



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.pass4itsure.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

Pass4itSure.com

<https://www.pass4itsure.com/professional-machine-learning-engineer.html>

2024 Latest pass4itsure 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 developed an ML model with AI Platform, and you want to move it to production. You serve a few thousand queries per second and are experiencing latency issues. Incoming requests are served by a load balancer that distributes them across multiple Kubeflow CPU-only pods running on Google Kubernetes Engine (GKE). Your goal is to improve the serving latency without changing the underlying infrastructure. What should you do?

- A. Significantly increase the max_batch_size TensorFlow Serving parameter.
- B. Switch to the tensorflow-model-server-universal version of TensorFlow Serving.
- C. Significantly increase the max_enqueued_batches TensorFlow Serving parameter.
- D. Recompile TensorFlow Serving using the source to support CPU-specific optimizations. Instruct GKE to choose an appropriate baseline minimum CPU platform for serving nodes.

Correct Answer: D

QUESTION 2

You have trained a text classification model in TensorFlow using AI Platform. You want to use the trained model for batch predictions on text data stored in BigQuery while minimizing computational overhead. What should you do?

- A. Export the model to BigQuery ML.
- B. Deploy and version the model on AI Platform.
- C. Use Dataflow with the SavedModel to read the data from BigQuery.
- D. Submit a batch prediction job on AI Platform that points to the model location in Cloud Storage.

Correct Answer: D

<https://cloud.google.com/vertex-ai/docs/tabular-data/classification-regression/get-batch-predictions>

QUESTION 3

You have deployed multiple versions of an image classification model on AI Platform. You want to monitor the performance of the model versions over time. How should you perform this comparison?

- A. Compare the loss performance for each model on a held-out dataset.
- B. Compare the loss performance for each model on the validation data.
- C. Compare the receiver operating characteristic (ROC) curve for each model using the What-If Tool.
- D. Compare the mean average precision across the models using the Continuous Evaluation feature.

Correct Answer: D

Compare the mean average precision across the models using the Continuous Evaluation feature



<https://cloud.google.com/vertex-ai/docs/evaluation/introduction>

Vertex AI provides model evaluation metrics, such as precision and recall, to help you determine the performance of your models...

Vertex AI supports evaluation of the following model types:

AuPRC: The area under the precision-recall (PR) curve, also referred to as average precision. This value ranges from zero to one, where a higher value indicates a higher-quality model.

QUESTION 4

You work for a magazine distributor and need to build a model that predicts which customers will renew their subscriptions for the upcoming year. Using your company's historical data as your training set, you created a TensorFlow model and deployed it to AI Platform. You need to determine which customer attribute has the most predictive power for each prediction served by the model. What should you do?

- A. Use AI Platform notebooks to perform a Lasso regression analysis on your model, which will eliminate features that do not provide a strong signal.
- B. Stream prediction results to BigQuery. Use BigQuery's `CORR(X1, X2)` function to calculate the Pearson correlation coefficient between each feature and the target variable.
- C. Use the AI Explanations feature on AI Platform. Submit each prediction request with the `'explain'` keyword to retrieve feature attributions using the sampled Shapley method.
- D. Use the What-If tool in Google Cloud to determine how your model will perform when individual features are excluded. Rank the feature importance in order of those that caused the most significant performance drop when removed from the model.

Correct Answer: C

<https://cloud.google.com/ai-platform/prediction/docs/ai-explanations/overview>

QUESTION 5

You have been asked to develop an input pipeline for an ML training model that processes images from disparate sources at a low latency. You discover that your input data does not fit in memory. How should you create a dataset following Google-recommended best practices?

- A. Create a `tf.data.Dataset.prefetch` transformation.
- B. Convert the images to `tf.Tensor` objects, and then run `Dataset.from_tensor_slices()`.
- C. Convert the images to `tf.Tensor` objects, and then run `tf.data.Dataset.from_tensors()`.
- D. Convert the images into `TFRecords`, store the images in Cloud Storage, and then use the `tf.data` API to read the images for training.

Correct Answer: D

Cite from Google Pag: to construct a Dataset from data in memory, use `tf.data.Dataset.from_tensors()` or `tf.data.Dataset.from_tensor_slices()`. When input data is stored in a file (not in memory), the recommended `TFRecord`



format, you can use `tf.data.TFRecordDataset()`.

`tf.data.Dataset` is for data in memory. `tf.data.TFRecordDataset` is for data in non-memory storage.

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

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

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