

# Amazon-Web-Services

## Exam Questions MLA-C01

AWS Certified Machine Learning Engineer - Associate



#### NEW QUESTION 1

A company is building a deep learning model on Amazon SageMaker. The company uses a large amount of data as the training dataset. The company needs to optimize the model's hyperparameters to minimize the loss function on the validation dataset.

Which hyperparameter tuning strategy will accomplish this goal with the LEAST computation time?

- A. Hyperbaric!
- B. Grid search
- C. Bayesian optimization
- D. Random search

**Answer: A**

#### NEW QUESTION 2

A company uses Amazon Athena to query a dataset in Amazon S3. The dataset has a target variable that the company wants to predict. The company needs to use the dataset in a solution to determine if a model can predict the target variable.

Which solution will provide this information with the LEAST development effort?

- A. Create a new model by using Amazon SageMaker Autopilo
- B. Report the model's achieved performance.
- C. Implement custom scripts to perform data pre-processing, multiple linear regression, and performance evaluatio
- D. Run the scripts on Amazon EC2 instances.
- E. Configure Amazon Macie to analyze the dataset and to create a mode
- F. Report the model's achieved performance.
- G. Select a model from Amazon Bedroc
- H. Tune the model with the dat
- I. Report the model's achieved performance.

**Answer: A**

#### NEW QUESTION 3

An advertising company uses AWS Lake Formation to manage a data lake. The data lake contains structured data and unstructured data. The company's ML engineers are assigned to specific advertisement campaigns.

The ML engineers must interact with the data through Amazon Athena and by browsing the data directly in an Amazon S3 bucket. The ML engineers must have access to only the resources that are specific to their assigned advertisement campaigns.

Which solution will meet these requirements in the MOST operationally efficient way?

- A. Configure IAM policies on an AWS Glue Data Catalog to restrict access to Athena based on the ML engineers' campaigns.
- B. Store users and campaign information in an Amazon DynamoDB tabl
- C. Configure DynamoDB Streams to invoke an AWS Lambda function to update S3 bucket policies.
- D. Use Lake Formation to authorize AWS Glue to access the S3 bucke
- E. Configure Lake Formation tags to map ML engineers to their campaigns.
- F. Configure S3 bucket policies to restrict access to the S3 bucket based on the ML engineers' campaigns.

**Answer: C**

#### NEW QUESTION 4

A company has trained and deployed an ML model by using Amazon SageMaker. The company needs to implement a solution to record and monitor all the API call events for the SageMaker endpoint. The solution also must provide a notification when the number of API call events breaches a threshold.

Use SageMaker Debugger to track the inferences and to report metrics. Create a custom rule to provide a notification when the threshold is breached.

Which solution will meet these requirements?

- A. Use SageMaker Debugger to track the inferences and to report metric
- B. Create a custom rule to provide a notification when the threshold is breached.
- C. Use SageMaker Debugger to track the inferences and to report metric
- D. Use the tensor\_variance built-in rule to provide a notification when the threshold is breached.
- E. Log all the endpoint invocation API events by using AWS CloudTrai
- F. Use an Amazon CloudWatch dashboard for monitorin
- G. Set up a CloudWatch alarm to provide notification when the threshold is breached.
- H. Add the Invocations metric to an Amazon CloudWatch dashboard for monitorin
- I. Set up a CloudWatch alarm to provide notification when the threshold is breached.

**Answer: D**

#### NEW QUESTION 5

An ML engineer has an Amazon Comprehend custom model in Account A in the us-east-1 Region. The ML engineer needs to copy the model to Account in the same Region.

Which solution will meet this requirement with the LEAST development effort?

- A. Use Amazon S3 to make a copy of the mode
- B. Transfer the copy to Account B.
- C. Create a resource-based IAM polic
- D. Use the Amazon Comprehend ImportModel API operation to copy the model to Account B.
- E. Use AWS DataSync to replicate the model from Account A to Account B.
- F. Create an AWS Site-to-Site VPN connection between Account A and Account to transfer the model.

**Answer: B**

#### NEW QUESTION 6

A company has a conversational AI assistant that sends requests through Amazon Bedrock to an Anthropic Claude large language model (LLM). Users report that when they ask similar questions multiple times, they sometimes receive different answers. An ML engineer needs to improve the responses to be more consistent and less random.

Which solution will meet these requirements?

- A. Increase the temperature parameter and the top\_k parameter.
- B. Increase the temperature parameter
- C. Decrease the top\_k parameter.
- D. Decrease the temperature parameter
- E. Increase the top\_k parameter.
- F. Decrease the temperature parameter and the top\_k parameter.

**Answer: D**

#### NEW QUESTION 7

A company has used Amazon SageMaker to deploy a predictive ML model in production. The company is using SageMaker Model Monitor on the model. After a model update, an ML engineer notices data quality issues in the Model Monitor checks.

What should the ML engineer do to mitigate the data quality issues that Model Monitor has identified?

- A. Adjust the model's parameters and hyperparameters.
- B. Initiate a manual Model Monitor job that uses the most recent production data.
- C. Create a new baseline from the latest dataset
- D. Update Model Monitor to use the new baseline for evaluations.
- E. Include additional data in the existing training set for the model
- F. Retrain and redeploy the model.

**Answer: C**

#### NEW QUESTION 8

A company has a Retrieval Augmented Generation (RAG) application that uses a vector database to store embeddings of documents. The company must migrate the application to AWS and must implement a solution that provides semantic search of text files. The company has already migrated the text repository to an Amazon S3 bucket.

Which solution will meet these requirements?

- A. Use an AWS Batch job to process the files and generate embedding
- B. Use AWS Glue to store the embedding
- C. Use SQL queries to perform the semantic searches.
- D. Use a custom Amazon SageMaker notebook to run a custom script to generate embedding
- E. Use SageMaker Feature Store to store the embedding
- F. Use SQL queries to perform the semantic searches.
- G. Use the Amazon Kendra S3 connector to ingest the documents from the S3 bucket into Amazon Kendra
- H. Query Amazon Kendra to perform the semantic searches.
- I. Use an Amazon Textract asynchronous job to ingest the documents from the S3 bucket
- J. Query Amazon Textract to perform the semantic searches.

**Answer: C**

#### NEW QUESTION 9

A company is using Amazon SageMaker to create ML models. The company's data scientists need fine-grained control of the ML workflows that they orchestrate. The data scientists also need the ability to visualize SageMaker jobs and workflows as a directed acyclic graph (DAG). The data scientists must keep a running history of model discovery experiments and must establish model governance for auditing and compliance verifications.

Which solution will meet these requirements?

- A. Use AWS CodePipeline and its integration with SageMaker Studio to manage the entire ML workflow
- B. Use SageMaker ML Lineage Tracking for the running history of experiments and for auditing and compliance verifications.
- C. Use AWS CodePipeline and its integration with SageMaker Experiments to manage the entire ML workflow
- D. Use SageMaker Experiments for the running history of experiments and for auditing and compliance verifications.
- E. Use SageMaker Pipelines and its integration with SageMaker Studio to manage the entire ML workflow
- F. Use SageMaker ML Lineage Tracking for the running history of experiments and for auditing and compliance verifications.
- G. Use SageMaker Pipelines and its integration with SageMaker Experiments to manage the entire ML workflow
- H. Use SageMaker Experiments for the running history of experiments and for auditing and compliance verifications.

**Answer: C**

#### NEW QUESTION 10

An ML engineer is developing a fraud detection model by using the Amazon SageMaker XGBoost algorithm. The model classifies transactions as either fraudulent or legitimate.

During testing, the model excels at identifying fraud in the training dataset. However, the model is inefficient at identifying fraud in new and unseen transactions.

What should the ML engineer do to improve the fraud detection for new transactions?

- A. Increase the learning rate.
- B. Remove some irrelevant features from the training dataset.
- C. Increase the value of the max\_depth hyperparameter.
- D. Decrease the value of the max\_depth hyperparameter.

**Answer: D**

**NEW QUESTION 10**

Case study

An ML engineer is developing a fraud detection model on AWS. The training dataset includes transaction logs, customer profiles, and tables from an on-premises MySQL database. The transaction logs and customer profiles are stored in Amazon S3.

The dataset has a class imbalance that affects the learning of the model's algorithm. Additionally, many of the features have interdependencies. The algorithm is not capturing all the desired underlying patterns in the data.

Before the ML engineer trains the model, the ML engineer must resolve the issue of the imbalanced data.

Which solution will meet this requirement with the LEAST operational effort?

- A. Use Amazon Athena to identify patterns that contribute to the imbalance
- B. Adjust the dataset accordingly.
- C. Use Amazon SageMaker Studio Classic built-in algorithms to process the imbalanced dataset.
- D. Use AWS Glue DataBrew built-in features to oversample the minority class.
- E. Use the Amazon SageMaker Data Wrangler balance data operation to oversample the minority class.

**Answer: D**

**NEW QUESTION 11**

An ML engineer has trained a neural network by using stochastic gradient descent (SGD). The neural network performs poorly on the test set. The values for training loss and validation loss remain high and show an oscillating pattern. The values decrease for a few epochs and then increase for a few epochs before repeating the same cycle.

What should the ML engineer do to improve the training process?

- A. Introduce early stopping.
- B. Increase the size of the test set.
- C. Increase the learning rate.
- D. Decrease the learning rate.

**Answer: D**

**NEW QUESTION 12**

HOTSPOT

An ML engineer is building a generative AI application on Amazon Bedrock by using large language models (LLMs).

Select the correct generative AI term from the following list for each description. Each term should be selected one time or not at all. (Select three.)

- Embedding
- Retrieval Augmented Generation (RAG)
- Temperature
- Token

Text representation of basic units of data processed by LLMs

Select...

Select...

Embedding

Retrieval Augmented Generation (RAG)

Temperature

Token

High-dimensional vectors that contain the semantic meaning of text

Select...

Select...

Embedding

Retrieval Augmented Generation (RAG)

Temperature

Token

Enrichment of information from additional data sources to improve a generated response

Select...

Select...

Embedding

Retrieval Augmented Generation (RAG)

Temperature

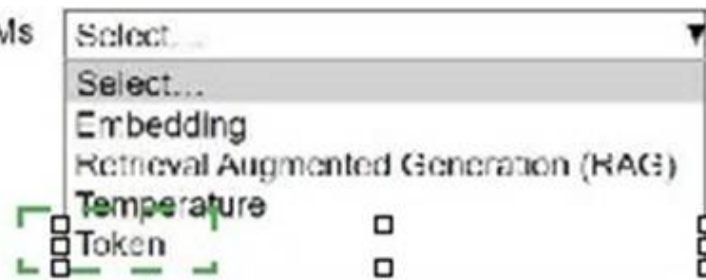
Token

- A. Mastered
- B. Not Mastered

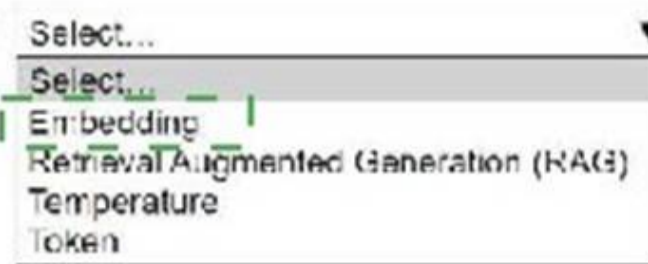
**Answer: A**

**Explanation:**

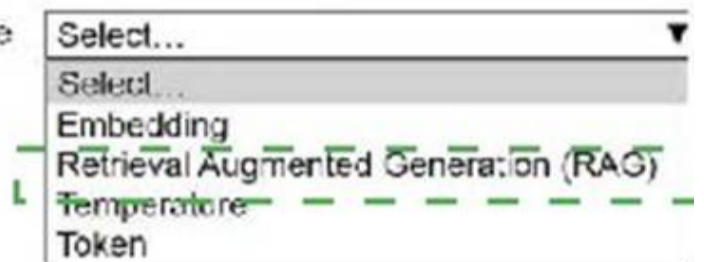
Text representation of basic units of data processed by LLMs



High-dimensional vectors that contain the semantic meaning of text



Enrichment of information from additional data sources to improve a generated response



**NEW QUESTION 17**

An ML engineer normalized training data by using min-max normalization in AWS Glue DataBrew. The ML engineer must normalize the production inference data in the same way as the training data before passing the production inference data to the model for predictions. Which solution will meet this requirement?

- A. Apply statistics from a well-known dataset to normalize the production samples.
- B. Keep the min-max normalization statistics from the training set
- C. Use these values to normalize the production samples.
- D. Calculate a new set of min-max normalization statistics from a batch of production sample
- E. Use these values to normalize all the production samples.
- F. Calculate a new set of min-max normalization statistics from each production sample
- G. Use these values to normalize all the production samples.

**Answer: B**

**NEW QUESTION 22**

An ML engineer needs to use AWS CloudFormation to create an ML model that an Amazon SageMaker endpoint will host. Which resource should the ML engineer declare in the CloudFormation template to meet this requirement?

- A. AWS::SageMaker::Model
- B. AWS::SageMaker::Endpoint
- C. AWS::SageMaker::NotebookInstance
- D. AWS::SageMaker::Pipeline

**Answer: A**

**NEW QUESTION 27**

An ML engineer needs to use an ML model to predict the price of apartments in a specific location. Which metric should the ML engineer use to evaluate the model's performance?

- A. Accuracy
- B. Area Under the ROC Curve (AUC)
- C. F1 score
- D. Mean absolute error (MAE)

**Answer: D**

**NEW QUESTION 29**

A company needs to host a custom ML model to perform forecast analysis. The forecast analysis will occur with predictable and sustained load during the same 2-hour period every day. Multiple invocations during the analysis period will require quick responses. The company needs AWS to manage the underlying infrastructure and any auto scaling activities. Which solution will meet these requirements?

- A. Schedule an Amazon SageMaker batch transform job by using AWS Lambda.
- B. Configure an Auto Scaling group of Amazon EC2 instances to use scheduled scaling.
- C. Use Amazon SageMaker Serverless Inference with provisioned concurrency.
- D. Run the model on an Amazon Elastic Kubernetes Service (Amazon EKS) cluster on Amazon EC2 with pod auto scaling.

**Answer: C**

### NEW QUESTION 32

Case study

An ML engineer is developing a fraud detection model on AWS. The training dataset includes transaction logs, customer profiles, and tables from an on-premises MySQL database. The transaction logs and customer profiles are stored in Amazon S3.

The dataset has a class imbalance that affects the learning of the model's algorithm. Additionally, many of the features have interdependencies. The algorithm is not capturing all the desired underlying patterns in the data.

The training dataset includes categorical data and numerical data. The ML engineer must prepare the training dataset to maximize the accuracy of the model.

Which action will meet this requirement with the LEAST operational overhead?

- A. Use AWS Glue to transform the categorical data into numerical data.
- B. Use AWS Glue to transform the numerical data into categorical data.
- C. Use Amazon SageMaker Data Wrangler to transform the categorical data into numerical data.
- D. Use Amazon SageMaker Data Wrangler to transform the numerical data into categorical data.

**Answer: C**

### NEW QUESTION 34

An ML engineer needs to use an Amazon EMR cluster to process large volumes of data in batches. Any data loss is unacceptable.

Which instance purchasing option will meet these requirements MOST cost-effectively?

- A. Run the primary node, core nodes, and task nodes on On-Demand Instances.
- B. Run the primary node, core nodes, and task nodes on Spot Instances.
- C. Run the primary node on an On-Demand Instance
- D. Run the core nodes and task nodes on Spot Instances.
- E. Run the primary node and core nodes on On-Demand Instance
- F. Run the task nodes on Spot Instances.

**Answer: D**

### NEW QUESTION 36

.....

## **Thank You for Trying Our Product**

### **We offer two products:**

1st - We have Practice Tests Software with Actual Exam Questions

2nd - Questions and Answers in PDF Format

### **MLA-C01 Practice Exam Features:**

- \* MLA-C01 Questions and Answers Updated Frequently
- \* MLA-C01 Practice Questions Verified by Expert Senior Certified Staff
- \* MLA-C01 Most Realistic Questions that Guarantee you a Pass on Your FirstTry
- \* MLA-C01 Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year

**100% Actual & Verified — Instant Download, Please Click**  
**[Order The MLA-C01 Practice Test Here](#)**