



DUMPSsheet

**Free Questions for Professional-Machine-Learning-Engineer
by dumpsheet**

Shared by Nieves on 20-10-2022

For More Free Questions and Preparation Resources

Check the Links on Last Page

Question 1

Question Type: MultipleChoice

You have built a model that is trained on data stored in Parquet files. You access the data through a Hive table hosted on Google Cloud. You preprocessed these data with PySpark and exported it as a CSV file into Cloud Storage. After preprocessing, you execute additional steps to train and evaluate your model. You want to parametrize this model training in Kubeflow Pipelines. What should you do?

Options:

- A-** Remove the data transformation step from your pipeline.
- B-** Containerize the PySpark transformation step, and add it to your pipeline.
- C-** Add a ContainerOp to your pipeline that spins a Dataproc cluster, runs a transformation, and then saves the transformed data in Cloud Storage.
- D-** Deploy Apache Spark at a separate node pool in a Google Kubernetes Engine cluster. Add a ContainerOp to your pipeline that invokes a corresponding transformation job for this Spark instance.

Answer:

D

Question 2

Question Type: MultipleChoice

Your data science team has requested a system that supports scheduled model retraining, Docker containers, and a service that supports autoscaling and monitoring for online prediction requests. Which platform components should you choose for this system?

Options:

- A- Kubeflow Pipelines and App Engine
- B- Kubeflow Pipelines and AI Platform Prediction
- C- Cloud Composer, BigQuery ML , and AI Platform Prediction
- D- Cloud Composer, AI Platform Training with custom containers , and App Engine

Answer:

B

Question 3

Question Type: MultipleChoice

You need to train a regression model based on a dataset containing 50,000 records that is stored in BigQuery. The data includes a total of 20 categorical and numerical features with a target variable that can include negative values. You need to minimize effort and training time while maximizing model performance. What approach should you take to train this regression model?

Options:

- A- Create a custom TensorFlow DNN model.
- B- Use BQML XGBoost regression to train the model
- C- Use AutoML Tables to train the model without early stopping.
- D- Use AutoML Tables to train the model with RMSLE as the optimization objective

Answer:

B

Explanation:

<https://cloud.google.com/bigquery-ml/docs/introduction>

Question 4

Question Type: MultipleChoice

Your task is classify if a company logo is present on an image. You found out that 96% of a data does not include a logo. You are dealing with data imbalance problem. Which metric do you use to evaluate to model?

Options:

- A- F1 Score
- B- RMSE
- C- F Score with higher precision weighting than recall
- D- F Score with higher recall weighted than precision

Answer:

D

Question 5

Question Type: MultipleChoice

You are developing an ML model intended to classify whether X-Ray images indicate bone fracture risk. You have trained on Api Resnet architecture on Vertex AI using a TPU as an accelerator, however you are unsatisfied with the training time and use memory usage. You want to quickly iterate your training code but make minimal changes to the code. You also want to minimize impact on the models accuracy. What should you do?

Options:

- A- Configure your model to use bfloat16 instead float32
- B- Reduce the global batch size from 1024 to 256
- C- Reduce the number of layers in the model architecture
- D- Reduce the dimensions of the images used un the model

Answer:

B

Question 6

Question Type: MultipleChoice

Your organization wants to make its internal shuttle service route more efficient. The shuttles currently stop at all pick-up points across the city every 30 minutes between 7 am and 10 am. The development team has already built an application on Google Kubernetes Engine that requires users to confirm their presence and shuttle station one day in advance. What approach should you take?

Options:

- A-** 1. Build a tree-based regression model that predicts how many passengers will be picked up at each shuttle station.
2. Dispatch an appropriately sized shuttle and provide the map with the required stops based on the prediction.
- B-** 1. Build a tree-based classification model that predicts whether the shuttle should pick up passengers at each shuttle station.
2. Dispatch an available shuttle and provide the map with the required stops based on the prediction
- C-** 1. Define the optimal route as the shortest route that passes by all shuttle stations with confirmed attendance at the given time under capacity constraints.
2 Dispatch an appropriately sized shuttle and indicate the required stops on the map
- D-** 1. Build a reinforcement learning model with tree-based classification models that predict the presence of passengers at shuttle stops as agents and a reward function around a distance-based metric
2. Dispatch an appropriately sized shuttle and provide the map with the required stops based on the simulated outcome.

Answer:

C

Explanation:

This is a case where machine learning would be terrible, as it would not be 100% accurate and some passengers would not get picked up. A simple algorithm works better here, and the question confirms customers will be indicating when they are at the stop so no ML required.

**To Get Premium Files for Professional-Machine-Learning-Engineer
Visit**

<https://www.p2pexams.com/products/professional-machine-learning-engineer>

For More Free Questions Visit

<https://www.p2pexams.com/google/pdf/professional-machine-learning-engineer>

