DUMPSARENA

Google Professional Machine Learning Engineer

Google Professional-Machine-Learning-Engineer

Version Demo

Total Demo Questions: 7

Total Premium Questions: 145

Buy Premium PDF

https://dumpsarena.co

sales@dumpsarena.co

sales@dumpsarena.co
dumpsarena.co

QUESTION NO: 1

You are an ML engineer at a mobile gaming company. A data scientist on your team recently trained a TensorFlow model, and you are responsible for deploying this model into a mobile application. You discover that the inference latency of the current model doesn't meet production requirements. You need to reduce the inference time by 50%, and you are willing to accept a small decrease in model accuracy in order to reach the latency requirement. Without training a new model, which model optimization technique for reducing latency should you try first?

A. Weight pruning

B. Dynamic range quantization

C. Model distillation

D. Dimensionality reduction

ANSWER: C

QUESTION NO: 2

You have a functioning end-to-end ML pipeline that involves tuning the hyperparameters of your ML model using AI Platform, and then using the best-tuned parameters for training. Hypertuning is taking longer than expected and is delaying the downstream processes. You want to speed up the tuning job without significantly compromising its effectiveness. Which actions should you take?

Choose 2 answers

A. Decrease the number of parallel trials

B. Decrease the range of floating-point values

C. Set the early stopping parameter to TRUE

D. Change the search algorithm from Bayesian search to random search.

E. Decrease the maximum number of trials during subsequent training phases.

ANSWER: C E

Explanation:

Reference: https://cloud.google.com/ai-platform/training/docs/hyperparameter-tuning-overview

https://cloud.google.com/ai-platform/training/docs/using-hyperparameter-tuning#early-stopping

QUESTION NO: 3

You want to rebuild your ML pipeline for structured data on Google Cloud. You are using PySpark to conduct data transformations at scale, but your pipelines are taking over 12 hours to run. To speed up development and pipeline run time, you want to use a serverless tool and SQL syntax. You have already moved your raw data into Cloud Storage. How should you build the pipeline on Google Cloud while meeting the speed and processing requirements?

A. Use Data Fusion's GUI to build the transformation pipelines, and then write the data into BigQuery

B. Convert your PySpark into SparkSQL queries to transform the data and then run your pipeline on Dataproc to write the data into BigQuery.

C. Ingest your data into Cloud SQL convert your PySpark commands into SQL queries to transform the data, and then use federated queries from BigQuery for machine learning

D. Ingest your data into BigQuery using BigQuery Load, convert your PySpark commands into BigQuery SQL queries to transform the data, and then write the transformations to a new table

ANSWER: D

Explanation:

Google has bought this software and support for this tool is not good. SQL can work in Cloud fusion pipelines too but I would prefer to use a single tool like Bigquery to both transform and store data.

QUESTION NO: 4

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

ANSWER: C

Explanation:

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

https://developers.google.com/machine-learning/recommendation/collaborative/basics

https://cloud.google.com/architecture/recommendations-using-machine-learning-on-compute-engine#filtering_the_data

QUESTION NO: 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

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: D E

Explanation:

https://towardsdatascience.com/overcoming-data-preprocessing-bottlenecks-with-tensorflow-data-service-nvidia-dali-and-other-d6321917f851

QUESTION NO: 6

You work for an advertising company and want to understand the effectiveness of your company's latest advertising campaign. You have streamed 500 MB of campaign data into BigQuery. You want to query the table, and then manipulate the results of that query with a pandas dataframe in an AI Platform notebook. What should you do?

A. Use AI Platform Notebooks' BigQuery cell magic to query the data, and ingest the results as a pandas dataframe

B. Export your table as a CSV file from BigQuery to Google Drive, and use the Google Drive API to ingest the file into your notebook instance

C. Download your table from BigQuery as a local CSV file, and upload it to your AI Platform notebook instance Use pandas. read_csv to ingest the file as a pandas dataframe

D. From a bash cell in your AI Platform notebook, use the bq extract command to export the table as a CSV file to Cloud Storage, and then use gsutii cp to copy the data into the notebook Use pandas. read_csv to ingest the file as a pandas dataframe

ANSWER: A

Explanation:

Refer to this link for details: https://cloud.google.com/bigquery/docs/bigquery-storage-python-pandas

First 2 points talks about querying the data.

Download query results to a pandas DataFrame by using the BigQuery Storage API from the IPython magics for BigQuery in a Jupyter notebook.

Download query results to a pandas DataFrame by using the BigQuery client library for Python.

Download BigQuery table data to a pandas DataFrame by using the BigQuery client library for Python.

Download BigQuery table data to a pandas DataFrame by using the BigQuery Storage API client library for Python.

https://googleapis.dev/python/bigquery/latest/magics.html#ipython-magics-for-bigquery

https://cloud.google.com/bigquery/docs/bigquery-storage-python-pandas

QUESTION NO: 7

You are going to train a DNN regression model with Keras APIs using this code:

```
model = tf.keras.Sequential()
model.add(tf.keras.layers.Dense(
    256,
    use_bias=True,
    activation='relu',
    kernel_initializer=None,
    kernel_regularizer=None,
    input_shape=(500,)))
model.add(tf.keras.layers.Dropout(rate=0.25))
model.add(tf.keras.layers.Dense(
    128, use_bias=True,
    activation='relu',
    kernel_initializer='uniform',
    kernel_regularizer='l2'))
model.add(tf.keras.layers.Dropout(rate=0.25))
model.add(tf.keras.layers.Dense(
    2, use_bias=False,
    activation='softmax'))
model.compile(loss='mse')
```

How many trainable weights does your model have? (The arithmetic below is correct.)

A. 501*256+257*128+2 = 161154

B. 500*256+256*128+128*2 = 161024

C. 501*256+257*128+128*2=161408

D. 500*256*0 25+256*128*0 25+128*2 = 40448

ANSWER: C