

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 machine learning (ML) engineer uses Bayesian optimization for a hyperpara meter tuning job in Amazon SageMaker. The ML engineer uses precision as the objective metric.

The ML engineer wants to use recall as the objective metric. The ML engineer also wants to expand the hyperparameter range for a new hyperparameter tuning job. The new hyperparameter range will include the range of the previously

performed tuning job.

Which approach will run the new hyperparameter tuning job in the LEAST amount of time?

A. Use a warm start hyperparameter tuning job.

B. Use a checkpointing hyperparameter tuning job.

C. Use the same random seed for the hyperparameter tuning job.

D. Use multiple jobs in parallel for the hyperparameter tuning job.

Correct Answer: A

QUESTION 2

A Data Scientist is developing a machine learning model to predict future patient outcomes based on information collected about each patient and their treatment plans. The model should output a continuous value as its prediction. The data

available includes labeled outcomes for a set of 4,000 patients. The study was conducted on a group of individuals over the age of 65 who have a particular disease that is known to worsen with age.

Initial models have performed poorly. While reviewing the underlying data, the Data Scientist notices that, out of 4,000 patient observations, there are 450 where the patient age has been input as 0. The other features for these observations

appear normal compared to the rest of the sample population.

How should the Data Scientist correct this issue?

A. Drop all records from the dataset where age has been set to 0.

B. Replace the age field value for records with a value of 0 with the mean or median value from the dataset.

C. Drop the age feature from the dataset and train the model using the rest of the features.

D. Use k-means clustering to handle missing features.

Correct Answer: D

Dropping the Age feature is a NOT ATOLL a good idea - as age plays a critical role in this disease as per the question Dropping 10% of data is NOT a good idea considering the fact that the number of observations is already low. The Mean or Median are a potential solutions But the question says that "Disease worsens after age 65 so there is a correlation between age and other symptoms related feature" So that means that using Unsupervised Learning we can



make pretty good prediction of "Age" So the answer is D Use K-Means clustering

QUESTION 3

A telecommunications company is developing a mobile app for its customers. The company is using an Amazon SageMaker hosted endpoint for machine learning model inferences.

Developers want to introduce a new version of the model for a limited number of users who subscribed to a preview feature of the app. After the new version of the model is tested as a preview, developers will evaluate its accuracy. If a new version of the model has better accuracy, developers need to be able to gradually release the new version for all users over a fixed period of time.

How can the company implement the testing model with the LEAST amount of operational overhead?

A. Update the ProductionVariant data type with the new version of the model by using the CreateEndpointConfigoperation with the InitialVariantWeightparameter set to 0. Specify the TargetVariantparameter for InvokeEndpoint calls for users who subscribed to the preview feature. When the new version of the model is ready for release, gradually increase InitialVariantWeightuntil all users have the updated version.

B. Configure two SageMaker hosted endpoints that serve the different versions of the model. Create an Application Load Balancer (ALB) to route traffic to both endpoints based on the TargetVariantquery string parameter. Reconfigure the app to send the TargetVariant query string parameter for users who subscribed to the preview feature. When the new version of the model is ready for release, change the ALB\\'s routing algorithm to weighted until all users have the

updated version.

C. Update the DesiredWeightsAndCapacity data type with the new version of the model by using the UpdateEndpointWeightsAndCapacitiesoperation with the DesiredWeightparameter set to 0. Specify the TargetVariantparameter for InvokeEndpoint calls for users who subscribed to the preview feature. When the new version of the model is ready for release, gradually increase DesiredWeightuntil all users have the updated version.

D. Configure two SageMaker hosted endpoints that serve the different versions of the model. Create an Amazon Route 53 record that is configured with a simple routing policy and that points to the current version of the model. Configure the mobile app to use the endpoint URL for users who subscribed to the preview feature and to use the Route 53 record for other users. When the new version of the model is ready for release, add a new model version endpoint to Route 53, and switch the policy to weighted until all users have the updated version.

Correct Answer: C

https://docs.aws.amazon.com/sagemaker/latest/dg/model-ab-testing.html

QUESTION 4

An online delivery company wants to choose the fastest courier for each delivery at the moment an order is placed. The company wants to implement this feature for existing users and new users of its application. Data scientists have trained separate models with XGBoost for this purpose, and the models are stored in Amazon S3. There is one model fof each city where the company operates.

The engineers are hosting these models in Amazon EC2 for responding to the web client requests, with one instance for each model, but the instances have only a 5% utilization in CPU and memory,operation engineers want to avoid managing unnecessary resources.

Which solution will enable the company to achieve its goal with the LEAST operational overhead?



A. Create an Amazon SageMaker notebook instance for pulling all the models from Amazon S3 using the boto3 library. Remove the existing instances and use the notebook to perform a SageMaker batch transform for performing inferences offline for all the possible users in all the cities. Store the results in different files in Amazon S3. Point the web client to the files.

B. Prepare an Amazon SageMaker Docker container based on the open-source multi- model server. Remove the existing instances and create a multi-model endpoint in SageMaker instead, pointing to the S3 bucket containing all the models Invoke the endpoint from the web client at runtime, specifying the TargetModel parameter according to the city of each request.

C. Keep only a single EC2 instance for hosting all the models. Install a model server in the instance and load each model by pulling it from Amazon S3. Integrate the instance with the web client using Amazon API Gateway for responding to the requests in real time, specifying the target resource according to the city of each request.

D. Prepare a Docker container based on the prebuilt images in Amazon SageMaker. Replace the existing instances with separate SageMaker endpoints. one for each city where the company operates. Invoke the endpoints from the web client, specifying the URL and EndpomtName parameter according to the city of each request.

Correct Answer: B

The best solution for this scenario is to use a multi-model endpoint in Amazon SageMaker, which allows hosting multiple models on the same endpoint and invoking them dynamically at runtime. This way, the company can reduce the operational overhead of managing multiple EC2 instances and model servers, and leverage the scalability, security, and performance of SageMaker hosting services. By using a multi- model endpoint, the company can also save on hosting costs by improving endpoint utilization and paying only for the models that are loaded in memory and the API calls that are made. To use a multi-model endpoint, the company needs to prepare a Docker container based on the open-source multi-model server, which is a framework-agnostic library that supports loading and serving multiple models from Amazon S3. The company can then create a multi-model endpoint in SageMaker, pointing to the S3 bucket containing all the models, and invoke the endpoint from the web client at runtime, specifying the TargetModel parameter according to the city of each request. This solution also enables the company to add or remove models from the S3 bucket without redeploying the endpoint, and to use different versions of the same model for different cities if needed. References: Use Docker containers to build models Host multiple models in one container behind one endpoint Multi-model endpoints using Scikit Learn Multi-model endpoints using XGBoost

QUESTION 5

A sports broadcasting company is planning to introduce subtitles in multiple languages for a live broadcast. The commentary is in English. The company needs the transcriptions to appear on screen in French or Spanish, depending on the broadcasting country. The transcriptions must be able to capture domain-specific terminology, names, and locations based on the commentary context. The company needs a solution that can support options to provide tuning data.

Which combination of AWS services and features will meet these requirements with the LEAST operational overhead? (Choose two.)

- A. Amazon Transcribe with custom vocabularies
- B. Amazon Transcribe with custom language models
- C. Amazon SageMaker Seq2Seq
- D. Amazon SageMaker with Hugging Face Speech2Text
- E. Amazon Translate



Correct Answer: BE

Latest MLS-C01 Dumps

MLS-C01 VCE Dumps

MLS-C01 Exam Questions