Hot AWS-Certified-Machine-Learning-Specialty Reliable Braindumps Book | Pass-Sure Amazon Reliable AWS-Certified-Machine-Learning-Specialty Test Dumps: AWS Certified Machine Learning - Specialty

P.S. Free & New AWS-Certified-Machine-Learning-Specialty dumps are available on Google Drive shared by Exam4Tests: https://drive.google.com/open?id=168IAIfu_PssnTX3Lr-nThRinToh6AHWS

Exam4Tests alerts you that the syllabus of the AWS Certified Machine Learning - Specialty (AWS-Certified-Machine-Learning-Specialty) certification exam changes from time to time. Therefore, keep checking the fresh updates released by the Amazon. It will save you from the unnecessary mental hassle of wasting your valuable money and time. Exam4Tests announces another remarkable feature to its users by giving them the AWS Certified Machine Learning - Specialty (AWS-Certified-Machine-Learning-Specialty) dumps updates until 1 year after purchasing the AWS Certified Machine Learning - Specialty (AWS-Certified-Machine-Learning-Specialty) certification exam pdf questions.

Understanding functional and technical aspects of AWS Certified Machine Learning - Specialty Modeling

The following will be discussed in **AMAZON MLS-C01 exam dumps**:

- Train machine learning models
- Evaluate machine learning models
- Perform hyperparameter optimization
- Frame business problems as machine learning problems
- Select the appropriate model(s) for a given machine learning problem

>> AWS-Certified-Machine-Learning-Specialty Reliable Braindumps Book <<

Amazon AWS-Certified-Machine-Learning-Specialty Reliable Braindumps

Book: AWS Certified Machine Learning - Specialty - Exam4Tests Precise Reliable Test Dumps for your free downloading

Our AWS-Certified-Machine-Learning-Specialty training materials are compiled by professional experts. All the necessary points have been mentioned in our AWS-Certified-Machine-Learning-Specialty practice engine particularly. About some tough questions or important points, they left notes under them. Besides, our experts will concern about changes happened in AWS-Certified-Machine-Learning-Specialty study prep all the time. Provided you have a strong determination, as well as the help of our AWS-Certified-Machine-Learning-Specialty learning guide, you can have success absolutely.

Amazon AWS Certified Machine Learning - Specialty Sample Questions (Q212-Q217):

NEW QUESTION # 212
A machine learning (ML) specialist uploads 5 TB of data to an Amazon SageMaker Studio environment. The ML specialist performs initial data cleansing. Before the ML specialist begins to train a model, the ML specialist needs to create and view an analysis report that details potential bias in the uploaded data.
Which combination of actions will meet these requirements with the LEAST operational overhead? (Choose two.)

- A. Turn on the bias detection option in SageMaker Ground Truth to automatically analyze data features.
- B. Configure SageMaker Data Wrangler to generate a bias report.
- C. Use SageMaker Experiments to perform a data check
- D. Use SageMaker Model Monitor to generate a bias drift report.
- E. Use SageMaker Clarify to automatically detect data bias

Answer: B,E

Explanation:
The combination of actions that will meet the requirements with the least operational overhead is to use SageMaker Clarify to automatically detect data bias and to configure SageMaker Data Wrangler to generate a bias report. SageMaker Clarify is a feature of Amazon SageMaker that provides machine learning (ML) developers with tools to gain greater insights into their ML training data and models. SageMaker Clarify can detect potential bias during data preparation, after model training, and in your deployed model. For instance, you can check for bias related to age in your dataset or in your trained model and receive a detailed report that quantifies different types of potential bias1. SageMaker Data Wrangler is another feature of Amazon SageMaker that enables you to prepare data for machine learning (ML) quickly and easily. You can use SageMaker Data Wrangler to identify potential bias during data preparation without having to write your own code. You specify input features, such as gender or age, and SageMaker Data Wrangler runs an analysis job to detect potential bias in those features. SageMaker Data Wrangler then provides a visual report with a description of the metrics and measurements of potential bias so that you can identify steps to remediate the bias2. The other actions either require more customization (such as using SageMaker Model Monitor or SageMaker Experiments) or do not meet the requirement of detecting data bias (such as using SageMaker Ground Truth). References:
1: Bias Detection and Model Explainability - Amazon Web Services
2: Amazon SageMaker Data Wrangler - Amazon Web Services

NEW QUESTION # 213
A manufacturer is operating a large number of factories with a complex supply chain relationship where unexpected downtime of a machine can cause production to stop at several factories. A data scientist wants to analyze sensor data from the factories to identify equipment in need of preemptive maintenance and then dispatch a service team to prevent unplanned downtime. The sensor readings from a single machine can include up to 200 data points including temperatures, voltages, vibrations, RPMs, and pressure readings. To collect this sensor data, the manufacturer deployed Wi-Fi and LANs across the factories. Even though many factory locations do not have reliable or high-speed internet connectivity, the manufacturer would like to maintain near-real-time inference capabilities.
Which deployment architecture for the model will address these business requirements?

- A. Deploy the model in Amazon SageMaker. Run sensor data through this model to predict which machines need maintenance.
- B. Deploy the model on AWS IoT Greengrass in each factory. Run sensor data through this model to infer which machines need maintenance.
- C. Deploy the model to an Amazon SageMaker batch transformation job. Generate inferences in a daily batch report to identify machines that need maintenance.
- D. Deploy the model in Amazon SageMaker and use an IoT rule to write data to an Amazon DynamoDB table. Consume a DynamoDB stream from the table with an AWS Lambda function to invoke the endpoint.

Answer: B

Explanation:
AWS IoT Greengrass is a service that extends AWS to edge devices, such as sensors and machines, so they can act locally on the data they generate, while still using the cloud for management, analytics, and durable storage. AWS IoT Greengrass enables local device messaging, secure data transfer, and local computing using AWS Lambda functions and machine learning models. AWS IoT Greengrass can run machine learning inference locally on devices using models that are created and trained in the cloud. This allows devices to respond quickly to local events, even when they are offline or have intermittent connectivity. Therefore, option B is the best deployment architecture for the model to address the business requirements of the manufacturer.
Option A is incorrect because deploying the model in Amazon SageMaker would require sending the sensor data to the cloud for inference, which would not work well for factory locations that do not have reliable or high-speed internet connectivity. Moreover, this option would not provide near-real-time inference capabilities, as there would be latency and bandwidth issues involved in transferring the data to and from the cloud. Option C is incorrect because deploying the model to an Amazon SageMaker batch transformation job would not provide near-real-time inference capabilities, as batch transformation is an asynchronous process that operates on large datasets. Batch transformation is not suitable for streaming data that requires low-latency responses. Option D is incorrect because deploying the model in Amazon SageMaker and using an IoT rule to write data to an Amazon DynamoDB table would also require sending the sensor data to the cloud for inference, which would have the same drawbacks as option A. Moreover, this option would introduce additional complexity and cost by involving multiple services, such as IoT Core, DynamoDB, and Lambda.
References:
AWS Greengrass Machine Learning Inference - Amazon Web Services
Machine learning components - AWS IoT Greengrass
What is AWS Greengrass? | AWS IoT Core | Onica
GitHub - aws-samples/aws-greengrass-ml-deployment-sample
AWS IoT Greengrass Architecture and Its Benefits | Quick Guide - XenonStack

NEW QUESTION # 214
A data scientist must build a custom recommendation model in Amazon SageMaker for an online retail company. Due to the nature of the company's products, customers buy only 4-5 products every 5-10 years. So, the company relies on a steady stream of new customers. When a new customer signs up, the company collects data on the customer's preferences. Below is a sample of the data available to the data scientist.

timestamp	user_id	product_id	preference_1	...	preference_10
2020-03-04	90	25	0	...	0.374
2020-03-04	90	61	0	...	0.374
2020-02-21	203	56	1	...	0.098

How should the data scientist split the dataset into a training and test set for this use case?

- A. Identify the 10% of users with the least interaction data. Split off all interaction data from these users for the test set.
- B. Shuffle all interaction data. Split off the last 10% of the interaction data for the test set.
- C. Randomly select 10% of the users. Split off all interaction data from these users for the test set.
- D. Identify the most recent 10% of interactions for each user. Split off these interactions for the test set.

Answer: C

Explanation:
The best way to split the dataset into a training and test set for this use case is to randomly select 10% of the users and split off all interaction data from these users for the test set. This is because the company relies on a steady stream of new customers, so the test set should reflect the behavior of new customers who have not been seen by the model before. The other options are not suitable because they either mix old and new customers in the test set (A and B), or they bias the test set towards users with less interaction data .
References:
* Amazon SageMaker Developer Guide: Train and Test Datasets
* Amazon Personalize Developer Guide: Preparing and Importing Data

NEW QUESTION # 215
A Machine Learning Specialist is using an Amazon SageMaker notebook instance in a private subnet of a corporate VPC. The ML Specialist has important data stored on the Amazon SageMaker notebook instance's Amazon EBS volume, and needs to take a snapshot of that EBS volume. However the ML Specialist cannot find the Amazon SageMaker notebook instance's EBS volume or

Amazon EC2 instance within the VPC.
Why is the ML Specialist not seeing the instance visible in the VPC?

- A. Amazon SageMaker notebook instances are based on AWS ECS instances running within AWS service accounts.
- B. Amazon SageMaker notebook instances are based on the Amazon ECS service within customer accounts.
- C. Amazon SageMaker notebook instances are based on EC2 instances running within AWS service accounts.
- D. Amazon SageMaker notebook instances are based on the EC2 instances within the customer account, but they run outside of VPCs.

Answer: C

NEW QUESTION # 216
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 # 217
......

If you get the AWS-Certified-Machine-Learning-Specialty certification, your working abilities will be proved and you will find an ideal job. We provide you with AWS-Certified-Machine-Learning-Specialty exam materials of high quality which can help you pass the exam easily. We provide you with AWS-Certified-Machine-Learning-Specialty exam materials of high quality which can help you pass the exam easily. It also saves your much time and energy that you only need little time to learn and prepare for exam. We also provide timely and free update for you to get more AWS-Certified-Machine-Learning-Specialty Questions torrent and follow the latest trend. The AWS-Certified-Machine-Learning-Specialty exam torrent is compiled by the experienced professionals and of great value.