



# 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.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



VCE & PDF

GeekCert.com

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

2024 Latest geekcert 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 are an ML engineer at a global shoe store. You manage the ML models for the company's website. You are asked to build a model that will recommend new products to the user based on their purchase behavior and similarity with other users. What should you do?

- A. Build a classification model
- B. Build a knowledge-based filtering model
- C. Build a collaborative-based filtering model
- D. Build a regression model using the features as predictors

Correct Answer: C

Reference: <https://cloud.google.com/solutions/recommendations-using-machine-learning-on-compute-engine>

---

### QUESTION 2

You are building a real-time prediction engine that streams files which may contain Personally Identifiable Information (PII) to Google Cloud. You want to use the Cloud Data Loss Prevention (DLP) API to scan the files. How should you ensure that the PII is not accessible by unauthorized individuals?

- A. Stream all files to Google Cloud, and then write the data to BigQuery. Periodically conduct a bulk scan of the table using the DLP API.
- B. Stream all files to Google Cloud, and write batches of the data to BigQuery. While the data is being written to BigQuery, conduct a bulk scan of the data using the DLP API.
- C. Create two buckets of data: Sensitive and Non-sensitive. Write all data to the Non-sensitive bucket. Periodically conduct a bulk scan of that bucket using the DLP API, and move the sensitive data to the Sensitive bucket.
- D. Create three buckets of data: Quarantine, Sensitive, and Non-sensitive. Write all data to the Quarantine bucket. Periodically conduct a bulk scan of that bucket using the DLP API, and move the data to either the Sensitive or Non-Sensitive bucket.

Correct Answer: D

[https://cloud.google.com/architecture/automating-classification-of-data-uploaded-to-cloud-storage#building\\_the\\_quarantine\\_and\\_classification\\_pipeline](https://cloud.google.com/architecture/automating-classification-of-data-uploaded-to-cloud-storage#building_the_quarantine_and_classification_pipeline)

---

### QUESTION 3

You have a large corpus of written support cases that can be classified into 3 separate categories: Technical Support, Billing Support, or Other Issues. You need to quickly build, test, and deploy a service that will automatically classify future written requests into one of the categories. How should you configure the pipeline?

- A. Use the Cloud Natural Language API to obtain metadata to classify the incoming cases.
- B. Use AutoML Natural Language to build and test a classifier. Deploy the model as a REST API.



C. Use BigQuery ML to build and test a logistic regression model to classify incoming requests. Use BigQuery ML to perform inference.

D. Create a TensorFlow model using Google's BERT pre-trained model. Build and test a classifier, and deploy the model using Vertex AI.

Correct Answer: B

#### QUESTION 4

You trained a text classification model. You have the following SignatureDefs:

```
signature_def['serving_default']:  
  The given SavedModel SignatureDef contains the following input(s):  
    inputs['text'] tensor_info:  
      dtype: DT_STRING  
      shape: (-1, 2)  
      name: serving_default_text: 0  
  The given SavedModel SignatureDef contains the following output(s):  
    outputs ['Softmax'] tensor_info:  
      dtype: DT_FLOAT  
      shape: (-1, 2)  
      name: StatefulPartitionedCall:0  
  Method name is: tensorflow/serving/predict
```

You started a TensorFlow-serving component server and tried to send an HTTP request to get a prediction using:

```
headers = {"content-type": "application/json"}
```

```
json_response = requests.post('http://localhost:8501/v1/models/text_model:predict', data=data, headers=headers)
```

What is the correct way to write the predict request?

- A. data = json.dumps({"signature\_name": "seving\_default", "instances" [['ab', 'bc', 'cd']]})
- B. data = json.dumps({"signature\_name": "serving\_default", "instances" [['a', 'b', 'c', 'd', 'e', 'f']]})
- C. data = json.dumps({"signature\_name": "serving\_default", "instances" [['a', 'b', 'c'], [d, 'e', 'f']]})
- D. data = json.dumps({"signature\_name": "serving\_default", "instances" [['a', 'b'], [c, 'd'], [e, 'f']]})

Correct Answer: D

A negative number in the shape enables auto expand (<https://stackoverflow.com/questions/37956197/what-is-the-negative-index-in-shape-arrays-used-for-tensorflow>).

Then the first number -1 out of the shape (-1, 2) speaks the number of 1 dimensional arrays within the tensor (and it can autoexpand) while the second number (2) sets the number of elements in the inner array at 2. Hence D.



## QUESTION 5

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

[Latest PROFESSIONAL-MA  
CHINE-LEARNING-  
ENGINEER Dumps](#)

[PROFESSIONAL-MACHIN  
E-LEARNING-ENGINEER  
VCE Dumps](#)

[PROFESSIONAL-MACHIN  
E-LEARNING-ENGINEER  
Practice Test](#)