Formats of Amazon AWS-Certified-Machine-Learning-Specialty Practice Exam Questions

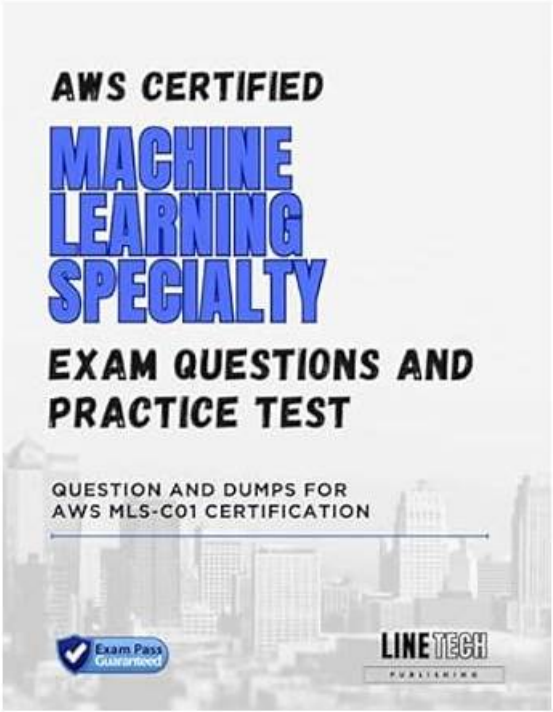

P.S. Free & New AWS-Certified-Machine-Learning-Specialty dumps are available on Google Drive shared by PracticeDump: https://drive.google.com/open?id=1h7arazcB4THJ3wy6tr2xgVlHhPcVN_Cs

With passing rate more than 98 percent from exam candidates who chose our Amazon AWS-Certified-Machine-Learning-Specialty Study Guide, we have full confidence that your AWS-Certified-Machine-Learning-Specialty actual test will be a piece of cake by them. Our AWS Certified Machine Learning - Specialty exam questions provide with the software which has a variety of self-study and self-assessment functions to detect learning results.

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.

The AWS Certified Machine Learning - Specialty certification exam is a highly sought-after credential for professionals seeking to demonstrate their expertise in implementing and deploying machine learning solutions on the Amazon Web Services (AWS) platform. AWS Certified Machine Learning - Specialty certification is designed for individuals with a background in data science, software engineering, and other related fields who are interested in leveraging AWS to build and deploy machine learning models.

The AWS Certified Machine Learning - Specialty exam is a certification offered by Amazon Web Services (AWS) for professionals who want to validate their expertise in machine learning. AWS-Certified-Machine-Learning-Specialty Exam is designed for individuals who have a solid understanding of machine learning concepts and techniques, as well as experience using AWS services to build and deploy machine learning solutions. AWS-Certified-Machine-Learning-Specialty exam covers a range of topics, including data preparation, model training and evaluation, deployment and implementation, and automation.

AWS-Certified-Machine-Learning-Specialty Test Fee, Latest AWS-Certified-Machine-Learning-Specialty Mock Exam

The AWS-Certified-Machine-Learning-Specialty exam prepare of our website is completed by experts who has a good understanding of real exams and have many years of experience writing AWS-Certified-Machine-Learning-Specialty study materials. They know very well what candidates really need most when they prepare for the exam. They also understand the real exam situation very well. So they compiled AWS-Certified-Machine-Learning-Specialty Exam prepare that they hope to do their utmost to help candidates pass the exam and get what job they want. They apply to exam candidates of different level of computer industry. So whichever degree you are at, you can utilize our AWS-Certified-Machine-Learning-Specialty study materials tool with following traits.

Amazon AWS Certified Machine Learning - Specialty Sample Questions (Q203-Q208):

NEW QUESTION # 203
Amazon Connect has recently been tolled out across a company as a contact call center The solution has been configured to store voice call recordings on Amazon S3 The content of the voice calls are being analyzed for the incidents being discussed by the call operators Amazon Transcribe is being used to convert the audio to text, and the output is stored on Amazon S3 Which approach will provide the information required for further analysis?

- A. Use Amazon Comprehend with the transcribed files to build the key topics
- B. Use the Amazon SageMaker k-Nearest-Neighbors (kNN) algorithm on the transcribed files to generate a word embeddings dictionary for the key topics
- C. Use the AWS Deep Learning AMI with Gluon Semantic Segmentation on the transcribed files to train and build a model for the key topics
- D. Use Amazon Translate with the transcribed files to train and build a model for the key topics

Answer: A

Explanation:
Amazon Comprehend is a natural language processing (NLP) service that uses machine learning to find insights and relationships in text. It can analyze text documents and identify the key topics, entities, sentiments, languages, and more. In this case, Amazon Comprehend can be used with the transcribed files from Amazon Transcribe to extract the main topics that are being discussed by the call operators. This can help to understand the common issues and concerns of the customers, and provide insights for further analysis and improvement. References:
Amazon Comprehend - Amazon Web Services
AWS Certified Machine Learning - Specialty Sample Questions

NEW QUESTION # 204
A Machine Learning Specialist discover the following statistics while experimenting on a model.

What can the Specialist from the experiments?

- A. The model in Experiment 1 had a high bias error and a high variance error that were reduced in Experiment 3 by regularization Experiment 2 shows thai high bias cannot be reduced by increasing layers and neurons in the model
- B. The model In Experiment 1 had a high variance error lhat was reduced in Experiment 3 by regularization Experiment 2 shows that there is minimal bias error in Experiment 1
- C. The model in Experiment 1 had a high bias error that was reduced in Experiment 3 by regularization Experiment 2 shows

that there is minimal variance error in Experiment 1

- D. The model in Experiment 1 had a high random noise error that was reduced in Experiment 3 by regularization Experiment 2 shows that random noise cannot be reduced by increasing layers and neurons in the model

Answer: B

Explanation:
Explanation
The model in Experiment 1 had a high variance error because it performed well on the training data (train error = 5%) but poorly on the test data (test error = 8%). This indicates that the model was overfitting the training data and not generalizing well to new data. The model in Experiment 3 had a lower variance error because it performed similarly on the training data (train error = 5.1%) and the test data (test error = 5.4%). This indicates that the model was more robust and less sensitive to the fluctuations in the training data. The model in Experiment 3 achieved this improvement by implementing regularization, which is a technique that reduces the complexity of the model and prevents overfitting by adding a penalty term to the loss function. The model in Experiment 2 had a minimal bias error because it performed similarly on the training data (train error = 5.2%) and the test data (test error = 5.7%) as the model in Experiment 1. This indicates that the model was not underfitting the data and capturing the true relationship between the input and output variables. The model in Experiment 2 increased the number of layers and neurons in the model, which is a way to increase the complexity and flexibility of the model. However, this did not improve the performance of the model, as the variance error remained high. This shows that increasing the complexity of the model is not always the best way to reduce the bias error, and may even increase the variance error if the model becomes too complex for the data.
References:
Bias Variance Tradeoff - Clearly Explained - Machine Learning Plus
The Bias-Variance Trade-off in Machine Learning - Stack Abuse

NEW QUESTION # 205
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 # 206
A large mobile network operating company is building a machine learning model to predict customers who are likely to unsubscribe from the service. The company plans to offer an incentive for these customers as the cost of churn is far greater than the cost of the incentive.
The model produces the following confusion matrix after evaluating on a test dataset of 100 customers:
Based on the model evaluation results, why is this a viable model for production?

- A. The model is 86% accurate and the cost incurred by the company as a result of false negatives is less than the false positives.
- B. The precision of the model is 86%, which is greater than the accuracy of the model.
- C. The precision of the model is 86%, which is less than the accuracy of the model.
- D. The model is 86% accurate and the cost incurred by the company as a result of false positives is less than the false negatives.

Answer: D

Explanation:
Based on the model evaluation results, this is a viable model for production because the model is 86% accurate and the cost incurred by the company as a result of false positives is less than the false negatives. The accuracy of the model is the proportion of correct predictions out of the total predictions, which can be calculated by adding the true positives and true negatives and dividing by the total number of observations. In this case, the accuracy of the model is (10 + 76) / 100 = 0.86, which means that the model correctly predicted
86% of the customers' churn status. The cost incurred by the company as a result of false positives and false negatives is the loss or damage that the company suffers when the model makes incorrect predictions. A false positive is when the model predicts that a customer will churn, but the customer actually does not churn. A false negative is when the model predicts that a customer will not churn, but the customer actually churns. In this case, the cost of a false positive is the incentive that the company offers to the customer who is predicted to churn, which is a relatively low cost. The cost of a false negative is the revenue that the company loses when the customer churns, which is a relatively high cost. Therefore, the cost of a false positive is less than the cost of a false negative, and the company would prefer to have more false positives than false negatives.
The model has 10 false positives and 4 false negatives, which means that the company's cost is lower than if the model had more false negatives and fewer false positives.

NEW QUESTION # 207
A global bank requires a solution to predict whether customers will leave the bank and choose another bank.
The bank is using a dataset to train a model to predict customer loss. The training dataset has 1,000 rows. The training dataset includes 100 instances of customers who left the bank.
A machine learning (ML) specialist is using Amazon SageMaker Data Wrangler to train a churn prediction model by using a SageMaker training job. After training, the ML specialist notices that the model returns only false results. The ML specialist must correct the model so that it returns more accurate predictions.
Which solution will meet these requirements?

- A. Apply undersampling to the training dataset before training.
- B. Apply anomaly detection to remove outliers from the training dataset before training.
- C. Apply normalization to the features of the training dataset before training.
- D. Apply Synthetic Minority Oversampling Technique (SMOTE) to the training dataset before training.

Answer: D

Explanation:
The best solution to meet the requirements is to apply Synthetic Minority Oversampling Technique (SMOTE) to the training dataset

before training. SMOTE is a technique that generates synthetic samples for the minority class by interpolating between existing samples. This can help balance the class distribution and provide more information to the model. SMOTE can improve the performance of the model on the minority class, which is the class of interest in churn prediction. SMOTE can be applied using the SageMaker Data Wrangler, which provides a built-in analysis for oversampling the minority class1.
The other options are not effective solutions for the problem. Applying anomaly detection to remove outliers from the training dataset before training may not improve the model's accuracy, as outliers may not be the main cause of the false results. Moreover, removing outliers may reduce the diversity of the data and make the model less robust. Applying normalization to the features of the training dataset before training may improve the model's convergence and stability, but it does not address the class imbalance issue. Normalization can also be applied using the SageMaker Data Wrangler, which provides a built-in transformation for scaling the features2. Applying undersampling to the training dataset before training may reduce the class imbalance, but it also discards potentially useful information from the majority class. Undersampling can also result in underfitting and high bias for the model.

NEW QUESTION # 208
.....

AWS-Certified-Machine-Learning-Specialty preparation materials will be the good helper for your qualification certification. We are concentrating on providing high-quality authorized AWS-Certified-Machine-Learning-Specialty study guide all over the world so that you can clear exam one time. AWS-Certified-Machine-Learning-Specialty reliable exam bootcamp materials contain three formats: PDF version, Soft test engine and APP test engine so that our products are enough to satisfy different candidates' habits and cover nearly full questions & answers of the real AWS-Certified-Machine-Learning-Specialty test.

□ Passing AWS-Certified-Machine-Learning-Specialty Score Feedback □ Open { www.pdfvce.com } enter ⇛ AWS-Certified-Machine-Learning-Specialty ⇚ and obtain a free download □AWS-Certified-Machine-Learning-Specialty New Learning Materials

- New AWS-Certified-Machine-Learning-Specialty Dumps Ppt □ New AWS-Certified-Machine-Learning-Specialty Dumps Ebook □ AWS-Certified-Machine-Learning-Specialty Valid Test Blueprint □ Download ➤ AWS-Certified-Machine-Learning-Specialty □ for free by simply entering [www.vceengine.com] website □AWS-Certified-Machine-Learning-Specialty Trustworthy Exam Content
- myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, ncon.edu.sa, www.stes.tyc.edu.tw, www.stes.tyc.edu.tw, ncon.edu.sa, 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, myportal.utt.edu.tt, pct.edu.pk, Disposable vapes

DOWNLOAD the newest PracticeDump AWS-Certified-Machine-Learning-Specialty PDF dumps from Cloud Storage for free: https://drive.google.com/open?id=1h7arazcB4THJ3wy6tr2xgVlHhPcVN_Cs