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.geekcert.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 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 2

A Machine Learning Specialist is developing recommendation engine for a photography blog Given a picture, the recommendation engine should show a picture that captures similar objects The Specialist would like to create a numerical representation feature to perform nearest-neighbor searches

What actions would allow the Specialist to get relevant numerical representations?

A. Reduce image resolution and use reduced resolution pixel values as features

B. Use Amazon Mechanical Turk to label image content and create a one-hot representation indicating the presence of specific labels

C. Run images through a neural network pie-trained on ImageNet, and collect the feature vectors from the penultimate layer

D. Average colors by channel to obtain three-dimensional representations of images.

Correct Answer: A

QUESTION 3

A large JSON dataset for a project has been uploaded to a private Amazon S3 bucket The Machine Learning Specialist wants to securely access and explore the data from an Amazon SageMaker notebook instance A new VPC was created and assigned to the Specialist

How can the privacy and integrity of the data stored in Amazon S3 be maintained while granting access to the Specialist for analysis?

A. Launch the SageMaker notebook instance within the VPC with SageMaker-provided internet access enabled Use an S3 ACL to open read privileges to the everyone group

B. Launch the SageMaker notebook instance within the VPC and create an S3 VPC endpoint for the notebook to access the data Copy the JSON dataset from Amazon S3 into the ML storage volume on the SageMaker notebook instance and work against the local dataset

C. Launch the SageMaker notebook instance within the VPC and create an S3 VPC endpoint for the notebook to access the data Define a custom S3 bucket policy to only allow requests from your VPC to access the S3 bucket

D. Launch the SageMaker notebook instance within the VPC with SageMaker-provided internet access enabled. Generate an S3 pre-signed URL for access to data in the bucket

Correct Answer: B

QUESTION 4

A company wants to classify user behavior as either fraudulent or normal. Based on internal research, a machine learning specialist will build a binary classifier based on two features: age of account, denoted by x, and transaction month, denoted by y. The class distributions are illustrated in the provided figure. The positive class is portrayed in red, while the negative class is portrayed in black.

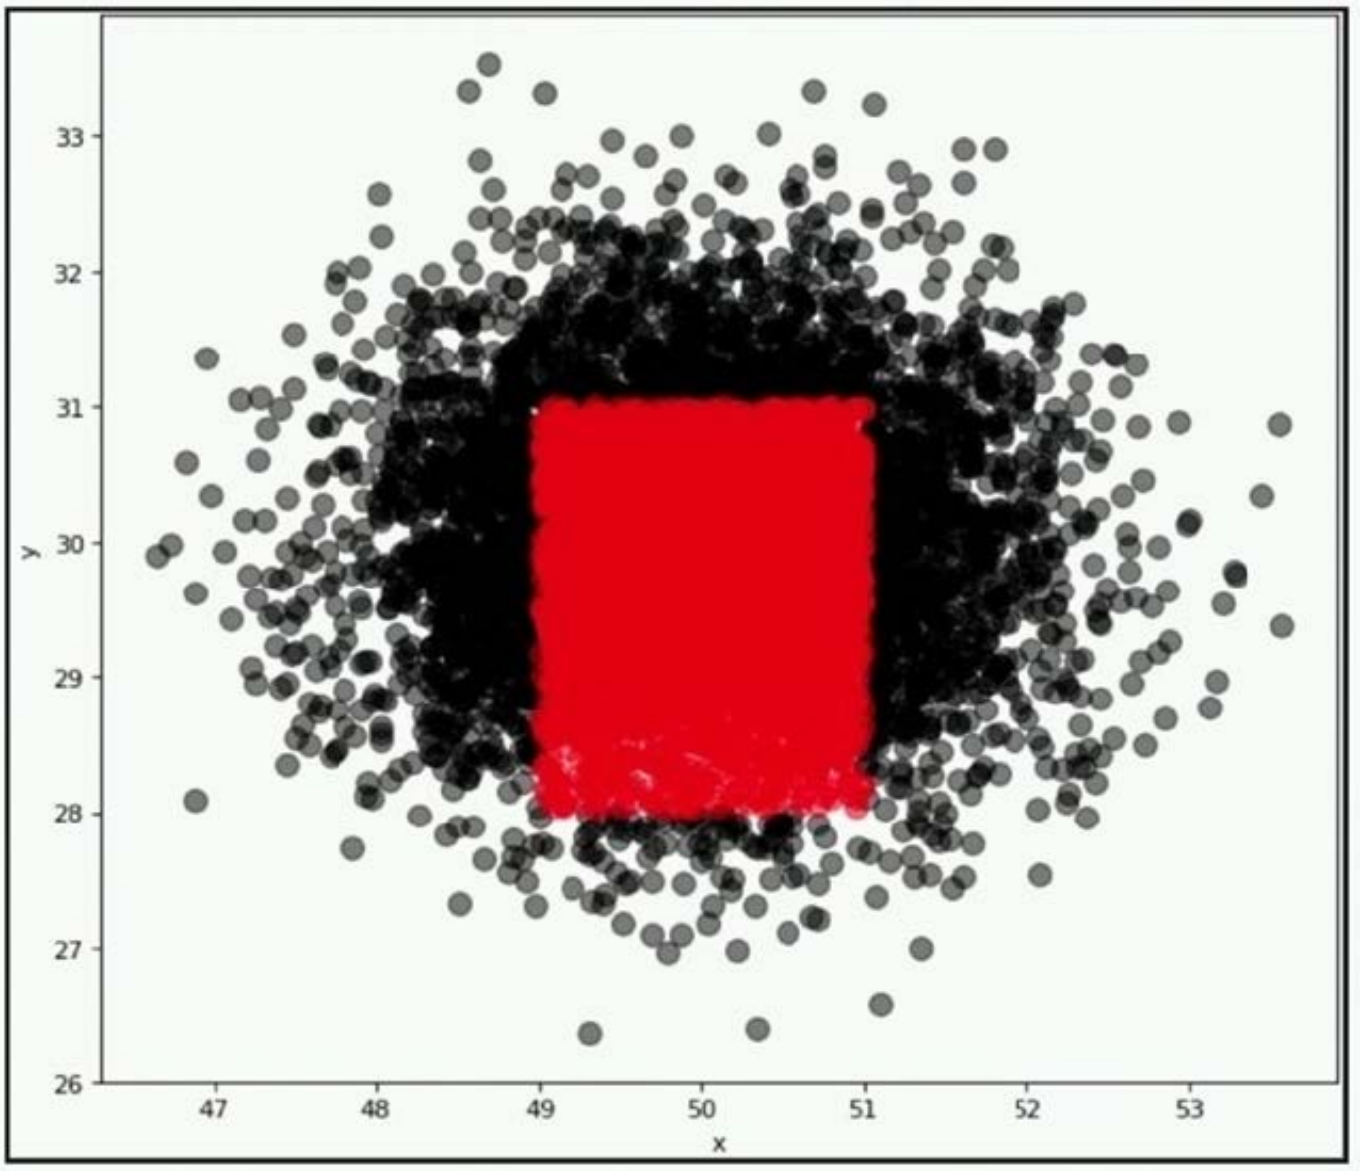

Which model would have the HIGHEST accuracy?

A. Linear support vector machine (SVM)

B. Decision tree

C. Support vector machine (SVM) with a radial basis function kernel

D. Single perceptron with a Tanh activation function

Correct Answer: B

A Decision tree produces a stepwise boundary that consists of rectilinear splits in the feature space that are perpendicular to the axes. The boundary is determined by the conditions specified in the tree.

Support Vector Machines (SVM) produce a non-linear boundary in the feature space that separates the classes. The boundary is defined as the maximum margin hyperplane, which is the line that maximally separates the classes while having the greatest margin between the classes. The boundary can be linear, polynomial, radial basis function (RBF) or other types of non-linear functions, depending on the choice of kernel and the configuration of the SVM.

QUESTION 5

A data scientist is working on a model to predict a company\\'s required inventory stock levels. All historical data is stored in .csv files in the company\\'s data lake on Amazon S3. The dataset consists of approximately 500 GB of data The data scientist wants to use SQL to explore the data before training the model. The company wants to minimize costs.

Which option meets these requirements with the LEAST operational overhead?

A. Create an Amazon EMR cluster. Create external tables in the Apache Hive metastore, referencing the data that is stored in the S3 bucket. Explore the data from the Hive console.

B. Use AWS Glue to crawl the S3 bucket and create tables in the AWS Glue Data Catalog. Use Amazon Athena to explore the data.

C. Create an Amazon Redshift cluster. Use the COPY command to ingest the data from Amazon S3. Explore the data from the Amazon Redshift query editor GUI.

D. Create an Amazon Redshift cluster. Create external tables in an external schema, referencing the S3 bucket that contains the data. Explore the data from the Amazon Redshift query editor GUI.

Correct Answer: B