AWS-Certified-Machine-Learning-Specialty Exam Braindumps: AWS Certified Machine Learning - Specialty & AWS-Certified-Machine-Learning-Specialty Certification Training

DOWNLOAD the newest ActualtestPDF AWS-Certified-Machine-Learning-Specialty PDF dumps from Cloud Storage for free: https://drive.google.com/open?id=1tM6o26HHoKH3sZay2BidsjBiLUDfnJ5E

There are many advantages of our product and it is worthy for you to buy it. You can download and try out our AWS-Certified-Machine-Learning-Specialty guide questions demo before the purchase and use them immediately after you pay for them successfully. Once you pay for it, we will send to you within 5-10 minutes. Then you can learn and practice it. We update the AWS-Certified-Machine-Learning-Specialty Torrent question frequently and provide the discounts to the old client. We check the update every day, once we update, we will send it to you as soon as possible. There are many benefits to buy AWS-Certified-Machine-Learning-Specialty guide torrent such as after the client pass the exam they can enter in the big company and double their wages.

To take the AWS Certified Machine Learning - Specialty exam, candidates must have a minimum of one year of experience in designing and implementing machine learning models on the AWS platform. They should also have a strong understanding of machine learning algorithms, data modeling, and data processing techniques. AWS-Certified-Machine-Learning-Specialty Exam consists of 65 multiple-choice and multiple-response questions, and candidates have 3 hours to complete it.

>> New AWS-Certified-Machine-Learning-Specialty Test Book <<

Professional New AWS-Certified-Machine-Learning-Specialty Test Book Provide Prefect Assistance in AWS-Certified-Machine-Learning-Specialty Preparation

Please select our ActualtestPDF to achieve good results in order to pass Amazon certification AWS-Certified-Machine-Learning-Specialty exam, and you will not regret doing so. It is worth spending a little money to get so much results. Our ActualtestPDF can not only give you a good exam preparation, allowing you to pass Amazon Certification AWS-Certified-Machine-Learning-Specialty Exam, but also provide you with one-year free update service.

Amazon AWS Certified Machine Learning - Specialty Sample Questions (Q125-Q130):

NEW QUESTION # 125
A manufacturing company has structured and unstructured data stored in an Amazon S3 bucket. A Machine Learning Specialist wants to use SQL to run queries on this data.
Which solution requires the LEAST effort to be able to query this data?

- A. Use AWS Lambda to transform the data and Amazon Kinesis Data Analytics to run queries.
- B. Use AWS Glue to catalogue the data and Amazon Athena to run queries.
- C. Use AWS Batch to run ETL on the data and Amazon Aurora to run the queries.
- D. Use AWS Data Pipeline to transform the data and Amazon RDS to run queries.

Answer: B

Explanation:
Using AWS Glue to catalogue the data and Amazon Athena to run queries is the solution that requires the least effort to be able to query the data stored in an Amazon S3 bucket using SQL. AWS Glue is a service that provides a serverless data integration platform for data preparation and transformation. AWS Glue can automatically discover, crawl, and catalogue the data stored in various sources, such as Amazon S3, Amazon RDS, Amazon Redshift, etc. AWS Glue can also use AWS KMS to encrypt the data at rest on the Glue Data Catalog and Glue ETL jobs. AWS Glue can handle both structured and unstructured data, and support various data formats, such as CSV, JSON, Parquet, etc. AWS Glue can also use built-in or custom classifiers to identify and parse the data schema and format1 Amazon Athena is a service that provides an interactive query engine that can run SQL queries directly on data stored in Amazon S3. Amazon Athena can integrate with AWS Glue to use the Glue Data Catalog as a central metadata repository for the data sources and tables. Amazon Athena can also use AWS KMS to encrypt the data at rest on Amazon S3 and the query results. Amazon Athena can query both structured and unstructured data, and support various data formats, such as CSV, JSON, Parquet, etc. Amazon Athena can also use partitions and compression to optimize the query performance and reduce the query cost23 The other options are not valid or require more effort to query the data stored in an Amazon S3 bucket using SQL. Using AWS Data Pipeline to transform the data and Amazon RDS to run queries is not a good option, as it involves moving the data from Amazon S3 to Amazon RDS, which can incur additional time and cost. AWS Data Pipeline is a service that can orchestrate and automate data movement and transformation across various AWS services and on-premises data sources. AWS Data Pipeline can be integrated with Amazon EMR to run ETL jobs on the data stored in Amazon S3. Amazon RDS is a service that provides a managed relational database service that can run various database engines, such as MySQL, PostgreSQL, Oracle, etc. Amazon RDS can use AWS KMS to encrypt the data at rest and in transit. Amazon RDS can run SQL queries on the data stored in the database tables45 Using AWS Batch to run ETL on the data and Amazon Aurora to run the queries is not a good option, as it also involves moving the data from Amazon S3 to Amazon Aurora, which can incur additional time and cost. AWS Batch is a service that can run batch computing workloads on AWS. AWS Batch can be integrated with AWS Lambda to trigger ETL jobs on the data stored in Amazon S3. Amazon Aurora is a service that provides a compatible and scalable relational database engine that can run MySQL or PostgreSQL. Amazon Aurora can use AWS KMS to encrypt the data at rest and in transit. Amazon Aurora can run SQL queries on the data stored in the database tables. Using AWS Lambda to transform the data and Amazon Kinesis Data Analytics to run queries is not a good option, as it is not suitable for querying data stored in Amazon S3 using SQL. AWS Lambda is a service that can run serverless functions on AWS. AWS Lambda can be integrated with Amazon S3 to trigger data transformation functions on the data stored in Amazon S3. Amazon Kinesis Data Analytics is a service that can analyze streaming data using SQL or Apache Flink. Amazon Kinesis Data Analytics can be integrated with Amazon Kinesis Data Streams or Amazon Kinesis Data Firehose to ingest streaming data sources, such as web logs, social media, IoT devices, etc. Amazon Kinesis Data Analytics is not designed for querying data stored in Amazon S3 using SQL.

NEW QUESTION # 126
A beauty supply store wants to understand some characteristics of visitors to the store. The store has security video recordings from the past several years. The store wants to generate a report of hourly visitors from the recordings. The report should group visitors by hair style and hair color.
Which solution will meet these requirements with the LEAST amount of effort?

- A. Use a semantic segmentation algorithm to identify a visitor's hair in video frames. Pass the identified hair to an XGBoost algorithm to determine hair style and hair.
- B. Use an object detection algorithm to identify a visitor's hair in video frames. Pass the identified hair to an ResNet-50 algorithm to determine hair style and hair color.
- C. Use an object detection algorithm to identify a visitor's hair in video frames. Pass the identified hair to an XGBoost algorithm to determine hair style and hair color.

- D. Use a semantic segmentation algorithm to identify a visitor's hair in video frames. Pass the identified hair to an ResNet-50 algorithm to determine hair style and hair color.

Answer: D

Explanation:
Explanation
The solution that will meet the requirements with the least amount of effort is to use a semantic segmentation algorithm to identify a visitor's hair in video frames, and pass the identified hair to an ResNet-50 algorithm to determine hair style and hair color. This solution can leverage the existing Amazon SageMaker algorithms and frameworks to perform the tasks of hair segmentation and classification.
Semantic segmentation is a computer vision technique that assigns a class label to every pixel in an image, such that pixels with the same label share certain characteristics. Semantic segmentation can be used to identify and isolate different objects or regions in an image, such as a visitor's hair in a video frame. Amazon SageMaker provides a built-in semantic segmentation algorithm that can train and deploy models for semantic segmentation tasks. The algorithm supports three state-of-the-art network architectures: Fully Convolutional Network (FCN), Pyramid Scene Parsing Network (PSP), and DeepLab v3. The algorithm can also use pre-trained or randomly initialized ResNet-50 or ResNet-101 as the backbone network. The algorithm can be trained using P2/P3 type Amazon EC2 instances in single machine configurations1.
ResNet-50 is a convolutional neural network that is 50 layers deep and can classify images into 1000 object categories. ResNet-50 is trained on more than a million images from the ImageNet database and can achieve high accuracy on various image recognition tasks. ResNet-50 can be used to determine hair style and hair color from the segmented hair regions in the video frames. Amazon SageMaker provides a built-in image classification algorithm that can use ResNet-50 as the network architecture. The algorithm can also perform transfer learning by fine-tuning the pre-trained ResNet-50 model with new data. The algorithm can be trained using P2/P3 type Amazon EC2 instances in single or multiple machine configurations2.
The other options are either less effective or more complex to implement. Using an object detection algorithm to identify a visitor's hair in video frames would not segment the hair at the pixel level, but only draw bounding boxes around the hair regions. This could result in inaccurate or incomplete hair segmentation, especially if the hair is occluded or has irregular shapes. Using an XGBoost algorithm to determine hair style and hair color would require transforming the segmented hair images into numerical features, which could lose some information or introduce noise. XGBoost is also not designed for image classification tasks, and may not achieve high accuracy or performance.
References:
1: Semantic Segmentation Algorithm - Amazon SageMaker
2: Image Classification Algorithm - Amazon SageMaker

NEW QUESTION # 127
An insurance company is developing a new device for vehicles that uses a camera to observe drivers' behavior and alert them when they appear distracted The company created approximately 10,000 training images in a controlled environment that a Machine Learning Specialist will use to train and evaluate machine learning models During the model evaluation the Specialist notices that the training error rate diminishes faster as the number of epochs increases and the model is not accurately inferring on the unseen test images Which of the following should be used to resolve this issue? (Select TWO)

- A. Use gradient checking in the model
- B. Perform data augmentation on the training data
- C. Add vanishing gradient to the model
- D. Make the neural network architecture complex.
- E. Add L2 regularization to the model

Answer: A,B

NEW QUESTION # 128
A Data Scientist is building a model to predict customer churn using a dataset of 100 continuous numerical features. The Marketing team has not provided any insight about which features are relevant for churn prediction. The Marketing team wants to interpret the model and see the direct impact of relevant features on the model outcome. While training a logistic regression model, the Data Scientist observes that there is a wide gap between the training and validation set accuracy.
Which methods can the Data Scientist use to improve the model performance and satisfy the Marketing team's needs? (Choose two.)

- A. Perform linear discriminant analysis
- B. Perform recursive feature elimination

- C. Add features to the dataset
- D. Perform t-distributed stochastic neighbor embedding (t-SNE)
- E. Add L1 regularization to the classifier

Answer: B,E

Explanation:
Explanation
The Data Scientist is building a model to predict customer churn using a dataset of 100 continuous numerical features. The Marketing team wants to interpret the model and see the direct impact of relevant features on the model outcome. However, the Data Scientist observes that there is a wide gap between the training and validation set accuracy, which indicates that the model is overfitting the data and generalizing poorly to new data.
To improve the model performance and satisfy the Marketing team's needs, the Data Scientist can use the following methods:
Add L1 regularization to the classifier: L1 regularization is a technique that adds a penalty term to the loss function of the logistic regression model, proportional to the sum of the absolute values of the coefficients. L1 regularization can help reduce overfitting by shrinking the coefficients of the less important features to zero, effectively performing feature selection. This can simplify the model and make it more interpretable, as well as improve the validation accuracy.
Perform recursive feature elimination: Recursive feature elimination (RFE) is a feature selection technique that involves training a model on a subset of the features, and then iteratively removing the least important features one by one until the desired number of features is reached. The idea behind RFE is to determine the contribution of each feature to the model by measuring how well the model performs when that feature is removed. The features that are most important to the model will have the greatest impact on performance when they are removed. RFE can help improve the model performance by eliminating the irrelevant or redundant features that may cause noise or multicollinearity in the data. RFE can also help the Marketing team understand the direct impact of the relevant features on the model outcome, as the remaining features will have the highest weights in the model.
References:
Regularization for Logistic Regression
Recursive Feature Elimination

NEW QUESTION # 129
A company is using Amazon Textract to extract textual data from thousands of scanned text-heavy legal documents daily. The company uses this information to process loan applications automatically. Some of the documents fail business validation and are returned to human reviewers, who investigate the errors. This activity increases the time to process the loan applications.
What should the company do to reduce the processing time of loan applications?

- A. Use an Amazon Textract synchronous operation instead of an asynchronous operation.
- B. Configure Amazon Textract to route low-confidence predictions to Amazon SageMaker Ground Truth. Perform a manual review on those words before performing a business validation.
- C. Use Amazon Rekognition's feature to detect text in an image to extract the data from scanned images. Use this information to process the loan applications.
- D. Configure Amazon Textract to route low-confidence predictions to Amazon Augmented AI (Amazon A2I). Perform a manual review on those words before performing a business validation.

Answer: D

Explanation:
The company should configure Amazon Textract to route low-confidence predictions to Amazon Augmented AI (Amazon A2I). Amazon A2I is a service that allows you to implement human review of machine learning (ML) predictions. It also comes integrated with some of the Artificial Intelligence (AI) services such as Amazon Textract. By using Amazon A2I, the company can perform a manual review on those words that have low confidence scores before performing a business validation. This will help reduce the processing time of loan applications by avoiding errors and rework.
Option A is incorrect because Amazon SageMaker Ground Truth is not a suitable service for human review of Amazon Textract predictions. Amazon SageMaker Ground Truth is a service that helps you build highly accurate training datasets for machine learning. It allows you to label your own data or use a workforce of human labelers. However, it does not provide an easy way to integrate with Amazon Textract and route low-confidence predictions for human review.
Option B is incorrect because using an Amazon Textract synchronous operation instead of an asynchronous operation will not reduce the processing time of loan applications. A synchronous operation is a request-response operation that returns the results immediately. An asynchronous operation is a start-and-check operation that returns a job identifier that you can use to check the status and results later. The choice of operation depends on the size and complexity of the document, not on the confidence of the predictions.
Option D is incorrect because using Amazon Rekognition's feature to detect text in an image to extract the data from scanned images

is not a better alternative than using Amazon Textract. Amazon Rekognition is a service that provides computer vision capabilities, such as face recognition, object detection, and scene analysis. It can also detect text in an image, but it does not provide the same level of accuracy and functionality as Amazon Textract. Amazon Textract can not only detect text, but also extract data from tables and forms, and understand the layout and structure of the document.
References:
Amazon Augmented AI
Amazon SageMaker Ground Truth
Amazon Textract Operations
Amazon Rekognition

NEW QUESTION # 130
.....

Our AWS-Certified-Machine-Learning-Specialty study materials are very popular in the international market and enjoy wide praise by the people in and outside the circle. We have shaped our AWS-Certified-Machine-Learning-Specialty exam questions into a famous and top-ranking brand and we enjoy well-deserved reputation among the clients. Our AWS-Certified-Machine-Learning-Specialty study materials boost many outstanding and superior advantages which other same kinds of products don't have. The clients can try out and download our study materials before their purchase. They can immediately use our AWS-Certified-Machine-Learning-Specialty training guide after they pay successfully.

- AWS Certified Machine Learning - Specialty latest study torrent - AWS Certified Machine Learning - Specialty reliable vce pdf - AWS Certified Machine Learning - Specialty valid training dumps □ Search for ➽ AWS-Certified-Machine-Learning-Specialty □ and download exam materials for free through □ www.troytecdumps.com □ □AWS-Certified-Machine-Learning-Specialty Questions Answers
- www.stes.tyc.edu.tw, www.stes.tyc.edu.tw, azrasehovic.com, 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, myportal.utt.edu.tt, www.stes.tyc.edu.tw, www.stes.tyc.edu.tw, 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, Disposable vapes

DOWNLOAD the newest ActualtestPDF AWS-Certified-Machine-Learning-Specialty PDF dumps from Cloud Storage for free: https://drive.google.com/open?id=1tM6o26HHoKH3sZay2BidsjBiLUDfnJ5E