



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



VCE & PDF

PassApply.com

<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
- ⚙️ **800,000+** Satisfied Customers





QUESTION 1

You were asked to investigate failures of a production line component based on sensor readings. After receiving the dataset, you discover that less than 1% of the readings are positive examples representing failure incidents. You have tried to train several classification models, but none of them converge. How should you resolve the class imbalance problem?

- A. Use the class distribution to generate 10% positive examples.
- B. Use a convolutional neural network with max pooling and softmax activation.
- C. Downsample the data with upweighting to create a sample with 10% positive examples.
- D. Remove negative examples until the numbers of positive and negative examples are equal.

Correct Answer: C

<https://developers.google.com/machine-learning/data-prep/construct/sampling-splitting/imbalanced-data#downsampling-and-upweighting>

-less than 1% of the readings are positive

-none of them converge.

Downsampling (in this context) means training on a disproportionately low subset of the majority class examples.

QUESTION 2

You have trained an XGBoost model that you plan to deploy on Vertex AI for online prediction. You are now uploading your model to Vertex AI Model Registry, and you need to configure the explanation method that will serve online prediction requests to be returned with minimal latency. You also want to be alerted when feature attributions of the model meaningfully change over time. What should you do?

- A. 1. Specify sampled Shapley as the explanation method with a path count of 5.

2.

Deploy the model to Vertex AI Endpoints.

3.

Create a Model Monitoring job that uses prediction drift as the monitoring objective.

- B. 1. Specify Integrated Gradients as the explanation method with a path count of 5.

2.

Deploy the model to Vertex AI Endpoints.

3.

Create a Model Monitoring job that uses prediction drift as the monitoring objective.



C. 1. Specify sampled Shapley as the explanation method with a path count of 50.

2.

Deploy the model to Vertex AI Endpoints.

3.

Create a Model Monitoring job that uses training-serving skew as the monitoring objective.

D. 1. Specify Integrated Gradients as the explanation method with a path count of 50.

2.

Deploy the model to Vertex AI Endpoints.

3.

Create a Model Monitoring job that uses training-serving skew as the monitoring objective.

Correct Answer: A

QUESTION 3

You are developing ML models with AI Platform for image segmentation on CT scans. You frequently update your model architectures based on the newest available research papers, and have to rerun training on the same dataset to benchmark their performance. You want to minimize computation costs and manual intervention while having version control for your code. What should you do?

A. Use Cloud Functions to identify changes to your code in Cloud Storage and trigger a retraining job.

B. Use the gcloud command-line tool to submit training jobs on AI Platform when you update your code.

C. Use Cloud Build linked with Cloud Source Repositories to trigger retraining when new code is pushed to the repository.

D. Create an automated workflow in Cloud Composer that runs daily and looks for changes in code in Cloud Storage using a sensor.

Correct Answer: C

At the heart of this architecture is Cloud Build, infrastructure. Cloud Build can import source from Cloud Source Repositories, GitHub, or Bitbucket, and then execute a build to your specifications, and produce artifacts such as Docker containers or Python tar files.

QUESTION 4

You lead a data science team at a large international corporation. Most of the models your team trains are large-scale models using high-level TensorFlow APIs on AI Platform with GPUs. Your team usually takes a few weeks or months to iterate on a new version of a model. You were recently asked to review your team's spending. How should you reduce your Google Cloud compute costs without impacting the model's performance?



- A. Use AI Platform to run distributed training jobs with checkpoints.
- B. Use AI Platform to run distributed training jobs without checkpoints.
- C. Migrate to training with Kuberflow on Google Kubernetes Engine, and use preemptible VMs with checkpoints.
- D. Migrate to training with Kuberflow on Google Kubernetes Engine, and use preemptible VMs without checkpoints.

Correct Answer: C

<https://cloud.google.com/blog/products/ai-machine-learning/reduce-the-costs-of-ml-workflows-with-preemptible-vms-and-gpus?hl=en>

QUESTION 5

Your data science team has requested a system that supports scheduled model retraining, Docker containers, and a service that supports autoscaling and monitoring for online prediction requests. Which platform components should you choose for this system?

- A. Vertex AI Pipelines and App Engine
- B. Vertex AI Pipelines, Vertex AI Prediction, and Vertex AI Model Monitoring
- C. Cloud Composer, BigQuery ML, and Vertex AI Prediction
- D. Cloud Composer, Vertex AI Training with custom containers, and App Engine

Correct Answer: B

<https://cloud.google.com/vertex-ai/docs/training/containers-overview>

[Latest PROFESSIONAL-MACHINE-LEARNING-ENGINEER Dumps](#)

[PROFESSIONAL-MACHINE-LEARNING-ENGINEER Practice Test](#)

[PROFESSIONAL-MACHINE-LEARNING-ENGINEER Study Guide](#)