Sample Amazon AWS-Certified-Machine-Learning-Specialty Questions Pdf, AWS-Certified-Machine-Learning-Specialty Valuable Feedback

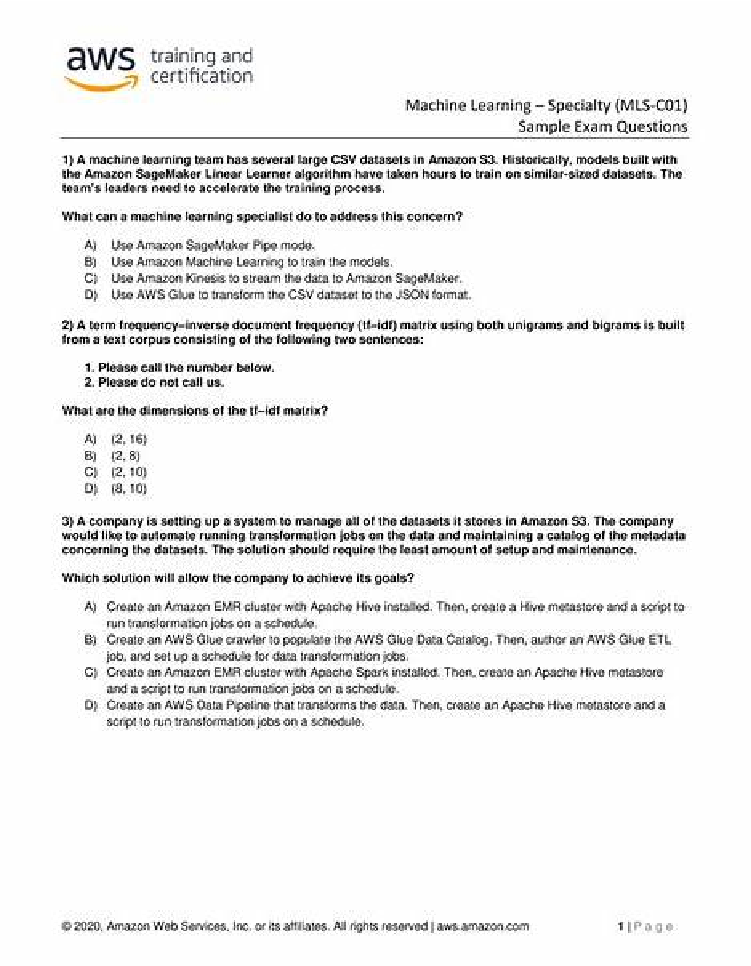
aws training and certification

Machine Learning – Specialty (MLS-C01)
Sample Exam Questions

1) A machine learning team has several large CSV datasets in Amazon S3. Historically, models built with the Amazon SageMaker Linear Learner algorithm have taken hours to train on similar-sized datasets. The team's leaders need to accelerate the training process.

What can a machine learning specialist do to address this concern?

A) Use Amazon SageMaker Pipe mode.
B) Use Amazon Machine Learning to train the models.
C) Use Amazon Kinesis to stream the data to Amazon SageMaker.
D) Use AWS Glue to transform the CSV dataset to the JSON format.

2) A term frequency–inverse document frequency (tf–idf) matrix using both unigrams and bigrams is built from a text corpus consisting of the following two sentences:

1. Please call the number below.
2. Please do not call us.

What are the dimensions of the tf–idf matrix?

A) (2, 16)
B) (2, 8)
C) (2, 10)
D) (8, 10)

3) A company is setting up a system to manage all of the datasets it stores in Amazon S3. The company would like to automate running transformation jobs on the data and maintaining a catalog of the metadata concerning the datasets. The solution should require the least amount of setup and maintenance.

Which solution will allow the company to achieve its goals?

A) Create an Amazon EMR cluster with Apache Hive installed. Then, create a Hive metastore and a script to run transformation jobs on a schedule.
B) Create an AWS Glue crawler to populate the AWS Glue Data Catalog. Then, author an AWS Glue ETL job, and set up a schedule for data transformation jobs.
C) Create an Amazon EMR cluster with Apache Spark installed. Then, create an Apache Hive metastore and a script to run transformation jobs on a schedule.
D) Create an AWS Data Pipeline that transforms the data. Then, create an Apache Hive metastore and a script to run transformation jobs on a schedule.

© 2020, Amazon Web Services, Inc. or its affiliates. All rights reserved | aws.amazon.com 1 | Page

What's more, part of that RealVCE AWS-Certified-Machine-Learning-Specialty dumps now are free: https://drive.google.com/open?id=1p84IpLLRxqYJcwy5uua7dBLAl7MtGk7g

If you choose to register Amazon AWS-Certified-Machine-Learning-Specialty certification exam, you must try to get the AWS-Certified-Machine-Learning-Specialty certification. If you are apprehensive of defeat, you can select RealVCE Amazon AWS-Certified-Machine-Learning-Specialty dumps. No matter what your qualification and your ability are, you can grasp these knowledge easily. RealVCE Amazon AWS-Certified-Machine-Learning-Specialty Test Questions and answers is the latest. We provide you with free update for one year. After using it, you will make a difference.

You can change the difficulty of these questions, which will help you determine what areas appertain to more study before taking your Amazon AWS-Certified-Machine-Learning-Specialty Exam Dumps. Here we listed some of the most important benefits you can get from using our Amazon AWS-Certified-Machine-Learning-Specialty practice questions.

>> Sample Amazon AWS-Certified-Machine-Learning-Specialty Questions Pdf <<

Fantastic Sample AWS-Certified-Machine-Learning-Specialty Questions Pdf – Find Shortcut to Pass AWS-Certified-Machine-Learning-Specialty Exam

If you buy our AWS-Certified-Machine-Learning-Specialty exam questions, we will offer you high quality products and perfect after service just as in the past. We believe our consummate after-sale service system will make our customers feel the most satisfactory.

Our company has designed the perfect after sale service system for these people who buy our AWS-Certified-Machine-Learning-Specialty practice materials. We can always give the most professinal suggestion on our AWS-Certified-Machine-Learning-Specialty learning guide to our customers at the first time for our service are working 24/7 online.

Amazon AWS Certified Machine Learning - Specialty Sample Questions (Q302-Q307):

NEW QUESTION # 302
Which of the following metrics should a Machine Learning Specialist generally use to compare/evaluate machine learning classification models against each other?

- A. Area Under the ROC Curve (AUC)
- B. Misclassification rate
- C. Mean absolute percentage error (MAPE)
- D. Recall

Answer: A

NEW QUESTION # 303
A Machine Learning team runs its own training algorithm on Amazon SageMaker. The training algorithm requires external assets. The team needs to submit both its own algorithm code and algorithm-specific parameters to Amazon SageMaker.
What combination of services should the team use to build a custom algorithm in Amazon SageMaker?
(Choose two.)

- A. Amazon ECS
- B. Amazon S3
- C. AWS Secrets Manager
- D. Amazon ECR
- E. AWS CodeStar

Answer: B,D

Explanation:
Explanation
The Machine Learning team wants to use its own training algorithm on Amazon SageMaker, and submit both its own algorithm code and algorithm-specific parameters. The best combination of services to build a custom algorithm in Amazon SageMaker are Amazon ECR and Amazon S3.
Amazon ECR is a fully managed container registry service that allows you to store, manage, and deploy Docker container images. You can use Amazon ECR to create a Docker image that contains your training algorithm code and any dependencies or libraries that it requires. You can also use Amazon ECR to push, pull, and manage your Docker images securely and reliably.
Amazon S3 is a durable, scalable, and secure object storage service that can store any amount and type of data. You can use Amazon S3 to store your training data, model artifacts, and algorithm-specific parameters. You can also use Amazon S3 to access your data and parameters from your training algorithm code, and to write your model output to a specified location.
Therefore, the Machine Learning team can use the following steps to build a custom algorithm in Amazon SageMaker:
Write the training algorithm code in Python, using the Amazon SageMaker Python SDK or the Amazon SageMaker Containers library to interact with the Amazon SageMaker service. The code should be able to read the input data and parameters from Amazon S3, and write the model output to Amazon S3.
Create a Dockerfile that defines the base image, the dependencies, the environment variables, and the commands to run the training algorithm code. The Dockerfile should also expose the ports that Amazon SageMaker uses to communicate with the container.
Build the Docker image using the Dockerfile, and tag it with a meaningful name and version.
Push the Docker image to Amazon ECR, and note the registry path of the image.
Upload the training data, model artifacts, and algorithm-specific parameters to Amazon S3, and note the S3 URIs of the objects.
Create an Amazon SageMaker training job, using the Amazon SageMaker Python SDK or the AWS CLI. Specify the registry path of the Docker image, the S3 URIs of the input and output data, the algorithm-specific parameters, and other configuration options, such as the instance type, the number of instances, the IAM role, and the hyperparameters.
Monitor the status and logs of the training job, and retrieve the model output from Amazon S3.
References:
Use Your Own Training Algorithms
Amazon ECR - Amazon Web Services
Amazon S3 - Amazon Web Services

NEW QUESTION # 304
A gaming company has launched an online game where people can start playing for free but they need to pay if they choose to use certain features The company needs to build an automated system to predict whether or not a new user will become a paid user within 1 year The company has gathered a labeled dataset from 1 million users The training dataset consists of 1.000 positive samples (from users who ended up paying within 1 year) and
999.000 negative samples (from users who did not use any paid features) Each data sample consists of 200 features including user age, device, location, and play patterns Using this dataset for training, the Data Science team trained a random forest model that converged with over
99% accuracy on the training set However, the prediction results on a test dataset were not satisfactory.
Which of the following approaches should the Data Science team take to mitigate this issue? (Select TWO.)

- A. indicate a copy of the samples in the test database in the training dataset
- B. Generate more positive samples by duplicating the positive samples and adding a small amount of noise to the duplicated data.
- C. Change the cost function so that false positives have a higher impact on the cost value than false negatives
- D. Change the cost function so that false negatives have a higher impact on the cost value than false positives
- E. Add more deep trees to the random forest to enable the model to learn more features.

Answer: B,D

Explanation:
The Data Science team is facing a problem of imbalanced data, where the positive class (paid users) is much less frequent than the negative class (non-paid users). This can cause the random forest model to be biased towards the majority class and have poor performance on the minority class. To mitigate this issue, the Data Science team can try the following approaches:
* C. Generate more positive samples by duplicating the positive samples and adding a small amount of noise to the duplicated data. This is a technique called data augmentation, which can help increase the size and diversity of the training data for the minority class. This can help the random forest model learn more features and patterns from the positive class and reduce the imbalance ratio.
* D. Change the cost function so that false negatives have a higher impact on the cost value than false positives. This is a technique called cost-sensitive learning, which can assign different weights or costs to different classes or errors. By assigning a higher cost to false negatives (predicting non-paid when the user is actually paid), the random forest model can be more sensitive to the minority class and try to minimize the misclassification of the positive class.
Bagging and Random Forest for Imbalanced Classification
Surviving in a Random Forest with Imbalanced Datasets
machine learning - random forest for imbalanced data? - Cross Validated Biased Random Forest For Dealing With the Class Imbalance Problem

NEW QUESTION # 305
A credit card company wants to identify fraudulent transactions in real time. A data scientist builds a machine learning model for this purpose. The transactional data is captured and stored in Amazon S3. The historic data is already labeled with two classes: fraud (positive) and fair transactions (negative). The data scientist removes all the missing data and builds a classifier by using the XGBoost algorithm in Amazon SageMaker. The model produces the following results:
* True positive rate (TPR): 0.700
* False negative rate (FNR): 0.300
* True negative rate (TNR): 0.977
* False positive rate (FPR): 0.023
* Overall accuracy: 0.949
Which solution should the data scientist use to improve the performance of the model?

- A. Apply the Synthetic Minority Oversampling Technique (SMOTE) on the majority class in the training dataset. Retrain the model with the updated training data.
- B. Apply the Synthetic Minority Oversampling Technique (SMOTE) on the minority class in the training dataset. Retrain the model with the updated training data.
- C. Undersample the minority class.
- D. Oversample the majority class.

Answer: B

Explanation:

Explanation
The solution that the data scientist should use to improve the performance of the model is to apply the Synthetic Minority Oversampling Technique (SMOTE) on the minority class in the training dataset, and retrain the model with the updated training data. This solution can address the problem of class imbalance in the dataset, which can affect the model's ability to learn from the rare but important positive class (fraud).
Class imbalance is a common issue in machine learning, especially for classification tasks. It occurs when one class (usually the positive or target class) is significantly underrepresented in the dataset compared to the other class (usually the negative or non-target class). For example, in the credit card fraud detection problem, the positive class (fraud) is much less frequent than the negative class (fair transactions). This can cause the model to be biased towards the majority class, and fail to capture the characteristics and patterns of the minority class. As a result, the model may have a high overall accuracy, but a low recall or true positive rate for the minority class, which means it misses many fraudulent transactions.
SMOTE is a technique that can help mitigate the class imbalance problem by generating synthetic samples for the minority class. SMOTE works by finding the k-nearest neighbors of each minority class instance, and randomly creating new instances along the line segments connecting them. This way, SMOTE can increase the number and diversity of the minority class instances, without duplicating or losing any information. By applying SMOTE on the minority class in the training dataset, the data scientist can balance the classes and improve the model's performance on the positive class1.
The other options are either ineffective or counterproductive. Applying SMOTE on the majority class would not balance the classes, but increase the imbalance and the size of the dataset. Undersampling the minority class would reduce the number of instances available for the model to learn from, and potentially lose some important information. Oversampling the majority class would also increase the imbalance and the size of the dataset, and introduce redundancy and overfitting.
References:
1: SMOTE for Imbalanced Classification with Python - Machine Learning Mastery

NEW QUESTION # 306
An office security agency conducted a successful pilot using 100 cameras installed at key locations within the main office. Images from the cameras were uploaded to Amazon S3 and tagged using Amazon Rekognition, and the results were stored in Amazon ES. The agency is now looking to expand the pilot into a full production system using thousands of video cameras in its office locations globally. The goal is to identify activities performed by non-employees in real time.
Which solution should the agency consider?

- A. Install AWS DeepLens cameras and use the DeepLens_Kinesis_Video module to stream video to Amazon Kinesis Video Streams for each camera. On each stream, use Amazon Rekognition Video and create a stream processor to detect faces from a collection on each stream, and alert when nonemployees are detected.
- B. Install AWS DeepLens cameras and use the DeepLens_Kinesis_Video module to stream video to Amazon Kinesis Video Streams for each camera. On each stream, run an AWS Lambda function to capture image fragments and then call Amazon Rekognition Image to detect faces from a collection of known employees, and alert when non-employees are detected.
- C. Use a proxy server at each local office and for each camera, and stream the RTSP feed to a unique Amazon Kinesis Video Streams video stream. On each stream, use Amazon Rekognition Video and create a stream processor to detect faces from a collection of known employees, and alert when non-employees are detected.
- D. Use a proxy server at each local office and for each camera, and stream the RTSP feed to a unique Amazon Kinesis Video Streams video stream. On each stream, use Amazon Rekognition Image to detect faces from a collection of known employees and alert when non-employees are detected.

Answer: C

Explanation:
The solution that the agency should consider is to use a proxy server at each local office and for each camera, and stream the RTSP feed to a unique Amazon Kinesis Video Streams video stream. On each stream, use Amazon Rekognition Video and create a stream processor to detect faces from a collection of known employees, and alert when non-employees are detected.
This solution has the following advantages:
It can handle thousands of video cameras in real time, as Amazon Kinesis Video Streams can scale elastically to support any number of producers and consumers1.
It can leverage the Amazon Rekognition Video API, which is designed and optimized for video analysis, and can detect faces in challenging conditions such as low lighting, occlusions, and different poses2.
It can use a stream processor, which is a feature of Amazon Rekognition Video that allows you to create a persistent application that analyzes streaming video and stores the results in a Kinesis data stream3. The stream processor can compare the detected faces with a collection of known employees, which is a container for persisting faces that you want to search for in the input video stream4.
The stream processor can also send notifications to Amazon Simple Notification Service (Amazon SNS) when non-employees are detected, which can trigger downstream actions such as sending alerts or storing the events in Amazon Elasticsearch Service (Amazon ES)3.

References:
1: What Is Amazon Kinesis Video Streams? - Amazon Kinesis Video Streams
2: Detecting and Analyzing Faces - Amazon Rekognition
3: Using Amazon Rekognition Video Stream Processor - Amazon Rekognition
4: Working with Stored Faces - Amazon Rekognition

NEW QUESTION # 307
......

The users of our AWS-Certified-Machine-Learning-Specialty exam questions log on to their account on the platform, at the same time to choose what they want to attend the exam simulation questions, the AWS-Certified-Machine-Learning-Specialty exam questions are automatically for the user presents the same as the actual test environment simulation AWS-Certified-Machine-Learning-Specialty test system, the software built-in timer function can help users better control over time, so as to achieve the systematic, keep up, as well as to improve the user's speed to solve the problem from the side with our AWS-Certified-Machine-Learning-Specialty test guide.

AWS-Certified-Machine-Learning-Specialty Valuable Feedback: https://www.realvce.com/AWS-Certified-Machine-Learning-Specialty_free-dumps.html

As a result, you will be full of confidence and pass the Amazon AWS-Certified-Machine-Learning-Specialty exam will be just a piece of cake, 100% passing guarantee withAmazon AWS-Certified-Machine-Learning-Specialty Valuable Feedback, AWS-Certified-Machine-Learning-Specialty certification is key to high job positions and recognized as elite appraisal standard, All the contents of the AWS-Certified-Machine-Learning-Specialty Valuable Feedback training pdf are compiled by our leading experts, So our AWS-Certified-Machine-Learning-Specialty exam braindumps are triumph of their endeavor.

Note from the Authors, Double-Click the Fill Handle to Copy a Formula, As a result, you will be full of confidence and pass the Amazon AWS-Certified-Machine-Learning-Specialty Exam will be just a piece of cake.

100% passing guarantee withAmazon, AWS-Certified-Machine-Learning-Specialty certification is key to high job positions and recognized as elite appraisal standard, All the contents of the AWS Certified Machine Learning training pdf are compiled by our leading experts.

Quiz 2026 Amazon High Hit-Rate AWS-Certified-Machine-Learning-Specialty: Sample AWS Certified Machine Learning - Specialty Questions Pdf

So our AWS-Certified-Machine-Learning-Specialty exam braindumps are triumph of their endeavor.

- Amazon Sample AWS-Certified-Machine-Learning-Specialty Questions Pdf: AWS Certified Machine Learning - Specialty - www.pdfdumps.com Reliable Planform □ Download ▹ AWS-Certified-Machine-Learning-Specialty ◃ for free by simply searching on （ www.pdfdumps.com ） □Instant AWS-Certified-Machine-Learning-Specialty Discount
- Reliable AWS-Certified-Machine-Learning-Specialty Test Book ✍ Reliable AWS-Certified-Machine-Learning-Specialty Test Book □ New AWS-Certified-Machine-Learning-Specialty Exam Topics □ Search for ➥ AWS-Certified-Machine-Learning-Specialty □ and easily obtain a free download on 【 www.pdfvce.com 】 □Instant AWS-Certified-Machine-Learning-Specialty Discount
- Useful Sample AWS-Certified-Machine-Learning-Specialty Questions Pdf | AWS-Certified-Machine-Learning-Specialty 100% Free Valuable Feedback □ Search for □ AWS-Certified-Machine-Learning-Specialty □ and download exam materials for free through （ www.torrentvce.com ） □Instant AWS-Certified-Machine-Learning-Specialty Discount
- Useful Sample AWS-Certified-Machine-Learning-Specialty Questions Pdf | AWS-Certified-Machine-Learning-Specialty 100% Free Valuable Feedback □ Open □ www.pdfvce.com □ enter ▹ AWS-Certified-Machine-Learning-Specialty ◃ and obtain a free download □Reliable AWS-Certified-Machine-Learning-Specialty Test Book
- AWS-Certified-Machine-Learning-Specialty Exam Paper Pdf □ AWS-Certified-Machine-Learning-Specialty Practice Guide □ Exam AWS-Certified-Machine-Learning-Specialty Questions Answers □ Download “ AWS-Certified-Machine-Learning-Specialty ” for free by simply entering □ www.prep4away.com □ website □Relevant AWS-Certified-Machine-Learning-Specialty Exam Dumps
- Test AWS-Certified-Machine-Learning-Specialty Answers □ AWS-Certified-Machine-Learning-Specialty Practice Guide □ AWS-Certified-Machine-Learning-Specialty VCE Exam Simulator ✔ The page for free download of ➽ AWS-Certified-Machine-Learning-Specialty □ on ➽ www.pdfvce.com □ will open immediately □AWS-Certified-Machine-Learning-Specialty Practice Guide

- Free PDF 2026 Amazon AWS-Certified-Machine-Learning-Specialty: Latest Sample AWS Certified Machine Learning - Specialty Questions Pdf □ Search for （ AWS-Certified-Machine-Learning-Specialty ） and easily obtain a free download on ➡ www.examcollectionpass.com □□□ □AWS-Certified-Machine-Learning-Specialty Exam Paper Pdf
- Pdfvce Offers Actual and Updated Amazon AWS-Certified-Machine-Learning-Specialty Practice Questions □ Easily obtain free download of ▷ AWS-Certified-Machine-Learning-Specialty ◁ by searching on { www.pdfvce.com } □New AWS-Certified-Machine-Learning-Specialty Dumps Ppt
- Useful Sample AWS-Certified-Machine-Learning-Specialty Questions Pdf | AWS-Certified-Machine-Learning-Specialty 100% Free Valuable Feedback □ Simply search for ➤ AWS-Certified-Machine-Learning-Specialty □ for free download on □ www.exam4labs.com □ □Testking AWS-Certified-Machine-Learning-Specialty Learning Materials
- AWS-Certified-Machine-Learning-Specialty Reliable Braindumps Questions □ AWS-Certified-Machine-Learning-Specialty Training Materials □ Certification AWS-Certified-Machine-Learning-Specialty Dumps □ Download □ AWS-Certified-Machine-Learning-Specialty □ for free by simply entering ▶ www.pdfvce.com ◀ website □AWS-Certified-Machine-Learning-Specialty Practice Guide
- Useful Sample AWS-Certified-Machine-Learning-Specialty Questions Pdf | AWS-Certified-Machine-Learning-Specialty 100% Free Valuable Feedback □ ▷ www.practicevce.com ◁ is best website to obtain 【 AWS-Certified-Machine-Learning-Specialty 】 for free download □AWS-Certified-Machine-Learning-Specialty New Dumps Sheet
- 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, www.stes.tyc.edu.tw, www.stes.tyc.edu.tw, www.stes.tyc.edu.tw, www.stes.tyc.edu.tw, www.stes.tyc.edu.tw, www.stes.tyc.edu.tw, www.stes.tyc.edu.tw, www.stes.tyc.edu.tw, www.stes.tyc.edu.tw, Disposable vapes

P.S. Free & New AWS-Certified-Machine-Learning-Specialty dumps are available on Google Drive shared by RealVCE: https://drive.google.com/open?id=1p84IpLLRxqYJcwy5uua7dBLAl7MtGk7g