SageMaker built-in XGBoost model. Use an Amazon Kinesis Data Stream to process the incoming web traffic data. Attach a preprocessing AWS Lambda function to perform data enrichment by calling the XGBoost model to calculate the anomaly score for each record.
- C. Collect the streaming data using Amazon Kinesis Data Firehose. Map the delivery stream as an input source for Amazon Kinesis Data Analytics. Write a SQL query to run in real time against the streaming data with the k-Nearest Neighbors

(kNN) SQL extension to calculate anomaly scores for each record using a tumbling window.

- D. Use historic web traffic data to train an anomaly detection model using the Amazon SageMaker Random Cut Forest (RCF) built-in model. Use an Amazon Kinesis Data Stream to process the incoming web traffic data. Attach a preprocessing AWS Lambda function to perform data enrichment by calling the RCF model to calculate the anomaly score for each record.

Answer: A

Explanation:
Amazon Kinesis Data Analytics is a service that allows users to analyze streaming data in real time using SQL queries. Amazon Random Cut Forest (RCF) is a SQL extension that enables anomaly detection on streaming data. RCF is an unsupervised machine learning algorithm that assigns an anomaly score to each data point based on how different it is from the rest of the data. A sliding window is a type of window that moves along with the data stream, so that the anomaly detection model can adapt to changing patterns over time. A tumbling window is a type of window that has a fixed size and does not overlap with other windows, so that the anomaly detection model is based on a fixed period of time. Therefore, option D is the best approach to meet the requirements of the question, as it uses RCF to calculate anomaly scores for each web traffic entry and uses a sliding window to adapt to changing web patterns over time.
Option A is incorrect because Amazon SageMaker Random Cut Forest (RCF) is a built-in model that can be used to train and deploy anomaly detection models on batch or streaming data, but it requires more steps and resources than using the RCF SQL extension in Amazon Kinesis Data Analytics. Option B is incorrect because Amazon SageMaker XGBoost is a built-in model that can be used for supervised learning tasks such as classification and regression, but not for unsupervised learning tasks such as anomaly detection. Option C is incorrect because k-Nearest Neighbors (kNN) is a SQL extension that can be used for classification and regression tasks on streaming data, but not for anomaly detection. Moreover, using a tumbling window would not allow the anomaly detection model to adapt to changing web patterns over time.
Using CloudWatch anomaly detection
Anomaly Detection With CloudWatch
Performing Real-time Anomaly Detection using AWS
What Is AWS Anomaly Detection? (And Is There A Better Option?)

NEW QUESTION # 167
A telecommunications company is developing a mobile app for its customers. The company is using an Amazon SageMaker hosted endpoint for machine learning model inferences.
Developers want to introduce a new version of the model for a limited number of users who subscribed to a preview feature of the app. After the new version of the model is tested as a preview, developers will evaluate its accuracy. If a new version of the model has better accuracy, developers need to be able to gradually release the new version for all users over a fixed period of time.
How can the company implement the testing model with the LEAST amount of operational overhead?

- A. Update the ProductionVariant data type with the new version of the model by using the CreateEndpointConfig operation with the InitialVariantWeight parameter set to 0. Specify the TargetVariant parameter for InvokeEndpoint calls for users who subscribed to the preview feature.
  When the new version of the model is ready for release, gradually increase InitialVariantWeight until all users have the updated version.
- B. Update the DesiredWeightsAndCapacity data type with the new version of the model by using the UpdateEndpointWeightsAndCapacities operation with the DesiredWeight parameter set to 0. Specify the TargetVariant parameter for InvokeEndpoint calls for users who subscribed to the preview feature.
  When the new version of the model is ready for release, gradually increase DesiredWeight until all users have the updated version.
- C. Configure two SageMaker hosted endpoints that serve the different versions of the model. Create an Application Load Balancer (ALB) to route traffic to both endpoints based on the TargetVariant query string parameter. Reconfigure the app to send the TargetVariant query string parameter for users who subscribed to the preview feature. When the new version of the model is ready for release, change the ALB's routing algorithm to weighted until all users have the updated version.
- D. Configure two SageMaker hosted endpoints that serve the different versions of the model. Create an Amazon Route 53 record that is configured with a simple routing policy and that points to the current version of the model. Configure the mobile app to use the endpoint URL for users who subscribed to the preview feature and to use the Route 53 record for other users. When the new version of the model is ready for release, add a new model version endpoint to Route 53, and switch the policy to weighted until all users have the updated version.

Answer: B

Explanation:
Explanation

The best solution for implementing the testing model with the least amount of operational overhead is to use the following steps: Update the DesiredWeightsAndCapacity data type with the new version of the model by using the UpdateEndpointWeightsAndCapacities operation with the DesiredWeight parameter set to 0. This operation allows the developers to update the variant weights and capacities of an existing SageMaker endpoint without deleting and recreating the endpoint. Setting the DesiredWeight parameter to 0 means that the new version of the model will not receive any traffic initially1 Specify the TargetVariant parameter for InvokeEndpoint calls for users who subscribed to the preview feature. This parameter allows the developers to override the variant weights and direct a request to a specific variant. This way, the developers can test the new version of the model for a limited number of users who opted in for the preview feature2 When the new version of the model is ready for release, gradually increase DesiredWeight until all users have the updated version. This operation allows the developers to perform a gradual rollout of the new version of the model and monitor its performance and accuracy. The developers can adjust the variant weights and capacities as needed until the new version of the model serves all the traffic1 The other options are incorrect because they either require more operational overhead or do not support the desired use cases. For example:
Option A uses the CreateEndpointConfig operation with the InitialVariantWeight parameter set to 0.
This operation creates a new endpoint configuration, which requires deleting and recreating the endpoint to apply the changes. This adds extra overhead and downtime for the endpoint. It also does not support the gradual rollout of the new version of the model3 Option B uses two SageMaker hosted endpoints that serve the different versions of the model and an Application Load Balancer (ALB) to route traffic to both endpoints based on the TargetVariant query string parameter. This option requires creating and managing additional resources and services, such as the second endpoint and the ALB. It also requires changing the app code to send the query string parameter for the preview feature4 Option D uses the access key and secret key of the IAM user with appropriate KMS and ECR permissions. This is not a secure way to pass credentials to the Processing job. It also requires the ML specialist to manage the IAM user and the keys.
References:
1: UpdateEndpointWeightsAndCapacities - Amazon SageMaker
2: InvokeEndpoint - Amazon SageMaker
3: CreateEndpointConfig - Amazon SageMaker
4: Application Load Balancer - Elastic Load Balancing

NEW QUESTION # 168
A company operates large cranes at a busy port. The company plans to use machine learning (ML) for predictive maintenance of the cranes to avoid unexpected breakdowns and to improve productivity.
The company already uses sensor data from each crane to monitor the health of the cranes in real time. The sensor data includes rotation speed, tension, energy consumption, vibration, pressure, and ...perature for each crane. The company contracts AWS ML experts to implement an ML solution.
Which potential findings would indicate that an ML-based solution is suitable for this scenario? (Select TWO.)

- A. The historical sensor data from the cranes are available with high granularity for the last 3 years.
- B. The historical sensor data contains most common types of crane failures that the company wants to predict.
- C. The historical sensor data shows that simple rule-based thresholds can predict crane failures.
- D. The historical sensor data does not include a significant number of data points and attributes for certain time periods.
- E. The historical sensor data contains failure data for only one type of crane model that is in operation and lacks failure data of most other types of crane that are in operation.

Answer: A,B

Explanation:
The best indicators that an ML-based solution is suitable for this scenario are D and E, because they imply that the historical sensor data is sufficient and relevant for building a predictive maintenance model. This model can use machine learning techniques such as regression, classification, or anomaly detection to learn from the past data and forecast future failures or issues12. Having high granularity and diversity of data can improve the accuracy and generalization of the model, as well as enable the detection of complex patterns and relationships that are not captured by simple rule-based thresholds3.
The other options are not good indicators that an ML-based solution is suitable, because they suggest that the historical sensor data is incomplete, inconsistent, or inadequate for building a predictive maintenance model. These options would require additional data collection, preprocessing, or augmentation to overcome the data quality issues and ensure that the model can handle different scenarios and types of cranes4 .
1: Machine Learning Techniques for Predictive Maintenance
2: A Guide to Predictive Maintenance & Machine Learning
3: Machine Learning for Predictive Maintenance: Reinventing Asset Upkeep
4: Predictive Maintenance with Machine Learning: A Complete Guide
[Machine Learning for Predictive Maintenance - AWS Online Tech Talks]

NEW QUESTION # 169
A Machine Learning Specialist is designing a system for improving sales for a company. The objective is to use the large amount of information the company has on users' behavior and product preferences to predict which products users would like based on the users' similarity to other users.
What should the Specialist do to meet this objective?

- A. Build a model-based filtering recommendation engine with Apache Spark ML on Amazon EMR.
- B. Build a collaborative filtering recommendation engine with Apache Spark ML on Amazon EMR.
- C. Build a combinative filtering recommendation engine with Apache Spark ML on Amazon EMR.
- D. Build a content-based filtering recommendation engine with Apache Spark ML on Amazon EMR.

Answer: B

Explanation:
Many developers want to implement the famous Amazon model that was used to power the "People who bought this also bought these items" feature on Amazon.com. This model is based on a method called Collaborative Filtering. It takes items such as movies, books, and products that were rated highly by a set of users and recommending them to other users who also gave them high ratings. This method works well in domains where explicit ratings or implicit user actions can be gathered and analyzed.

NEW QUESTION # 170
......