

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.passapply.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 https://www.passapply.com/professional-machine-learning-engineer.html 2024 Latest passapply PROFESSIONAL-MACHINE-LEARNING-ENGINEER PDF and VCE dumps Download

- Instant Download After Purchase
- 100% Money Back Guarantee
- 😳 365 Days Free Update

VCE & PDF

PassApply.com

800,000+ Satisfied Customers





QUESTION 1

You are training and deploying updated versions of a regression model with tabular data by using Vertex AI Pipelines, Vertex AI Training, Vertex AI Experiments, and Vertex AI Endpoints. The model is deployed in a Vertex AI endpoint, and your users call the model by using the Vertex AI endpoint. You want to receive an email when the feature data distribution changes significantly, so you can retrigger the training pipeline and deploy an updated version of your model. What should you do?

A. Use Vertex Al Model Monitoring. Enable prediction drift monitoring on the endpoint, and specify a notification email.

B. In Cloud Logging, create a logs-based alert using the logs in the Vertex AI endpoint. Configure Cloud Logging to send an email when the alert is triggered.

C. In Cloud Monitoring create a logs-based metric and a threshold alert for the metric. Configure Cloud Monitoring to send an email when the alert is triggered.

D. Export the container logs of the endpoint to BigQuery. Create a Cloud Function to run a SQL query over the exported logs and send an email. Use Cloud Scheduler to trigger the Cloud Function.

Correct Answer: A

QUESTION 2

You are creating an ML pipeline for data processing, model training, and model deployment that uses different Google Cloud services. You have developed code for each individual task, and you expect a high frequency of new files. You now need to create an orchestration layer on top of these tasks. You only want this orchestration pipeline to run if new files are present in your dataset in a Cloud Storage bucket. You also want to minimize the compute node costs. What should you do?

A. Create a pipeline in Vertex AI Pipelines. Configure the first step to compare the contents of the bucket to the last time the pipeline was run. Use the scheduler API to run the pipeline periodically.

B. Create a Cloud Function that uses a Cloud Storage trigger and deploys a Cloud Composer directed acyclic graph (DAG).

C. Create a pipeline in Vertex AI Pipelines. Create a Cloud Function that uses a Cloud Storage trigger and deploys the pipeline.

D. Deploy a Cloud Composer directed acyclic graph (DAG) with a GCSObjectUpdateSensor class that detects when a new file is added to the Cloud Storage bucket.

Correct Answer: C

QUESTION 3

Your organization manages an online message board. A few months ago, you discovered an increase in toxic language and bullying on the message board. You deployed an automated text classifier that flags certain comments as toxic or harmful. Now some users are reporting that benign comments referencing their religion are being misclassified as abusive. Upon further inspection, you find that your classifier\\'s false positive rate is higher for comments that reference certain underrepresented religious groups. Your team has a limited budget and is already overextended. What should you do?



- A. Add synthetic training data where those phrases are used in non-toxic ways.
- B. Remove the model and replace it with human moderation.
- C. Replace your model with a different text classifier.
- D. Raise the threshold for comments to be considered toxic or harmful.

Correct Answer: D

QUESTION 4

You work for a social media company. You want to create a no-code image classification model for an iOS mobile application to identify fashion accessories. You have a labeled dataset in Cloud Storage. You need to configure a training workflow that minimizes cost and serves predictions with the lowest possible latency. What should you do?

A. Train the model by using AutoML, and register the model in Vertex Al Model Registry. Configure your mobile application to send batch requests during prediction.

B. Train the model by using AutoML Edge, and export it as a Core ML model. Configure your mobile application to use the .mlmodel file directly.

C. Train the model by using AutoML Edge, and export the model as a TFLite model. Configure your mobile application to use the .tflite file directly.

D. Train the model by using AutoML, and expose the model as a Vertex AI endpoint. Configure your mobile application to invoke the endpoint during prediction.

Correct Answer: B

QUESTION 5

You work for a bank and are building a random forest model for fraud detection. You have a dataset that includes transactions, of which 1% are identified as fraudulent. Which data transformation strategy would likely improve the performance of your classifier?

A. Modify the target variable using the Box-Cox transformation.

- B. Z-normalize all the numeric features.
- C. Oversample the fraudulent transaction 10 times.
- D. Log transform all numeric features.

Correct Answer: C

https://towardsdatascience.com/how-to-build-a-machine-learning-model-to-identify-credit-card-fraud-in-5-stepsa-hands-on-modeling-5140b3bd19f1



Latest PROFESSIONAL-MA PROFESSIONAL-MACHIN **CHINE-LEARNING-ENGINEER Dumps**

E-LEARNING-ENGINEER PDF Dumps

PROFESSIONAL-MACHIN E-LEARNING-ENGINEER Study Guide