

# 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.pass2lead.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

- ⚙️ **Instant Download** After Purchase
- ⚙️ **100% Money Back** Guarantee
- ⚙️ **365 Days** Free Update
- ⚙️ **800,000+** Satisfied Customers



### QUESTION 1

You work on the data science team for a multinational beverage company. You need to develop an ML model to predict the company's profitability for a new line of naturally flavored bottled waters in different locations. You are provided with historical data that includes product types, product sales volumes, expenses, and profits for all regions. What should you use as the input and output for your model?

- A. Use latitude, longitude, and product type as features. Use profit as model output.
- B. Use latitude, longitude, and product type as features. Use revenue and expenses as model outputs.
- C. Use product type and the feature cross of latitude with longitude, followed by binning, as features. Use profit as model output.
- D. Use product type and the feature cross of latitude with longitude, followed by binning, as features. Use revenue and expenses as model outputs.

Correct Answer: C

<https://developers.google.com/machine-learning/crash-course/feature-crosses/video-lecture>  
<https://developers.google.com/machine-learning/crash-course/regularization-for-sparsity/l1-regularization>

---

### QUESTION 2

You are developing a Kubeflow pipeline on Google Kubernetes Engine. The first step in the pipeline is to issue a query against BigQuery. You plan to use the results of that query as the input to the next step in your pipeline. You want to achieve this in the easiest way possible. What should you do?

- A. Use the BigQuery console to execute your query, and then save the query results into a new BigQuery table.
- B. Write a Python script that uses the BigQuery API to execute queries against BigQuery. Execute this script as the first step in your Kubeflow pipeline.
- C. Use the Kubeflow Pipelines domain-specific language to create a custom component that uses the Python BigQuery client library to execute queries.
- D. Locate the Kubeflow Pipelines repository on GitHub. Find the BigQuery Query Component, copy that component's URL, and use it to load the component into your pipeline. Use the component to execute queries against BigQuery.

Correct Answer: D

Kubeflow pipelines have different types of components, ranging from low- to high-level. They have a ComponentStore that allows you to access prebuilt functionality from GitHub.

<https://github.com/kubeflow/pipelines/blob/master/components/gcp/bigquery/query/sample.ipynb>;  
<https://v0-5.kubeflow.org/docs/pipelines/reusable-components/>

---

### QUESTION 3

You work for a toy manufacturer that has been experiencing a large increase in demand. You need to build an ML model to reduce the amount of time spent by quality control inspectors checking for product defects. Faster defect detection is a priority. The factory does not have reliable Wi-Fi. Your company wants to implement the new ML model as

soon as possible. Which model should you use?

- A. AutoML Vision Edge mobile-high-accuracy-1 model
- B. AutoML Vision Edge mobile-low-latency-1 model
- C. AutoML Vision model
- D. AutoML Vision Edge mobile-versatile-1 model

Correct Answer: B

Hence faster defect detection is a priority, AutoML Vision Edge mobile-low-latency-1 model should be the choice. This model is designed to run efficiently on mobile devices and prioritize low latency, which means that it can provide fast defect detection without requiring a connection to the cloud. <https://cloud.google.com/vision/automl/docs/train-edge>

---

#### QUESTION 4

Your team is building a convolutional neural network (CNN)-based architecture from scratch. The preliminary experiments running on your on-premises CPU-only infrastructure were encouraging, but have slow convergence. You have been asked to speed up model training to reduce time-to-market. You want to experiment with virtual machines (VMs) on Google Cloud to leverage more powerful hardware. Your code does not include any manual device placement and has not been wrapped in Estimator model-level abstraction. Which environment should you train your model on?

- A. AVM on Compute Engine and 1 TPU with all dependencies installed manually.
- B. AVM on Compute Engine and 8 GPUs with all dependencies installed manually.
- C. A Deep Learning VM with an n1-standard-2 machine and 1 GPU with all libraries pre-installed.
- D. A Deep Learning VM with more powerful CPU e2-highcpu-16 machines with all libraries pre-installed.

Correct Answer: C

[https://cloud.google.com/deep-learning-vm/docs/cli#creating\\_an\\_instance\\_with\\_one\\_or\\_more\\_gpus](https://cloud.google.com/deep-learning-vm/docs/cli#creating_an_instance_with_one_or_more_gpus)  
[https://cloud.google.com/deep-learning-vm/docs/introduction#pre-installed\\_packages](https://cloud.google.com/deep-learning-vm/docs/introduction#pre-installed_packages)

---

#### QUESTION 5

As the lead ML Engineer for your company, you are responsible for building ML models to digitize scanned customer forms. You have developed a TensorFlow model that converts the scanned images into text and stores them in Cloud Storage. You need to use your ML model on the aggregated data collected at the end of each day with minimal manual intervention. What should you do?

- A. Use the batch prediction functionality of AI Platform.
- B. Create a serving pipeline in Compute Engine for prediction.
- C. Use Cloud Functions for prediction each time a new data point is ingested.
- D. Deploy the model on AI Platform and create a version of it for online inference.

Correct Answer: A

<https://cloud.google.com/ai-platform/prediction/docs/batch-predict>

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

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

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