# Google Google Professional-Machine-Learning-Engineer PDF

## Google Google Professional-Machine-Learning-Engineer PDF Questions Available Here at:

https://www.certification-questions.com//google-dumps/professional-machine-learning-engineer.html

Enrolling now you will get access to 71 questions in a unique set of Google Professional-Machine-Learning-Engineer

## Question 1

Your team needs to build a model that predicts whether images contain a driver's license, passport, or credit card. The data engineering team already built the pipeline and generated a dataset composed of 10,000 images with driver's licenses, 1,000 images with passports, and 1,000 images with credit cards. You now have to train a model with the following label map: ['driversjicense', 'passport', 'credit_card']. Which loss function should you use?

**Options:**

A. Categorical hinge

B. Binary cross-entropy

C. Categorical cross-entropy

D. Sparse categorical cross-entropy

**Answer: D**

**Explanation:**
se sparse_categorical_crossentropy. Examples for above 3-class classification problem: [1] , [2], [3]

## Question 2

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?

**Options:**

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

**Answer: B**

# Question 3

You are an ML engineer at a bank that has a mobile application. Management has asked you to build an ML-based biometric authentication for the app that verifies a customer's identity based on their fingerprint. Fingerprints are considered highly sensitive personal information and cannot be downloaded and stored into the bank databases. Which learning strategy should you recommend to train and deploy this ML model?

**Options:**

A. Differential privacy

B. Federated learning

C. MD5 to encrypt data

D. Data Loss Prevention API

**Answer: B**

# Question 4

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?

**Options:**

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

**Answer: A**

## Question 5

You are training a Resnet model on AI Platform using TPUs to visually categorize types of
defects in automobile engines. You capture the training profile using the Cloud TPU profiler plugin
and observe that it is highly input-bound. You want to reduce the bottleneck and speed up your
model training process. Which modifications should you make to the tf .data dataset?
Choose 2 answers

**Options:**

A. Use the interleave option for reading data

B. Reduce the value of the repeat parameter

C. Increase the buffer size for the shuffle option.

D. Set the prefetch option equal to the training batch size

E. Decrease the batch size argument in your transformation

**Answer: A, E**

## Question 6

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 models 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?

**Options:**

A. Classification

B. Reinforcement Learning

C. Recurrent Neural Networks (RNN)

D. Convolutional Neural Networks (CNN)

**Answer: B**

# Question 7

You have a demand forecasting pipeline in production that uses Dataflow to preprocess raw data prior to model training and prediction. During preprocessing, you employ Z-score normalization on data stored in BigQuery and write it back to BigQuery. New training data is added every week. You want to make the process more efficient by minimizing computation time and manual intervention. What should you do?

**Options:**

A. Normalize the data using Google Kubernetes Engine

B. Translate the normalization algorithm into SQL for use with BigQuery

C. Use the normalizer_fn argument in TensorFlow's Feature Column API

D. Normalize the data with Apache Spark using the Dataproc connector for BigQuery

**Answer: B**

# Question 8

You built and manage a production system that is responsible for predicting sales numbers. Model accuracy is crucial, because the production model is required to keep up with market changes. Since being deployed to production, the model hasn't changed; however the accuracy of the model has steadily deteriorated. What issue is most likely causing the steady decline in model accuracy?

**Options:**

A. Poor data quality

B. Lack of model retraining

C. Too few layers in the model for capturing information

D. Incorrect data split ratio during model training, evaluation, validation, and test

**Answer: D**

## Question  9

During batch training of a neural network, you notice that there is an oscillation in the loss.
How should you adjust your model to ensure that it converges?

**Options:**

A. Increase the size of the training batch

B. Decrease the size of the training batch

C. Increase the learning rate hyperparameter

D. Decrease the learning rate hyperparameter

**Answer: C**

## Question  10

You have written unit tests for a Kubeflow Pipeline that require custom libraries. You want to

automate the execution of unit tests with each new push to your development branch in Cloud
Source Repositories. What should you do?

**Options:**

A. Write a script that sequentially performs the push to your development branch and executes the
unit tests on Cloud Run

B. Using Cloud Build, set an automated trigger to execute the unit tests when changes are pushed
to
your development branch.

C. Set up a Cloud Logging sink to a Pub/Sub topic that captures interactions with Cloud Source
Repositories Configure a Pub/Sub trigger for Cloud Run, and execute the unit tests on Cloud Run.

D. Set up a Cloud Logging sink to a Pub/Sub topic that captures interactions with Cloud Source
Repositories. Execute the unit tests using a Cloud Function that is triggered when messages are
sent
to the Pub/Sub topic

**Answer: B**

# Would you like to see more? Don't miss our Google Professional-Machine-Learning-Engineer PDF file at:

https://www.certification-questions.com/google-pdf/professional-machine-learning-engineer-pdf.html