PROFESSIONAL-MACHINE-LEARNING-ENGINEER Q&As

Professional Machine Learning Engineer

Pass Google PROFESSIONAL-MACHINE-LEARNING-ENGINEER Exam with 100% Guarantee

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

https://www.geekcert.com/professional-machine-learning-engineer.html

100% Passing Guarantee
100% Money Back Assurance

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

QUESTION 1

You have deployed a model on Vertex AI for real-time inference. During an online prediction request, you get an "Out of Memory" error. What should you do?

A. Use batch prediction mode instead of online mode.

B. Send the request again with a smaller batch of instances.

C. Use base64 to encode your data before using it for prediction.

D. Apply for a quota increase for the number of prediction requests.

Correct Answer: B

https://cloud.google.com/ai-platform/training/docs/troubleshooting

QUESTION 2

You deployed an ML model into production a year ago. Every month, you collect all raw requests that were sent to your model prediction service during the previous month. You send a subset of these requests to a human labeling service to evaluate your model\\'s performance. After a year, you notice that your model\\'s performance sometimes degrades significantly after a month, while other times it takes several months to notice any decrease in performance. The labeling service is costly, but you also need to avoid large performance degradations. You want to determine how often you should retrain your model to maintain a high level of performance while minimizing cost. What should you do?

A. Train an anomaly detection model on the training dataset, and run all incoming requests through this model. If an anomaly is detected, send the most recent serving data to the labeling service.

B. Identify temporal patterns in your model\\'s performance over the previous year. Based on these patterns, create a schedule for sending serving data to the labeling service for the next year.

C. Compare the cost of the labeling service with the lost revenue due to model performance degradation over the past year. If the lost revenue is greater than the cost of the labeling service, increase the frequency of model retraining; otherwise, decrease the model retraining frequency.

D. Run training-serving skew detection batch jobs every few days to compare the aggregate statistics of the features in the training dataset with recent serving data. If skew is detected, send the most recent serving data to the labeling service.

Correct Answer: D

https://cloud.google.com/blog/topics/developers-practitioners/monitor-models-training-serving-skew-vertex-aiew-vertex-aiandved=2ahUKEwiRg_aoj9n8AhWb7TgGHcGCDREQFnoECAwQAQandusg=AOvVaw197NneIJM0ra7fLq2zsOin

QUESTION 3

One of your models is trained using data provided by a third-party data broker. The data broker does not reliably notify you of formatting changes in the data. You want to make your model training pipeline more robust to issues like this. What should you do?

A. Use TensorFlow Data Validation to detect and flag schema anomalies.

B. Use TensorFlow Transform to create a preprocessing component that will normalize data to the expected distribution, and replace values that don\\'t match the schema with 0.

C. Use tf.math to analyze the data, compute summary statistics, and flag statistical anomalies.

D. Use custom TensorFlow functions at the start of your model training to detect and flag known formatting errors.

Correct Answer: A

https://www.tensorflow.org/tfx/guide/tfdv#schema_based_example_validation

QUESTION 4

You have been given a dataset with sales predictions based on your company\\'s marketing activities. The data is structured and stored in BigQuery, and has been carefully managed by a team of data analysts. You need to prepare a report providing insights into the predictive capabilities of the data. You were asked to run several ML models with different levels of sophistication, including simple models and multilayered neural networks. You only have a few hours to gather the results of your experiments. Which Google Cloud tools should you use to complete this task in the most efficient and self-serviced way?

A. Use BigQuery ML to run several regression models, and analyze their performance.

B. Read the data from BigQuery using Dataproc, and run several models using SparkML.

C. Use Vertex AI Workbench user-managed notebooks with scikit-learn code for a variety of ML algorithms and performance metrics.

D. Train a custom TensorFlow model with Vertex AI, reading the data from BigQuery featuring a variety of ML algorithms.

Correct Answer: A

QUESTION 5

You work on a data science team at a bank and are creating an ML model to predict loan default risk. You have collected and cleaned hundreds of millions of records worth of training data in a BigQuery table, and you now want to develop and compare multiple models on this data using TensorFlow and Vertex AI. You want to minimize any bottlenecks during the data ingestion state while considering scalability. What should you do?

A. Use the BigQuery client library to load data into a dataframe, and use tf.data.Dataset.from_tensor_slices() to read it.

B. Export data to CSV files in Cloud Storage, and use tf.data.TextLineDataset() to read them.

C. Convert the data into TFRecords, and use tf.data.TFRecordDataset() to read them.

D. Use TensorFlow I/O\\'s BigQuery Reader to directly read the data.

Correct Answer: D

https://www.tensorflow.org/io/api_docs/python/tfio/bigquery

Latest PROFESSIONAL-MACHINE-LEARNING-ENGINEER Dumps	PROFESSIONAL-MACHINE-LEARNING-ENGINEER PDF Dumps	PROFESSIONAL-MACHINE-LEARNING-ENGINEER Study Guide