



Free Amazon MLS-C01 Practice Questions

Shared by Spence on 12-08-2026

For More Free Questions and Preparation Resources

Check the Links on Last Page



Question 1

Question Type: MultipleChoice

[Modeling]

An aircraft engine manufacturing company is measuring 200 performance metrics in a time-series. Engineers

want to detect critical manufacturing defects in near-real time during testing. All of the data needs to be stored

for offline analysis.

What approach would be the MOST effective to perform near-real time defect detection?

Options:

A- Use AWS IoT Analytics for ingestion, storage, and further analysis. Use Jupyter notebooks from within AWS IoT Analytics to carry out analysis for anomalies.

B- Use Amazon S3 for ingestion, storage, and further analysis. Use an Amazon EMR cluster to carry out Apache Spark ML k-means clustering to determine anomalies.

C- Use Amazon S3 for ingestion, storage, and further analysis. Use the Amazon SageMaker Random Cut Forest (RCF) algorithm to determine anomalies.

D- Use Amazon Kinesis Data Firehose for ingestion and Amazon Kinesis Data Analytics Random Cut Forest (RCF) to perform anomaly detection. Use Kinesis Data Firehose to store data in Amazon S3 for further analysis.

Answer:

D

Question 2

Question Type: MultipleChoice

[Data Engineering]

A company stores its documents in Amazon S3 with no predefined product categories. A data scientist needs to build a machine learning model to categorize the documents for all the company's products.

Which solution will meet these requirements with the MOST operational efficiency?

Options:

- A- Build a custom clustering model. Create a Dockerfile and build a Docker image. Register the Docker image in Amazon Elastic Container Registry (Amazon ECR). Use the custom image in Amazon SageMaker to generate a trained model.
- B- Tokenize the data and transform the data into tabular data. Train an Amazon SageMaker k-means model to generate the product categories.
- C- Train an Amazon SageMaker Neural Topic Model (NTM) model to generate the product categories.
- D- Train an Amazon SageMaker Blazing Text model to generate the product categories.

Answer:

C



Question 3

Question Type: MultipleChoice

[Modeling]

A Machine Learning Specialist is assigned a TensorFlow project using Amazon SageMaker for training, and needs to continue working for an extended period with no Wi-Fi access.

Which approach should the Specialist use to continue working?

Options:

- A- Install Python 3 and boto3 on their laptop and continue the code development using that environment.
- B- Download the TensorFlow Docker container used in Amazon SageMaker from GitHub to their local environment, and use the Amazon SageMaker Python SDK to test the code.
- C- Download TensorFlow from tensorflow.org to emulate the TensorFlow kernel in the SageMaker environment.
- D- Download the SageMaker notebook to their local environment then install Jupyter Notebooks on their laptop and continue the development in a local notebook.

Answer:

B

Question 4

Question Type: MultipleChoice

[Modeling]

A company has set up and deployed its machine learning (ML) model into production with an endpoint using Amazon SageMaker hosting services. The ML team has configured automatic scaling for its SageMaker instances to support workload changes. During testing, the team notices that additional instances are being launched before the new instances are ready. This behavior needs to change as soon as possible.

How can the ML team solve this issue?

Options:

- A- Decrease the cooldown period for the scale-in activity. Increase the configured maximum capacity of instances.
- B- Replace the current endpoint with a multi-model endpoint using SageMaker.
- C- Set up Amazon API Gateway and AWS Lambda to trigger the SageMaker inference endpoint.
- D- Increase the cooldown period for the scale-out activity.

Answer:

D

Question 5

Question Type: MultipleChoice

[Modeling]

A retail company intends to use machine learning to categorize new products. A labeled dataset of current products was provided to the Data Science team. The dataset includes 1,200 products. The labeled dataset has 15 features for each product such as title, dimensions, weight, and price. Each product is labeled as belonging to one of six categories such as books, games, electronics, and movies.

Which model should be used for categorizing new products using the provided dataset for training?

Options:

- A- An XGBoost model where the objective parameter is set to multi: softmax
- B- A deep convolutional neural network (CNN) with a softmax activation function for the last layer
- C- A regression forest where the number of trees is set equal to the number of product

categories

D- A DeepAR forecasting model based on a recurrent neural network (RNN)

Answer:

A

Question 6

Question Type: MultipleChoice

[Machine Learning Implementation and Operations]

A company wants to create an artificial intelligence (AI) yoga instructor that can lead large classes of students. The company needs to create a feature that can accurately count the number of students who are in a class. The company also needs a feature that can differentiate students who are performing a yoga stretch correctly from students who are performing a stretch incorrectly.

...etermine whether students are performing a stretch correctly, the solution needs to measure the location and angle of each student's arms and legs A data scientist must use Amazon SageMaker to ...ss video footage of a yoga class by extracting image frames and applying computer vision models.

Which combination of models will meet these requirements with the LEAST effort? (Select TWO.)

Options:

- A- Image Classification
- B- Optical Character Recognition (OCR)
- C- Object Detection
- D- Pose estimation
- E- Image Generative Adversarial Networks (GANs)

Answer:

C, D

Question 7

Question Type: MultipleChoice

[Data Engineering]

A Machine Learning Specialist must build out a process to query a dataset on Amazon S3 using Amazon Athena. The dataset contains more than 800,000 records stored as plaintext CSV files. Each record contains 200 columns and is approximately 1.5 MB in size. Most queries will span 5 to 10 columns only.

How should the Machine Learning Specialist transform the dataset to minimize query runtime?

Options:

- A- Convert the records to Apache Parquet format
- B- Convert the records to JSON format
- C- Convert the records to GZIP CSV format
- D- Convert the records to XML format

Answer:

A

Question 8

Question Type: MultipleChoice

[Data Engineering]

A Data Scientist is building a linear regression model and will use resulting p-values to evaluate the statistical significance of each coefficient. Upon inspection of the dataset, the Data Scientist discovers that most of the features are normally distributed. The plot of one feature in the dataset is shown in the graphic.

What transformation should the Data Scientist apply to satisfy the statistical assumptions of the linear regression model?

Options:

- A- Exponential transformation
- B- Logarithmic transformation
- C- Polynomial transformation
- D- Sinusoidal transformation

Answer:

B

Question 9

Question Type: MultipleChoice

[Machine Learning Implementation and Operations]

A financial services company wants to adopt Amazon SageMaker as its default data science environment. The company's data scientists run machine learning (ML) models on confidential financial data.

a. The company is worried about data egress and wants an ML engineer to secure the environment.

Which mechanisms can the ML engineer use to control data egress from SageMaker? (Select three.)

Options:

- A- Connect to SageMaker by using a VPC interface endpoint powered by AWS PrivateLink.
- B- Use SCPs to restrict access to SageMaker.
- C- Disable root access on the SageMaker notebook instances.
- D- Enable network isolation for training jobs and models.
- E- Restrict notebook presigned URLs to specific IPs used by the company.
- F- Protect data with encryption at rest and in transit. Use AWS Key Management Service (AWS KMS) to manage encryption keys.

Answer:

A, D, F

To Get Premium Files for MLS-C01 Visit

<https://www.p2pexams.com/products/mls-c01>

For More Free Questions Visit

<https://www.p2pexams.com/amazon/pdf/mls-c01>

20%
DISCOUNT

P2P
exams