



# PROFESSIONAL-MACHINE- LEARNING-ENGINEER<sup>Q&As</sup>

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 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-steps-a-hands-on-modeling-5140b3bd19f1>

---

### QUESTION 2

You are an ML engineer at a large grocery retailer with stores in multiple regions. You have been asked to create an inventory prediction model. Your model's features include region, location, historical demand, and seasonal popularity. You want the algorithm to learn from new inventory data on a daily basis. Which algorithms should you use to build the model?

- A. Classification
- B. Reinforcement Learning
- C. Recurrent Neural Networks (RNN)
- D. Convolutional Neural Networks (CNN)

Correct Answer: C

### QUESTION 3

You work for an online retail company that is creating a visual search engine. You have set up an end-to-end ML pipeline on Google Cloud to classify whether an image contains your company's product. Expecting the release of new products in the near future, you configured a retraining functionality in the pipeline so that new data can be fed into your ML models. You also want to use AI Platform's continuous evaluation service to ensure that the models have high accuracy on your test dataset. What should you do?

- A. Keep the original test dataset unchanged even if newer products are incorporated into retraining.
- B. Extend your test dataset with images of the newer products when they are introduced to retraining.
- C. Replace your test dataset with images of the newer products when they are introduced to retraining.
- D. Update your test dataset with images of the newer products when your evaluation metrics drop below a pre-decided



threshold.

Correct Answer: B

You need to correctly classify newer products, so you need the new training data ==> A is wrong;

You need to keep doing a good job on older dataset, you can't just ignore it ==> C is wrong;

You know when you are introducing new products, there is no need to wait for a drop in preformaces ==> D is wrong;

B is correct

---

#### QUESTION 4

You have trained a model on a dataset that required computationally expensive preprocessing operations. You need to execute the same preprocessing at prediction time. You deployed the model on AI Platform for high-throughput online prediction. Which architecture should you use?

- A. Validate the accuracy of the model that you trained on preprocessed data. Create a new model that uses the raw data and is available in real time. Deploy the new model onto AI Platform for online prediction.
- B. Send incoming prediction requests to a Pub/Sub topic. Transform the incoming data using a Dataflow job. Submit a prediction request to AI Platform using the transformed data. Write the predictions to an outbound Pub/Sub queue.
- C. Stream incoming prediction request data into Cloud Spanner. Create a view to abstract your preprocessing logic. Query the view every second for new records. Submit a prediction request to AI Platform using the transformed data. Write the predictions to an outbound Pub/Sub queue.
- D. Send incoming prediction requests to a Pub/Sub topic. Set up a Cloud Function that is triggered when messages are published to the Pub/Sub topic. Implement your preprocessing logic in the Cloud Function. Submit a prediction request to AI Platform using the transformed data. Write the predictions to an outbound Pub/Sub queue.

Correct Answer: B

[https://cloud.google.com/architecture/data-preprocessing-for-ml-with-tf-transform-pt1#where\\_to\\_do\\_preprocessing](https://cloud.google.com/architecture/data-preprocessing-for-ml-with-tf-transform-pt1#where_to_do_preprocessing)

---

#### QUESTION 5

You are developing a model to predict whether a failure will occur in a critical machine part. You have a dataset consisting of a multivariate time series and labels indicating whether the machine part failed. You recently started experimenting with a few different preprocessing and modeling approaches in a Vertex AI Workbench notebook. You want to log data and track artifacts from each run. How should you set up your experiments?

- A. 1. Use the Vertex AI SDK to create an experiment and set up Vertex ML Metadata.  
2. Use the `log_time_series_metrics` function to track the preprocessed data, and use the `log_merriics` function to log loss values.
- B. 1. Use the Vertex AI SDK to create an experiment and set up Vertex ML Metadata.  
2. Use the `log_time_series_metrics` function to track the preprocessed data, and use the `log_metrics` function to log loss values.



C. 1. Create a Vertex AI TensorBoard instance and use the Vertex AI SDK to create an experiment and associate the TensorBoard instance.

2. Use the `assign_input_artifact` method to track the preprocessed data and use the `log_time_series_metrics` function to log loss values.

D. 1. Create a Vertex AI TensorBoard instance, and use the Vertex AI SDK to create an experiment and associate the TensorBoard instance.

2. Use the `log_time_series_metrics` function to track the preprocessed data, and use the `log_metrics` function to log loss values.

Correct Answer: B

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

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

[PROFESSIONAL-MACHINE-LEARNING-ENGINEER Exam Questions](#)