

Test AWS-Certified-Machine-Learning-Specialty Dumps Pdf | AWS-Certified-Machine-Learning-Specialty Certificate Exam



What's more, part of that UpdateDumps AWS-Certified-Machine-Learning-Specialty dumps now are free:
<https://drive.google.com/open?id=1LD6hMVxvOeRpkHdZDWJGZPJdJpSDtKb>

Our AWS-Certified-Machine-Learning-Specialty study question has high quality. So there is all effective and central practice for you to prepare for your test. With our professional ability, we can accord to the necessary testing points to edit AWS-Certified-Machine-Learning-Specialty exam questions. It points to the exam heart to solve your difficulty. So high quality materials can help you to pass your exam effectively, make you feel easy, to achieve your goal. With the AWS-Certified-Machine-Learning-Specialty Test Guide use feedback, it has 98%-100% pass rate. That's the truth from our customers. And it is easy for you to pass the AWS-Certified-Machine-Learning-Specialty exam after 20 hours' to 30 hours' practice.

To prepare for the AWS-Certified-Machine-Learning-Specialty exam, candidates are advised to take online courses, read books, and practice with sample questions. Amazon provides a range of training materials and resources to help candidates prepare for the exam. Candidates can also find a range of third-party resources, including study guides, practice exams, and online courses, to help them prepare for the exam.

Exam Details

The Amazon MLS-C01 exam consists of 65 questions that are presented in the multiple-response and multiple-choice formats. These items have to be answered within the allocated time of 170 minutes. It is required that the candidates get 750 points on a scale of 100-1000. This test is available in many languages, including English, Simplified Chinese, Japanese, and Korean. To schedule it, you have to pay the fee of \$300.

The AWS-Certified-Machine-Learning-Specialty Exam covers various topics, such as data engineering, exploratory data analysis, modeling, machine learning implementation and operations, and ethical and legal considerations. Candidates should be well-versed in these topics and should have hands-on experience using AWS services, such as Amazon SageMaker, Amazon S3, Amazon EC2, and Amazon Comprehend.

>> Test AWS-Certified-Machine-Learning-Specialty Dumps Pdf <<

Pass Guaranteed Quiz Professional Amazon - Test AWS-Certified-Machine-Learning-Specialty Dumps Pdf

Our company is a professional certification exam materials provider, we have occupied in this field for more than ten years, and therefore we have rich experience. AWS-Certified-Machine-Learning-Specialty exam braindumps are high quality, because we have a professional team to collect the first-hand information for the exam, we can ensure that you can get the latest information for the exam. In addition, our company is strict with the quality and answers for AWS-Certified-Machine-Learning-Specialty Exam Materials, and therefore you can use them at ease. Our AWS-Certified-Machine-Learning-Specialty exam braindumps are known as instant access to download, you can obtain the downloading link and password within ten minutes.

Amazon AWS Certified Machine Learning - Specialty Sample Questions (Q234-Q239):

NEW QUESTION # 234

A data scientist needs to identify fraudulent user accounts for a company's ecommerce platform. The company wants the ability to determine if a newly created account is associated with a previously known fraudulent user. The data scientist is using AWS Glue to cleanse the company's application logs during ingestion.

Which strategy will allow the data scientist to identify fraudulent accounts?

- A. Search for duplicate accounts in the AWS Glue Data Catalog.
- B. Execute the built-in FindDuplicates Amazon Athena query.
- C. Create an AWS Glue crawler to infer duplicate accounts in the source data.
- D. Create a FindMatches machine learning transform in AWS Glue.

Answer: D

Explanation:

Explanation

The best strategy to identify fraudulent accounts is to create a FindMatches machine learning transform in AWS Glue. The FindMatches transform enables you to identify duplicate or matching records in your dataset, even when the records do not have a common unique identifier and no fields match exactly. This can help you improve fraud detection by finding accounts that are associated with a previously known fraudulent user. You can teach the FindMatches transform your definition of a "duplicate" or a "match" through examples, and it will use machine learning to identify other potential duplicates or matches in your dataset. You can then use the FindMatches transform in your AWS Glue ETL jobs to cleanse your data.

Option A is incorrect because there is no built-in FindDuplicates Amazon Athena query. Amazon Athena is an interactive query service that makes it easy to analyze data in Amazon S3 using standard SQL. However, Amazon Athena does not provide a predefined query to find duplicate records in a dataset. You would have to write your own SQL query to perform this task, which might not be as effective or accurate as using the FindMatches transform.

Option C is incorrect because creating an AWS Glue crawler to infer duplicate accounts in the source data is not a valid strategy. An AWS Glue crawler is a program that connects to a data store, progresses through a prioritized list of classifiers to determine the schema for your data, and then creates metadata tables in the AWS Glue Data Catalog. A crawler does not perform any data cleansing or record matching tasks.

Option D is incorrect because searching for duplicate accounts in the AWS Glue Data Catalog is not a feasible strategy. The AWS Glue Data Catalog is a central repository to store structural and operational metadata for your data assets. The Data Catalog does not store the actual data, but rather the metadata that describes where the data is located, how it is formatted, and what it contains. Therefore, you cannot search for duplicate records in the Data Catalog.

References:

Record matching with AWS Lake Formation FindMatches - AWS Glue

Amazon Athena - Interactive SQL Queries for Data in Amazon S3

AWS Glue Crawlers - AWS Glue

AWS Glue Data Catalog - AWS Glue

NEW QUESTION # 235

During mini-batch training of a neural network for a classification problem, a Data Scientist notices that training accuracy oscillates. What is the MOST likely cause of this issue?

- A. The batch size is too big
- B. The learning rate is very high
- C. Dataset shuffling is disabled
- D. The class distribution in the dataset is imbalanced

Answer: C

NEW QUESTION # 236

A data scientist is training a large PyTorch model by using Amazon SageMaker. It takes 10 hours on average to train the model on GPU instances. The data scientist suspects that training is not converging and that resource utilization is not optimal.

What should the data scientist do to identify and address training issues with the LEAST development effort?

- A. Use the SageMaker Debugger confusion and feature_importance_overweight built-in rules to detect issues and to launch the StopTrainingJob action if issues are detected.
- B. Use high-resolution custom metrics that are captured in Amazon CloudWatch. Configure an AWS Lambda function to analyze the metrics and to stop the training job early if issues are detected.
- C. Use CPU utilization metrics that are captured in Amazon CloudWatch. Configure a CloudWatch alarm to stop the training job early if low CPU utilization occurs.
- D. Use the SageMaker Debugger vanishing_gradient and LowGPUUtilization built-in rules to detect issues and to launch the StopTrainingJob action if issues are detected.

Answer: D

Explanation:

The solution C is the best option to identify and address training issues with the least development effort. The solution C involves the following steps:

Use the SageMaker Debugger vanishing_gradient and LowGPUUtilization built-in rules to detect issues.

SageMaker Debugger is a feature of Amazon SageMaker that allows data scientists to monitor, analyze, and debug machine learning models during training. SageMaker Debugger provides a set of built-in rules that can automatically detect common issues and anomalies in model training, such as vanishing or exploding gradients, overfitting, underfitting, low GPU utilization, and more¹. The data scientist can use the vanishing_gradient rule to check if the gradients are becoming too small and causing the training to not converge. The data scientist can also use the LowGPUUtilization rule to check if the GPU resources are underutilized and causing the training to be inefficient².

Launch the StopTrainingJob action if issues are detected. SageMaker Debugger can also take actions based on the status of the rules. One of the actions is StopTrainingJob, which can terminate the training job if a rule is in an error state. This can help the data scientist to save time and money by stopping the training early if issues are detected³.

The other options are not suitable because:

Option A: Using CPU utilization metrics that are captured in Amazon CloudWatch and configuring a CloudWatch alarm to stop the training job early if low CPU utilization occurs will not identify and address training issues effectively. CPU utilization is not a good indicator of model training performance, especially for GPU instances. Moreover, CloudWatch alarms can only trigger actions based on simple thresholds, not complex rules or conditions⁴.

Option B: Using high-resolution custom metrics that are captured in Amazon CloudWatch and configuring an AWS Lambda function to analyze the metrics and to stop the training job early if issues are detected will incur more development effort than using SageMaker Debugger. The data scientist will have to write the code for capturing, sending, and analyzing the custom metrics, as well as for invoking the Lambda function and stopping the training job. Moreover, this solution may not be able to detect all the issues that SageMaker Debugger can⁵.

Option D: Using the SageMaker Debugger confusion and feature_importance_overweight built-in rules and launching the StopTrainingJob action if issues are detected will not identify and address training issues effectively. The confusion rule is used to monitor the confusion matrix of a classification model, which is not relevant for a regression model that predicts prices. The feature_importance_overweight rule is used to check if some features have too much weight in the model, which may not be related to the convergence or resource utilization issues².

1: Amazon SageMaker Debugger

2: Built-in Rules for Amazon SageMaker Debugger

3: Actions for Amazon SageMaker Debugger

4: Amazon CloudWatch Alarms

5: Amazon CloudWatch Custom Metrics

NEW QUESTION # 237

A Machine Learning Specialist is creating a new natural language processing application that processes a dataset comprised of 1 million sentences. The aim is to then run Word2Vec to generate embeddings of the sentences and enable different types of predictions - Here is an example from the dataset

"The quck BROWN FOX jumps over the lazy dog "

Which of the following are the operations the Specialist needs to perform to correctly sanitize and prepare the data in a repeatable manner? (Select THREE)

- A. Tokenize the sentence into words.
- B. Remove stop words using an English stopword dictionary.
- C. Perform part-of-speech tagging and keep the action verb and the nouns only
- D. Normalize all words by making the sentence lowercase
- E. Correct the typography on "quck" to "quick."
- F. One-hot encode all words in the sentence

Answer: A,B,D

Explanation:

To prepare the data for Word2Vec, the Specialist needs to perform some preprocessing steps that can help reduce the noise and complexity of the data, as well as improve the quality of the embeddings. Some of the common preprocessing steps for Word2Vec are:

Normalizing all words by making the sentence lowercase: This can help reduce the vocabulary size and treat words with different capitalizations as the same word. For example, "Fox" and "fox" should be considered as the same word, not two different words.

Removing stop words using an English stopword dictionary: Stop words are words that are very common and do not carry much semantic meaning, such as "the", "a", "and", etc. Removing them can help focus on the words that are more relevant and informative for the task.

Tokenizing the sentence into words: Tokenization is the process of splitting a sentence into smaller units, such as words or subwords. This is necessary for Word2Vec, as it operates on the word level and requires a list of words as input.

The other options are not necessary or appropriate for Word2Vec:

Performing part-of-speech tagging and keeping the action verb and the nouns only: Part-of-speech tagging is the process of assigning a grammatical category to each word, such as noun, verb, adjective, etc. This can be useful for some natural language processing tasks, but not for Word2Vec, as it can lose some important information and context by discarding other words.

Correcting the typography on "quck" to "quick": Typo correction can be helpful for some tasks, but not for Word2Vec, as it can introduce errors and inconsistencies in the data. For example, if the typo is intentional or part of a dialect, correcting it can change the meaning or style of the sentence. Moreover, Word2Vec can learn to handle typos and variations in spelling by learning similar embeddings for them.

One-hot encoding all words in the sentence: One-hot encoding is a way of representing words as vectors of 0s and 1s, where only one element is 1 and the rest are 0. The index of the 1 element corresponds to the word's position in the vocabulary. For example, if the vocabulary is ["cat", "dog", "fox"], then "cat" can be encoded as [1, 0, 0], "dog" as [0, 1, 0], and "fox" as [0, 0, 1]. This can be useful for some machine learning models, but not for Word2Vec, as it does not capture the semantic similarity and relationship between words. Word2Vec aims to learn dense and low-dimensional embeddings for words, where similar words have similar vectors.

NEW QUESTION # 238

A Machine Learning Specialist is configuring Amazon SageMaker so multiple Data Scientists can access notebooks, train models, and deploy endpoints. To ensure the best operational performance, the Specialist needs to be able to track how often the Scientists are deploying models, GPU and CPU utilization on the deployed SageMaker endpoints, and all errors that are generated when an endpoint is invoked.

Which services are integrated with Amazon SageMaker to track this information? (Select TWO.)

- A. AWS Config
- **B. Amazon CloudWatch**
- C. AWS Trusted Advisor
- **D. AWS CloudTrail**
- E. AWS Health

Answer: B,D

Explanation:

Explanation

The services that are integrated with Amazon SageMaker to track the information that the Machine Learning Specialist needs are AWS CloudTrail and Amazon CloudWatch. AWS CloudTrail is a service that records the API calls and events for AWS services, including Amazon SageMaker. AWS CloudTrail can track the actions performed by the Data Scientists, such as creating notebooks, training models, and deploying endpoints. AWS CloudTrail can also provide information such as the identity of the user, the time of the action, the parameters used, and the response elements returned. AWS CloudTrail can help the Machine Learning Specialist to monitor the usage and activity of Amazon SageMaker, as well as to audit and troubleshoot any issues. Amazon CloudWatch is a service that collects and analyzes the metrics and logs for AWS services, including Amazon SageMaker. Amazon CloudWatch can track the performance and utilization of the Amazon SageMaker endpoints, such as the CPU and GPU utilization, the inference latency, the number of invocations, etc. Amazon CloudWatch can also track the errors and alarms that are generated when an endpoint is invoked, such as the model errors, the throttling errors, the HTTP errors, etc. Amazon CloudWatch can help the Machine Learning Specialist to optimize the operational performance and reliability of Amazon SageMaker, as well as to set up notifications and actions based on the metrics and logs.

myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt,
myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt,
myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, Disposable vapes

P.S. Free 2026 Amazon AWS-Certified-Machine-Learning-Specialty dumps are available on Google Drive shared by
UpdateDumps: <https://drive.google.com/open?id=1LD6hMVxvOeRpkrHdZDWJGZPJJDjPsDtKb>