[2025-November-NewBraindump2go MLS-C01 Dumps with PDF and VCE Free[Q259-Q290

2025/November Latest Braindump2go MLS-C01 Exam Dumps with PDF and VCE Free Updated Today! Following are some new

Braindump2go MLS-C01 Real Exam Questions!QUESTION 259An ecommerce company is collecting structured data and unstructured data from its website, mobile apps, and IoT devices. The data is stored in several databases and Amazon S3 buckets. The company is implementing a scalable repository to store structured data and unstructured data. The company must implement a solution that provides a central data catalog, self-service access to the data, and granular data access policies and encryption to protect the data. Which combination of actions will meet these requirements with the LEAST amount of setup? (Choose three.) A. Identify the existing data in the databases and S3 buckets. Link the data to AWS Lake Formation.B. Identify the existing data in the databases and S3 buckets. Link the data to AWS Glue.C. Run AWS Glue crawlers on the linked data sources to create a central data catalog.D. Apply granular access policies by using AWS Identity and Access Management (IAM). Configure server-side encryption on each data source.E. Apply granular access policies and encryption by using AWS Lake Formation.F. Apply granular access policies and encryption by using AWS Glue. Answer: ACEExplanation: https://docs.aws.amazon.com/lake-formation/latest/dg/what-is-lake-formation.htmlQUESTION 260A machine learning (ML) specialist is developing a deep learning sentiment analysis model that is based on data from movie reviews. After the ML specialist trains the model and reviews the model results on the validation set, the ML specialist discovers that the model is overfitting. Which solutions will MOST improve the model generalization and reduce overfitting? (Choose three.)A. Shuffle the dataset with a different seed.B. Decrease the learning rate.C. Increase the number of layers in the network.D. Add L1 regularization and L2 regularization.E. Add dropout.F. Decrease the number of layers in the network.Answer: DEFQUESTION 261An online advertising company is developing a linear model to predict the bid price of advertisements in real time with low-latency predictions. A data scientist has trained the linear model by using many features, but the model is overfitting the training dataset. The data scientist needs to prevent overfitting and must reduce the number of features. Which solution will meet these requirements?A. Retrain the model with L1 regularization applied.B. Retrain the model with L2 regularization applied.C. Retrain the model with dropout regularization applied.D. Retrain the model by using more data. Answer: AQUESTION 262A 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: 0.949Which solution should the data scientist use to improve the performance of the model? A. Apply the Synthetic Minority Oversampling Technique (SMOTE) on the minority class in the training dataset. Retrain the model with the updated training data.B. Apply the Synthetic Minority Oversampling Technique (SMOTE) on the majority class in the training dataset. Retrain the model with the updated training data.C. Undersample the minority class.D. Oversample the majority class.Answer: AQUESTION 263A company is training machine learning (ML) models on Amazon SageMaker by using 200 TB of data that is stored in Amazon S3 buckets. The training data consists of individual files that are each larger than 200 MB in size. The company needs a data access solution that offers the shortest processing time and the least amount of setup. Which solution will meet these requirements? A. Use File mode in SageMaker to copy the dataset from the S3 buckets to the ML instance storage.B. Create an Amazon FSx for Lustre file system. Link the file system to the S3 buckets.C. Create an Amazon Elastic File System (Amazon EFS) file system. Mount the file system to the training instances.D. Use FastFile mode in SageMaker to stream the files on demand from the S3 buckets. Answer: DExplanation: For larger datasets with larger files (more than 50 MB per file), the first option is to try fast file mode, which is more straightforward to use than FSx for Lustre because it doesn't require creating a file system, or connecting to a VPC. Fast file mode is ideal for large file containers (more than 150 MB), and might also do well with files more than 50 MB. https://docs.aws.amazon.com/sagemaker/latest/dg/model-access-training-data.html#model-access-training-data-best-practices QUESTION 264An online store is predicting future book sales by using a linear regression model that is based on past sales data. The data includes duration, a numerical feature that represents the number of days that a book has been listed in the online store. A data scientist performs an exploratory data analysis and discovers that the relationship between book sales and duration is skewed and non-linear. Which data transformation step should the data scientist take to improve the predictions of the model? A. One-hot encodingB. Cartesian product transformationC. Quantile binningD. NormalizationAnswer: CExplanation: https://docs.aws.amazon.com/machine-learning/latest/dg/data-transformations-reference.htmlQUESTION 265A company's data

engineer wants to use Amazon S3 to share datasets with data scientists. The data scientists work in three departments: Finance. Marketing, and Human Resources. Each department has its own IAM user group. Some datasets contain sensitive information and should be accessed only by the data scientists from the Finance department. How can the data engineer set up access to meet these requirements? A. Create an S3 bucket for each dataset. Create an ACL for each S3 bucket. For each S3 bucket that contains a sensitive dataset, set the ACL to allow access only from the Finance department user group. Allow all three department user groups to access each S3 bucket that contains a non-sensitive dataset. B. Create an S3 bucket for each dataset. For each S3 bucket that contains a sensitive dataset, set the bucket policy to allow access only from the Finance department user group. Allow all three department user groups to access each S3 bucket that contains a non-sensitive dataset. C. Create a single S3 bucket that includes two folders to separate the sensitive datasets from the non-sensitive datasets. For the Finance department user group, attach an IAM policy that provides access to both folders. For the Marketing and Human Resources department user groups, attach an IAM policy that provides access to only the folder that contains the non-sensitive datasets. Set the policy for the S3 bucket that includes two folders to separate the sensitive datasets from the non-sensitive datasets. Set the policy for the S3 bucket to allow only the Finance department user group to access the folder that contains the sensitive datasets. Allow all three department user groups to access the folder that contains the non-sensitive datasets. Allow all three department user groups to access the folder that contains the non-sensitive datasets. Allow all three department user groups to access the folder that contains the non-sensitive datasets.

https://aws.amazon.com/blogs/security/how-to-restrict-amazon-s3-bucket-access-to-a-specific-iam-role/QUESTION 266A company operates an amusement park. The company wants to collect, monitor, and store real-time traffic data at several park entrances by using strategically placed cameras. The company's security team must be able to immediately access the data for viewing. Stored data must be indexed and must be accessible to the company's data science team. Which solution will meet these requirements MOST cost-effectively?A. Use Amazon Kinesis Video Streams to ingest, index, and store the data. Use the built-in integration with Amazon Rekognition for viewing by the security team.B. Use Amazon Kinesis Video Streams to ingest, index, and store the data. Use the built-in HTTP live streaming (HLS) capability for viewing by the security team.C. Use Amazon Rekognition Video and the GStreamer plugin to ingest the data for viewing by the security team. Use Amazon Kinesis Data Streams to index and store the data.D. Use Amazon Kinesis Data Firehose to ingest, index, and store the data. Use the built-in HTTP live streaming (HLS) capability for viewing by the security team. Answer: BExplanation:

https://aws.amazon.com/about-aws/whats-new/2018/07/kinesis-video-adds-hls-support/QUESTION 267An engraving company wants to automate its quality control process for plaques. The company performs the process before mailing each customized plaque to a customer. The company has created an Amazon S3 bucket that contains images of defects that should cause a plaque to be rejected. Low-confidence predictions must be sent to an internal team of reviewers who are using Amazon Augmented AI (Amazon A2I). Which solution will meet these requirements? A. Use Amazon Textract for automatic processing. Use Amazon A2I with Amazon Mechanical Turk for manual review.B. Use Amazon Rekognition for automatic processing. Use Amazon A2I with a private workforce option for manual review.C. Use Amazon Transcribe for automatic processing. Use Amazon A2I with a private workforce option for manual review.D. Use AWS Panorama for automatic processing. Use Amazon A2I with Amazon Mechanical Turk for manual review. Answer: BQUESTION 268A machine learning (ML) engineer at a bank is building a data ingestion solution to provide transaction features to financial ML models. Raw transactional data is available in an Amazon Kinesis data stream. The solution must compute rolling averages of the ingested data from the data stream and must store the results in Amazon SageMaker Feature Store. The solution also must serve the results to the models in near real time. Which solution will meet these requirements?A. Load the data into an Amazon S3 bucket by using Amazon Kinesis Data Firehose. Use a SageMaker Processing job to aggregate the data and to load the results into SageMaker Feature Store as an online feature group.B. Write the data directly from the data stream into SageMaker Feature Store as an online feature group. Calculate the rolling averages in place within SageMaker Feature Store by using the SageMaker GetRecord API operation.C. Consume the data stream by using an Amazon Kinesis Data Analytics SQL application that calculates the rolling averages. Generate a result stream. Consume the result stream by using a custom AWS Lambda function that publishes the results to SageMaker Feature Store as an online feature group.D. Load the data into an Amazon S3 bucket by using Amazon Kinesis Data Firehose. Use a SageMaker Processing job to load the data into SageMaker Feature Store as an offline feature group. Compute the rolling averages at query time. Answer: CQUESTION 269Each morning, a data scientist at a rental car company creates insights about the previous day's rental car reservation demands. The company needs to automate this process by streaming the data to Amazon S3 in near real time. The solution must detect high-demand rental cars at each of the company's locations. The solution also must create a visualization dashboard that automatically refreshes with the most recent data. Which solution will meet these requirements with the LEAST development time?A. Use Amazon Kinesis Data Firehose to stream the reservation data directly to Amazon S3. Detect high-demand outliers by using Amazon QuickSight ML Insights. Visualize the data in QuickSight.B. Use Amazon Kinesis Data Streams to stream the

reservation data directly to Amazon S3. Detect high-demand outliers by using the Random Cut Forest (RCF) trained model in Amazon SageMaker. Visualize the data in Amazon QuickSight.C. Use Amazon Kinesis Data Firehose to stream the reservation data directly to Amazon S3. Detect high-demand outliers by using the Random Cut Forest (RCF) trained model in Amazon SageMaker. Visualize the data in Amazon QuickSight.D. Use Amazon Kinesis Data Streams to stream the reservation data directly to Amazon S3. Detect high-demand outliers by using Amazon QuickSight ML Insights. Visualize the data in OuickSight.Answer: AOUESTION 270A company operates an amusement park. The company wants to collect, monitor, and store real-time traffic data at several park entrances by using strategically placed cameras. The company's security team must be able to immediately access the data for viewing. Stored data must be indexed and must be accessible to the company's data science team. Which solution will meet these requirements MOST cost-effectively? A. Use Amazon Kinesis Video Streams to ingest, index, and store the data. Use the built-in integration with Amazon Rekognition for viewing by the security team.B. Use Amazon Kinesis Video Streams to ingest, index, and store the data. Use the built-in HTTP live streaming (HLS) capability for viewing by the security team.C. Use Amazon Rekognition Video and the GStreamer plugin to ingest the data for viewing by the security team. Use Amazon Kinesis Data Streams to index and store the data.D. Use Amazon Kinesis Data Firehose to ingest, index, and store the data. Use the built-in HTTP live streaming (HLS) capability for viewing by the security team. Answer: BQUESTION 271An engraving company wants to automate its quality control process for plaques. The company performs the process before mailing each customized plaque to a customer. The company has created an Amazon S3 bucket that contains images of defects that should cause a plaque to be rejected. Low-confidence predictions must be sent to an internal team of reviewers who are using Amazon Augmented AI (Amazon A2I). Which solution will meet these requirements? A. Use Amazon Textract for automatic processing. Use Amazon A2I with Amazon Mechanical Turk for manual review.B. Use Amazon Rekognition for automatic processing. Use Amazon A2I with a private workforce option for manual review.C. Use Amazon Transcribe for automatic processing. Use Amazon A2I with a private workforce option for manual review.D. Use AWS Panorama for automatic processing. Use Amazon A2I with Amazon Mechanical Turk for manual review. Answer: BQUESTION 272A machine learning (ML) engineer at a bank is building a data ingestion solution to provide transaction features to financial ML models. Raw transactional data is available in an Amazon Kinesis data stream. The solution must compute rolling averages of the ingested data from the data stream and must store the results in Amazon SageMaker Feature Store. The solution also must serve the results to the models in near real time. Which solution will meet these requirements? A. Load the data into an Amazon S3 bucket by using Amazon Kinesis Data Firehose. Use a SageMaker Processing job to aggregate the data and to load the results into SageMaker Feature Store as an online feature group.B. Write the data directly from the data stream into SageMaker Feature Store as an online feature group. Calculate the rolling averages in place within SageMaker Feature Store by using the SageMaker GetRecord API operation.C. Consume the data stream by using an Amazon Kinesis Data Analytics SQL application that calculates the rolling averages. Generate a result stream. Consume the result stream by using a custom AWS Lambda function that publishes the results to SageMaker Feature Store as an online feature group.D. Load the data into an Amazon S3 bucket by using Amazon Kinesis Data Firehose. Use a SageMaker Processing job to load the data into SageMaker Feature Store as an offline feature group. Compute the rolling averages at query time. Answer: CQUESTION 273Each morning, a data scientist at a rental car company creates insights about the previous day's rental car reservation demands. The company needs to automate this process by streaming the data to Amazon S3 in near real time. The solution must detect high-demand rental cars at each of the company's locations. The solution also must create a visualization dashboard that automatically refreshes with the most recent data. Which solution will meet these requirements with the LEAST development time?A. Use Amazon Kinesis Data Firehose to stream the reservation data directly to Amazon S3. Detect high-demand outliers by using Amazon QuickSight ML Insights. Visualize the data in QuickSight.B. Use Amazon Kinesis Data Streams to stream the reservation data directly to Amazon S3. Detect high-demand outliers by using the Random Cut Forest (RCF) trained model in Amazon SageMaker. Visualize the data in Amazon QuickSight.C. Use Amazon Kinesis Data Firehose to stream the reservation data directly to Amazon S3. Detect high-demand outliers by using the Random Cut Forest (RCF) trained model in Amazon SageMaker. Visualize the data in Amazon QuickSight.D. Use Amazon Kinesis Data Streams to stream the reservation data directly to Amazon S3. Detect high-demand outliers by using Amazon QuickSight ML Insights. Visualize the data in QuickSight.Answer: AQUESTION 274A machine learning (ML) engineer is integrating a production model with a customer metadata repository for real-time inference. The repository is hosted in Amazon SageMaker Feature Store. The engineer wants to retrieve only the latest version of the customer metadata record for a single customer at a time. Which solution will meet these requirements?A. Use the SageMaker Feature Store BatchGetRecord API with the record identifier. Filter to find the latest record. B. Create an Amazon Athena query to retrieve the data from the feature table.C. Create an Amazon Athena query to retrieve the data from the feature table. Use the write_time value to find the latest record.D. Use the SageMaker Feature Store GetRecord API

with the record identifier. Answer: DExplanation: GetRecord API retrieves latest record whereas BatchGetRecrdbatch of records. https://docs.aws.amazon.com/sagemaker/latest/APIReference/API_feature_store_GetRecord.html QUESTION 275A company's data scientist has trained a new machine learning model that performs better on test data than the company's existing model performs in the production environment. The data scientist wants to replace the existing model that runs on an Amazon SageMaker endpoint in the production environment. However, the company is concerned that the new model might not work well on the production environment data. The data scientist needs to perform A/B testing in the production environment to evaluate whether the new model performs well on production environment data. Which combination of steps must the data scientist take to perform the A/B testing? (Choose two.)A. Create a new endpoint configuration that includes a production variant for each of the two models.B. Create a new endpoint configuration that includes two target variants that point to different endpoints.C. Deploy the new model to the existing endpoint.D. Update the existing endpoint to activate the new model.E. Update the existing endpoint to use the new endpoint configuration. Answer: ACExplanation:

https://aws.amazon.com/blogs/machine-learning/a-b-testing-ml-models-in-production-using-amazon-sagemaker/QUESTION 276A data scientist is working on a forecast problem by using a dataset that consists of .csv files that are stored in Amazon S3. The files contain a timestamp variable in the following format:March 1st, 2020, 08:14pmThere is a hypothesis about seasonal differences in the dependent variable. This number could be higher or lower for weekdays because some days and hours present varying values, so the day of the week, month, or hour could be an important factor. As a result, the data scientist needs to transform the timestamp into weekdays, month, and day as three separate variables to conduct an analysis. Which solution requires the LEAST operational overhead to create a new dataset with the added features?A. Create an Amazon EMR cluster. Develop PySpark code that can read the timestamp variable as a string, transform and create the new variables, and save the dataset as a new file in Amazon S3.B. Create a processing job in Amazon SageMaker. Develop Python code that can read the timestamp variable as a string, transform and create the new variables, and save the dataset as a new file in Amazon S3.C. Create a new flow in Amazon SageMaker Data Wrangler. Import the S3 file, use the Featurize date/time transform to generate the new variables, and save the dataset as a new file in Amazon S3.D. Create an AWS Glue job. Develop code that can read the timestamp variable as a string, transform and create the new variables, and save the dataset as a new file in Amazon S3.Answer: CExplanation:

https://aws.amazon.com/blogs/machine-learning/prepare-time-series-data-with-amazon-sagemaker-data-wrangler/"Featurize datetime time series transformation to add the month, day of the month, day of the year, week of the year, and quarter features to our dataset." OUESTION 277A manufacturing company has a production line with sensors that collect hundreds of quality metrics. The company has stored sensor data and manual inspection results in a data lake for several months. To automate quality control, the machine learning team must build an automated mechanism that determines whether the produced goods are good quality, replacement market quality, or scrap quality based on the manual inspection results. Which modeling approach will deliver the MOST accurate prediction of product quality? A. Amazon SageMaker DeepAR forecasting algorithm B. Amazon SageMaker XGBoost algorithmC. Amazon SageMaker Latent Dirichlet Allocation (LDA) algorithmD. A convolutional neural network (CNN) and ResNetAnswer: BQUESTION 278A healthcare company wants to create a machine learning (ML) model to predict patient outcomes. A data science team developed an ML model by using a custom ML library. The company wants to use Amazon SageMaker to train this model. The data science team creates a custom SageMaker image to train the model. When the team tries to launch the custom image in SageMaker Studio, the data scientists encounter an error within the application. Which service can the data scientists use to access the logs for this error?A. Amazon S3B. Amazon Elastic Block Store (Amazon EBS)C. AWS CloudTrailD. Amazon CloudWatchAnswer: DExplanation: Amazon CloudWatch is a monitoring and logging service provided by AWS. It collects and stores log files and metrics from various AWS services, including Amazon SageMaker. CloudWatch allows you to gain visibility into your applications and infrastructure by providing a unified view of logs, metrics, and events.QUESTION 279A data scientist wants to build a financial trading bot to automate investment decisions. The financial bot should recommend the quantity and price of an asset to buy or sell to maximize long-term profit. The data scientist will continuously stream financial transactions to the bot for training purposes. The data scientist must select the appropriate machine learning (ML) algorithm to develop the financial trading bot. Which type of ML algorithm will meet these requirements? A. Supervised learning B. Unsupervised learningC. Semi-supervised learningD. Reinforcement learningAnswer: DExplanation:Reinforcement learning allows the bot to continuously learn from its own experiences, adapt to changing market conditions, and optimize its decision-making process over time. It is well-suited for dynamic and uncertain environments like financial markets, where the optimal trading strategies may vary depending on various factors and trends.QUESTION 280A manufacturing company wants to create a machine learning (ML) model to predict when equipment is likely to fail. A data science team already constructed a deep learning model by using TensorFlow and a custom Python script in a local environment. The company wants to use Amazon

SageMaker to train the model. Which TensorFlow estimator configuration will train the model MOST cost-effectively? A. Turn on SageMaker Training Compiler by adding compiler_config=TrainingCompilerConfig() as a parameter. Pass the script to the estimator in the call to the TensorFlow fit() method.B. Turn on SageMaker Training Compiler by adding compiler_config=TrainingCompilerConfig() as a parameter. Turn on managed spot training by setting the use_spot_instances parameter to True. Pass the script to the estimator in the call to the TensorFlow fit() method.C. Adjust the training script to use distributed data parallelism. Specify appropriate values for the distribution parameter. Pass the script to the estimator in the call to the TensorFlow fit() method.D. Turn on SageMaker Training Compiler by adding compiler config=TrainingCompilerConfig() as a parameter. Set the MaxWaitTimeInSeconds parameter to be equal to the MaxRuntimeInSeconds parameter. Pass the script to the estimator in the call to the TensorFlow fit() method. Answer: BQUESTION 281An automotive company uses computer vision in its autonomous cars. The company trained its object detection models successfully by using transfer learning from a convolutional neural network (CNN). The company trained the models by using PyTorch through the Amazon SageMaker SDK. The vehicles have limited hardware and compute power. The company wants to optimize the model to reduce memory, battery, and hardware consumption without a significant sacrifice in accuracy. Which solution will improve the computational efficiency of the models? A. Use Amazon CloudWatch metrics to gain visibility into the SageMaker training weights, gradients, biases, and activation outputs. Compute the filter ranks based on the training information. Apply pruning to remove the low-ranking filters. Set new weights based on the pruned set of filters. Run a new training job with the pruned model.B. Use Amazon SageMaker Ground Truth to build and run data labeling workflows. Collect a larger labeled dataset with the labelling workflows. Run a new training job that uses the new labeled data with previous training data.C. Use Amazon SageMaker Debugger to gain visibility into the training weights, gradients, biases, and activation outputs. Compute the filter ranks based on the training information. Apply pruning to remove the low-ranking filters. Set the new weights based on the pruned set of filters. Run a new training job with the pruned model.D. Use Amazon SageMaker Model Monitor to gain visibility into the ModelLatency metric and OverheadLatency metric of the model after the company deploys the model. Increase the model learning rate. Run a new training job. Answer: CQUESTION 282A data scientist wants to improve the fit of a machine learning (ML) model that predicts house prices. The data scientist makes a first attempt to fit the model, but the fitted model has poor accuracy on both the training dataset and the test dataset. Which steps must the data scientist take to improve model accuracy? (Choose three.)A. Increase the amount of regularization that the model uses.B. Decrease the amount of regularization that the model uses.C. Increase the number of training examples that that model uses.D. Increase the number of test examples that the model uses.E. Increase the number of model features that the model uses.F. Decrease the number of model features that the model uses. Answer: ACEQUESTION 283A car company is developing a machine learning solution to detect whether a car is present in an image. The image dataset consists of one million images. Each image in the dataset is 200 pixels in height by 200 pixels in width. Each image is labeled as either having a car or not having a car. Which architecture is MOST likely to produce a model that detects whether a car is present in an image with the highest accuracy?A. Use a deep convolutional neural network (CNN) classifier with the images as input. Include a linear output layer that outputs the probability that an image contains a car.B. Use a deep convolutional neural network (CNN) classifier with the images as input. Include a softmax output layer that outputs the probability that an image contains a car.C. Use a deep multilayer perceptron (MLP) classifier with the images as input. Include a linear output layer that outputs the probability that an image contains a car.D. Use a deep multilayer perceptron (MLP) classifier with the images as input. Include a softmax output layer that outputs the probability that an image contains a car. Answer: BExplanation: Both MLP and CNN can process images, but CNN is more accurate and can be used for more complex images.QUESTION 284A company is creating an application to identify, count, and classify animal images that are uploaded to the company's website. The company is using the Amazon SageMaker image classification algorithm with an ImageNetV2 convolutional neural network (CNN). The solution works well for most animal images but does not recognize many animal species that are less common. The company obtains 10,000 labeled images of less common animal species and stores the images in Amazon S3. A machine learning (ML) engineer needs to incorporate the images into the model by using Pipe mode in SageMaker. Which combination of steps should the ML engineer take to train the model? (Choose two.) A. Use a ResNet model. Initiate full training mode by initializing the network with random weights.B. Use an Inception model that is available with the SageMaker image classification algorithm.C. Create a .lst file that contains a list of image files and corresponding class labels. Upload the .lst file to Amazon S3.D. Initiate transfer learning. Train the model by using the images of less common species.E. Use an augmented manifest file in JSON Lines format. Answer: DEQUESTION 285A music streaming company is building a pipeline to extract features. The company wants to store the features for offline model training and online inference. The company wants to track feature history and to give the company's data science teams access to the features. Which solution will meet these requirements with the MOST operational efficiency?A. Use Amazon SageMaker Feature Store to store features for model training

and inference. Create an online store for online inference. Create an offline store for model training. Create an IAM role for data scientists to access and search through feature groups.B. Use Amazon SageMaker Feature Store to store features for model training and inference. Create an online store for both online inference and model training. Create an IAM role for data scientists to access and search through feature groups.C. Create one Amazon S3 bucket to store online inference features. Create a second S3 bucket to store offline model training features. Turn on versioning for the S3 buckets and use tags to specify which tags are for online inference features and which are for offline model training features. Use Amazon Athena to query the S3 bucket for online inference. Connect the S3 bucket for offline model training to a SageMaker training job. Create an IAM policy that allows data scientists to access both buckets.D. Create two separate Amazon DynamoDB tables to store online inference features and offline model training features. Use time-based versioning on both tables. Query the DynamoDB table for online inference. Move the data from DynamoDB to Amazon S3 when a new SageMaker training job is launched. Create an IAM policy that allows data scientists to access both tables. Answer: AExplanation:SageMaker Feature Store consists of an online and an offline mode for managing features. The online store is used for low-latency real-time inference use cases. The offline store is primarily used for batch predictions and model training.

https://aws.amazon.com/blogs/machine-learning/speed-ml-development-using-sagemaker-feature-store-and-apache-iceberg-offline-store-compaction/QUESTION 286A 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 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. B. 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. C. 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. D.

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. Answer: AQUESTION 287A financial services company wants to automate its loan approval process by building a machine learning (ML) model. Each loan data point contains credit history from a third-party data source and demographic information about the customer. Each loan approval prediction must come with a report that contains an explanation for why the customer was approved for a loan or was denied for a loan. The company will use Amazon SageMaker to build the model. Which solution will meet these requirements with the LEAST development effort? A. Use SageMaker Model Debugger to automatically debug the predictions, generate the explanation, and attach the explanation report.B. Use AWS Lambda to provide feature importance and partial dependence plots. Use the plots to generate and attach the explanation report.C. Use SageMaker Clarify to generate the explanation report. Attach the report to the predicted results.D. Use custom Amazon CloudWatch metrics to generate the explanation report. Attach the report to the predicted results. Answer: CQUESTION 288A financial company sends special offers to customers through weekly email campaigns. A bulk email marketing system takes the list of email addresses as an input and sends the marketing campaign messages in batches. Few customers use the offers from the campaign messages. The company does not want to send irrelevant offers to customers. A machine learning (ML) team at the company is using Amazon SageMaker to build a model to recommend specific offers to each customer based on the customer's profile and the offers that the customer has accepted in the past. Which solution will meet these requirements with the MOST operational efficiency? A. Use the Factorization Machines algorithm to build a model that can generate personalized offer recommendations for customers. Deploy a SageMaker endpoint to generate offer recommendations. Feed the offer recommendations into the bulk email marketing system.B. Use the Neural Collaborative Filtering algorithm to build a model that can generate personalized offer recommendations for customers. Deploy a SageMaker endpoint to generate offer recommendations. Feed the offer recommendations into the bulk email marketing system.C. Use the Neural Collaborative Filtering algorithm to build a model that can generate personalized offer recommendations for customers. Deploy a SageMaker batch inference job to generate offer recommendations. Feed the offer recommendations into the bulk email marketing system.D. Use the Factorization Machines algorithm to build a model that can generate personalized offer recommendations for customers. Deploy a SageMaker batch inference job to generate offer recommendations. Feed the offer recommendations into the bulk email marketing system. Answer: CQUESTION 289A social media company wants to develop a machine learning (ML) model to detect inappropriate or offensive content in images. The company has collected a large dataset of labeled images and plans to use the built-in Amazon SageMaker image classification algorithm to train the model. The company also intends to use SageMaker pipe mode to speed up the training. The company splits the dataset into training, validation, and testing datasets. The company stores the training and validation images in folders that are named Training and Validation, respectively. The folders contain subfolders that correspond to the names of the dataset classes. The company resizes

Braindump2go, We Devoted to Helping You 100% Pass All Exams!

the images to the same size and generates two input manifest files named training.lst and validation.lst, for the training dataset and the validation dataset, respectively. Finally, the company creates two separate Amazon S3 buckets for uploads of the training dataset and the validation dataset. Which additional data preparation steps should the company take before uploading the files to Amazon S3?A. Generate two Apache Parquet files, training parquet and validation parquet, by reading the images into a Pandas data frame and storing the data frame as a Parquet file. Upload the Parquet files to the training S3 bucket.B. Compress the training and validation directories by using the Snappy compression library. Upload the manifest and compressed files to the training S3 bucket. Compress the training and validation directories by using the gzip compression library. Upload the manifest and compressed files to the training S3 bucket.D. Generate two RecordIO files, training.rec and validation.rec, from the manifest files by using the im2rec Apache MXNet utility tool. Upload the RecordIO files to the training S3 bucket. Answer: CQUESTION 290A media company wants to create a solution that identifies celebrities in pictures that users upload. The company also wants to identify the IP address and the timestamp details from the users so the company can prevent users from uploading pictures from unauthorized locations. Which solution will meet these requirements with LEAST development effort? A. Use AWS Panorama to identify celebrities in the pictures. Use AWS CloudTrail to capture IP address and timestamp details.B. Use AWS Panorama to identify celebrities in the pictures. Make calls to the AWS Panorama Device SDK to capture IP address and timestamp details.C. Use Amazon Rekognition to identify celebrities in the pictures. Use AWS CloudTrail to capture IP address and timestamp details.D. Use Amazon Rekognition to identify celebrities in the pictures. Use the text detection feature to capture IP address and timestamp details. Answer: CResources From: 1.2025 Latest Braindump2go MLS-C01 Exam Dumps (PDF & VCE) Free Share: https://www.braindump2go.com/mls-c01.html2.2025 Latest Braindump2go MLS-C01 PDF and MLS-C01 VCE Dumps Free Share:https://drive.google.com/drive/folders/1eX--L9LzE21hzqPIkigeo1QoAGNWL4vd?usp=sharing3.2025 Free Braindump2go MLS-C01 Exam Questions Download: https://www.braindump2go.com/free-online-pdf/MLS-C01-VCE-Dumps(259-290).pdfFree Resources from