



MLS-C01^{Q&As}

AWS Certified Machine Learning - Specialty (MLS-C01)

Pass Amazon MLS-C01 Exam with 100% Guarantee

Free Download Real Questions & Answers **PDF** and **VCE** file from:

<https://www.pass4itsure.com/aws-certified-machine-learning-specialty.html>

100% Passing Guarantee
100% Money Back Assurance

Following Questions and Answers are all new published by Amazon
Official Exam Center

- ⚙️ **Instant Download** After Purchase
- ⚙️ **100% Money Back** Guarantee
- ⚙️ **365 Days** Free Update
- ⚙️ **800,000+** Satisfied Customers





QUESTION 1

A financial company is trying to detect credit card fraud. The company observed that, on average, 2% of credit card transactions were fraudulent. A data scientist trained a classifier on a year's worth of credit card transactions data. The model needs to identify the fraudulent transactions (positives) from the regular ones (negatives). The company's goal is to accurately capture as many positives as possible.

Which metrics should the data scientist use to optimize the model? (Choose two.)

- A. Specificity
- B. False positive rate
- C. Accuracy
- D. Area under the precision-recall curve
- E. True positive rate

Correct Answer: DE

The goal is to accurately capture as many fraudulent transactions (positives) as possible. To optimize the model towards this goal, the data scientist should focus on metrics that emphasize the true positive rate and the area under the precision-recall curve.

True positive rate (TPR or sensitivity) is the proportion of actual positive cases that are correctly identified as positive by the model. A higher TPR means that more fraudulent transactions are being captured.


The precision-recall curve is a graph that shows the trade-off between precision and recall for different thresholds

QUESTION 2

This graph shows the training and validation loss against the epochs for a neural network.

The network being trained is as follows:

1.
Two dense layers, one output neuron
2.
100 neurons in each layer
3.
100 epochs
4.
Random initialization of weights



Which technique can be used to improve model performance in terms of accuracy in the validation set?

- A. Early stopping
- B. Random initialization of weights with appropriate seed
- C. Increasing the number of epochs
- D. Adding another layer with the 100 neurons

Correct Answer: A

The answer is Early Stopping. Stop the training before accuracy start do decrease.

QUESTION 3

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?

- 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.

Correct Answer: D

Reference: <https://aws.amazon.com/blogs/machine-learning/configuring-autoscaling-inference-endpoints-in-amazon-sagemaker/>

QUESTION 4

A company wants to use machine learning (ML) to improve its customer churn prediction model. The company stores data in an Amazon Redshift data warehouse.

A data science team wants to use Amazon Redshift machine learning (Amazon Redshift ML) to build a model and run predictions for new data directly within the data warehouse.

Which combination of steps should the company take to use Amazon Redshift ML to meet these requirements? (Choose three.)

- A. Define the feature variables and target variable for the churn prediction model.
- B. Use the SQL EXPLAIN_MODEL function to run predictions.
- C. Write a CREATE MODEL SQL statement to create a model.
- D. Use Amazon Redshift Spectrum to train the model.
- E. Manually export the training data to Amazon S3.
- F. Use the SQL prediction function to run predictions.

Correct Answer: ACF

QUESTION 5

An insurance company is creating an application to automate car insurance claims. A machine learning (ML) specialist used an Amazon SageMaker Object Detection - TensorFlow built-in algorithm to train a model to detect scratches and dents in images of cars. After the model was trained, the ML specialist noticed that the model performed better on the training dataset than on the testing dataset.

Which approach should the ML specialist use to improve the performance of the model on the testing data?

- A. Increase the value of the momentum hyperparameter.
- B. Reduce the value of the dropout_rate hyperparameter.
- C. Reduce the value of the learning_rate hyperparameter
- D. Increase the value of the L2 hyperparameter.

Correct Answer: D



Feature selection: consider using fewer feature combinations, decrease n-grams size, and decrease the number of numeric attribute bins. Increase the amount of regularization used. <https://docs.aws.amazon.com/machine-learning/latest/dg/model-fit-underfitting-vs-overfitting.html>

QUESTION 6

A Machine Learning Specialist is building a logistic regression model that will predict whether or not a person will order a pizza. The Specialist is trying to build the optimal model with an ideal classification threshold.

What model evaluation technique should the Specialist use to understand how different classification thresholds will impact the model's performance?

- A. Receiver operating characteristic (ROC) curve
- B. Misclassification rate
- C. Root Mean Square Error (RMand)
- D. L1 norm

Correct Answer: A

Reference: <https://docs.aws.amazon.com/machine-learning/latest/dg/binary-model-insights.html>

QUESTION 7

A sports analytics company is providing services at a marathon. Each runner in the marathon will have their race ID printed as text on the front of their shirt. The company needs to extract race IDs from images of the runners. Which solution will meet these requirements with the LEAST operational overhead?

- A. Use Amazon Rekognition.
- B. Use a custom convolutional neural network (CNN).
- C. Use the Amazon SageMaker Object Detection algorithm.
- D. Use Amazon Lookout for Vision.

Correct Answer: A

<https://docs.aws.amazon.com/rekognition/latest/dg/text-detection.html?pg=inandsec=ft>

QUESTION 8

A media company is building a computer vision model to analyze images that are on social media. The model consists of CNNs that the company trained by using images that the company stores in Amazon S3. The company used an Amazon SageMaker training job in File mode with a single Amazon EC2 On-Demand Instance.

Every day, the company updates the model by using about 10,000 images that the company has collected in the last 24 hours. The company configures training with only one epoch. The company wants to speed up training and lower costs without the need to make any code changes.



Which solution will meet these requirements?

- A. Instead of File mode, configure the SageMaker training job to use Pipe mode. Ingest the data from a pipe.
- B. Instead of File mode, configure the SageMaker training job to use FastFile mode with no other changes.
- C. Instead of On-Demand Instances, configure the SageMaker training job to use Spot Instances. Make no other changes,
- D. Instead of On-Demand Instances, configure the SageMaker training job to use Spot Instances, implement model checkpoints.

Correct Answer: A

QUESTION 9

A data scientist obtains a tabular dataset that contains 150 correlated features with different ranges to build a regression model. The data scientist needs to achieve more efficient model training by implementing a solution that minimizes impact on the model's performance. The data scientist decides to perform a principal component analysis (PCA) preprocessing step to reduce the number of features to a smaller set of independent features before the data scientist uses the new features in the regression model.

Which preprocessing step will meet these requirements?

- A. Use the Amazon SageMaker built-in algorithm for PCA on the dataset to transform the data
- B. Load the data into Amazon SageMaker Data Wrangler. Scale the data with a Min Max Scaler transformation step Use the SageMaker built-in algorithm for PCA on the scaled dataset to transform the data.
- C. Reduce the dimensionality of the dataset by removing the features that have the highest correlation Load the data into Amazon SageMaker Data Wrangler Perform a Standard Scaler transformation step to scale the data Use the SageMaker built-in algorithm for PCA on the scaled dataset to transform the data
- D. Reduce the dimensionality of the dataset by removing the features that have the lowest correlation. Load the data into Amazon SageMaker Data Wrangler. Perform a Min Max Scaler transformation step to scale the data. Use the SageMaker built-in algorithm for PCA on the scaled dataset to transform the data.

Correct Answer: B

Principal component analysis (PCA) is a technique for reducing the dimensionality of datasets, increasing interpretability but at the same time minimizing information loss. It does so by creating new uncorrelated variables that successively maximize variance. PCA is useful when dealing with datasets that have a large number of correlated features. However, PCA is sensitive to the scale of the features, so it is important to standardize or normalize the data before applying PCA. Amazon SageMaker provides a built-in algorithm for PCA that can be used to transform the data into a lower-dimensional representation. Amazon SageMaker Data Wrangler is a tool that allows data scientists to visually explore, clean, and prepare data for machine learning. Data Wrangler provides various transformation steps that can be applied to the data, such as scaling, encoding, imputing, etc. Data Wrangler also integrates with SageMaker built-in algorithms, such as PCA, to enable feature engineering and dimensionality reduction. Therefore, option B is the correct answer, as it involves scaling the data with a Min Max Scaler transformation step, which rescales the data to a range of [0, 1], and then using the SageMaker built-in algorithm for PCA on the scaled dataset to transform the data. Option A is incorrect, as it does not involve scaling the data before applying PCA, which can affect the results of the dimensionality reduction. Option C is incorrect, as it involves removing the features that have the highest correlation, which can lead to information loss and reduce the performance of the regression model. Option D is incorrect, as it involves removing the features that have the lowest correlation, which can also lead to information loss and reduce the performance of the regression model. References: Principal Component Analysis (PCA) - Amazon SageMaker Scale data with a Min Max



Scaler - Amazon SageMaker Data Wrangler Use Amazon SageMaker built-in algorithms - Amazon SageMaker Data Wrangler

QUESTION 10

A Machine Learning Specialist is given a structured dataset on the shopping habits of a company's customer base. The dataset contains thousands of columns of data and hundreds of numerical columns for each customer. The Specialist wants to identify whether there are natural groupings for these columns across all customers and visualize the results as quickly as possible.

What approach should the Specialist take to accomplish these tasks?

- A. Embed the numerical features using the t-distributed stochastic neighbor embedding (t-SNE) algorithm and create a scatter plot.
- B. Run k-means using the Euclidean distance measure for different values of k and create an elbow plot.
- C. Embed the numerical features using the t-distributed stochastic neighbor embedding (t-SNE) algorithm and create a line graph.
- D. Run k-means using the Euclidean distance measure for different values of k and create box plots for each numerical column within each cluster.

Correct Answer: A

<https://towardsdatascience.com/an-introduction-to-t-sne-with-python-example-5a3a293108d1>

QUESTION 11

A mining company wants to use machine learning (ML) models to identify mineral images in real time. A data science team built an image recognition model that is based on convolutional neural network (CNN). The team trained the model on Amazon SageMaker by using GPU instances. The team will deploy the model to a SageMaker endpoint.

The data science team already knows the workload traffic patterns. The team must determine instance type and configuration for the workloads.

Which solution will meet these requirements with the LEAST development effort?

- A. Register the model artifact and container to the SageMaker Model Registry. Use the SageMaker Inference Recommender Default job type. Provide the known traffic pattern for load testing to select the best instance type and configuration based on the workloads.
- B. Register the model artifact and container to the SageMaker Model Registry. Use the SageMaker Inference Recommender Advanced job type. Provide the known traffic pattern for load testing to select the best instance type and configuration based on the workloads.
- C. Deploy the model to an endpoint by using GPU instances. Use AWS Lambda and Amazon API Gateway to handle invocations from the web. Use open-source tools to perform load testing against the endpoint and to select the best instance type and configuration.
- D. Deploy the model to an endpoint by using CPU instances. Use AWS Lambda and Amazon API Gateway to handle invocations from the web. Use open-source tools to perform load testing against the endpoint and to select the best instance type and configuration.



Correct Answer: B

QUESTION 12

A financial services company is building a robust serverless data lake on Amazon S3. The data lake should be flexible and meet the following requirements:

1.

Support querying old and new data on Amazon S3 through Amazon Athena and Amazon Redshift Spectrum.

2.

Support event-driven ETL pipelines.

3.

Provide a quick and easy way to understand metadata. Which approach meets these requirements?

A. Use an AWS Glue crawler to crawl S3 data, an AWS Lambda function to trigger an AWS Glue ETL job, and an AWS Glue Data catalog to search and discover metadata.

B. Use an AWS Glue crawler to crawl S3 data, an AWS Lambda function to trigger an AWS Batch job, and an external Apache Hive metastore to search and discover metadata.

C. Use an AWS Glue crawler to crawl S3 data, an Amazon CloudWatch alarm to trigger an AWS Batch job, and an AWS Glue Data Catalog to search and discover metadata.

D. Use an AWS Glue crawler to crawl S3 data, an Amazon CloudWatch alarm to trigger an AWS Glue ETL job, and an external Apache Hive metastore to search and discover metadata.

Correct Answer: A

The AWS Glue Data Catalog is your persistent metadata store. It is a managed service that lets you store, annotate, and share metadata in the AWS Cloud in the same way you would in an Apache Hive metastore. The Data Catalog is a drop-in replacement for the Apache Hive Metastore

https://docs.aws.amazon.com/zh_tw/glue/latest/dg/components-overview.html

QUESTION 13

An insurance company developed a new experimental machine learning (ML) model to replace an existing model that is in production. The company must validate the quality of predictions from the new experimental model in a production environment before the company uses the new experimental model to serve general user requests.

Which one model can serve user requests at a time. The company must measure the performance of the new experimental model without affecting the current live traffic

Which solution will meet these requirements?

A. A/B testing

B. Canary release



C. Shadow deployment

D. Blue/green deployment

Correct Answer: C

The best solution for this scenario is to use shadow deployment, which is a technique that allows the company to run the new experimental model in parallel with the existing model, without exposing it to the end users. In shadow deployment, the company can route the same user requests to both models, but only return the responses from the existing model to the users. The responses from the new experimental model are logged and analyzed for quality and performance metrics, such as accuracy, latency, and resource consumption¹². This way, the company can validate the new experimental model in a production environment, without affecting the current live traffic or user experience. The other solutions are not suitable, because they have the following drawbacks:

A: A/B testing is a technique that involves splitting the user traffic between two or more models, and comparing their outcomes based on predefined metrics. However, this technique exposes the new experimental model to a portion of the end

users, which might affect their experience if the model is not reliable or consistent with the existing model³.

B: Canary release is a technique that involves gradually rolling out the new experimental model to a small subset of users, and monitoring its performance and feedback. However, this technique also exposes the new experimental model to

some end users, and requires careful selection and segmentation of the user groups⁴.

D: Blue/green deployment is a technique that involves switching the user traffic from the existing model (blue) to the new experimental model (green) at once, after testing and verifying the new model in a separate environment. However, this

technique does not allow the company to validate the new experimental model in a production environment, and might cause service disruption or inconsistency if the new model is not compatible or stable⁵.

References:

1: Shadow Deployment: A Safe Way to Test in Production | LaunchDarkly Blog

2: Shadow Deployment: A Safe Way to Test in Production | LaunchDarkly Blog

3: A/B Testing for Machine Learning Models | AWS Machine Learning Blog

4: Canary Releases for Machine Learning Models | AWS Machine Learning Blog

5: Blue-Green Deployments for Machine Learning Models | AWS Machine Learning Blog

QUESTION 14

A retail company wants to use Amazon Forecast to predict daily stock levels of inventory. The cost of running out of items in stock is much higher for the company than the cost of having excess inventory. The company has millions of data

samples for multiple years for thousands of items. The company's purchasing department needs to predict demand for 30-day cycles for each item to ensure that restocking occurs.

A machine learning (ML) specialist wants to use item-related features such as "category," "brand," and "safety stock



count." The ML specialist also wants to use a binary time series feature that has "promotion applied?" as its name. Future

promotion information is available only for the next 5 days.

The ML specialist must choose an algorithm and an evaluation metric for a solution to produce prediction results that will maximize company profit. Which solution will meet these requirements?

- A. Train a model by using the Autoregressive Integrated Moving Average (ARIMA) algorithm. Evaluate the model by using the Weighted Quantile Loss (wQL) metric at 0.75 (P75).
- B. Train a model by using the Autoregressive Integrated Moving Average (ARIMA) algorithm. Evaluate the model by using the Weighted Absolute Percentage Error (WAPE) metric.
- C. Train a model by using the Convolutional Neural Network - Quantile Regression (CNN-QR) algorithm. Evaluate the model by using the Weighted Quantile Loss (wQL) metric at 0.75 (P75).
- D. Train a model by using the Convolutional Neural Network - Quantile Regression (CNN-QR) algorithm. Evaluate the model by using the Weighted Absolute Percentage Error (WAPE) metric.

Correct Answer: C

<https://docs.aws.amazon.com/forecast/latest/dg/aws-forecast-algo-cnnqr.html>
<https://docs.aws.amazon.com/forecast/latest/dg/metrics.html#metrics-wQL>

QUESTION 15

A company wants to use automatic speech recognition (ASR) to transcribe messages that are less than 60 seconds long from a voicemail-style application. The company requires the correct identification of 200 unique product names, some of which have unique spellings or pronunciations.

The company has 4,000 words of Amazon SageMaker Ground Truth voicemail transcripts it can use to customize the chosen ASR model. The company needs to ensure that everyone can update their customizations multiple times each hour.

Which approach will maximize transcription accuracy during the development phase?

- A. Use a voice-driven Amazon Lex bot to perform the ASR customization. Create customer slots within the bot that specifically identify each of the required product names. Use the Amazon Lex synonym mechanism to provide additional variations of each product name as mis-transcriptions are identified in development.
- B. Use Amazon Transcribe to perform the ASR customization. Analyze the word confidence scores in the transcript, and automatically create or update a custom vocabulary file with any word that has a confidence score below an acceptable threshold value. Use this updated custom vocabulary file in all future transcription tasks.
- C. Create a custom vocabulary file containing each product name with phonetic pronunciations, and use it with Amazon Transcribe to perform the ASR customization. Analyze the transcripts and manually update the custom vocabulary file to include updated or additional entries for those names that are not being correctly identified.
- D. Use the audio transcripts to create a training dataset and build an Amazon Transcribe custom language model. Analyze the transcripts and update the training dataset with a manually corrected version of transcripts where product names are not being transcribed correctly. Create an updated custom language model.

Correct Answer: C



<https://aws.amazon.com/blogs/machine-learning/build-a-custom-vocabulary-to-enhance-speech-to-text-transcription-accuracy-with-amazon-transcribe/>

[MLS-C01 Practice Test](#)

[MLS-C01 Study Guide](#)

[MLS-C01 Exam Questions](#)