# Google

Exam Questions Professional-Machine-Learning-Engineer

Google Professional Machine Learning Engineer

NEW QUESTION 1
You are designing an architecture with a serveress ML system to enrich customer support tickets with informative metadata before they are routed to a support agent. You need a set of models to predict ticket priority, predict ticket resolution time, and perform sentiment analysis to help agents make strategic decisions when they process support requests. Tickets are not expected to have any domain-specific terms or jargon.
The proposed architecture has the following flow:
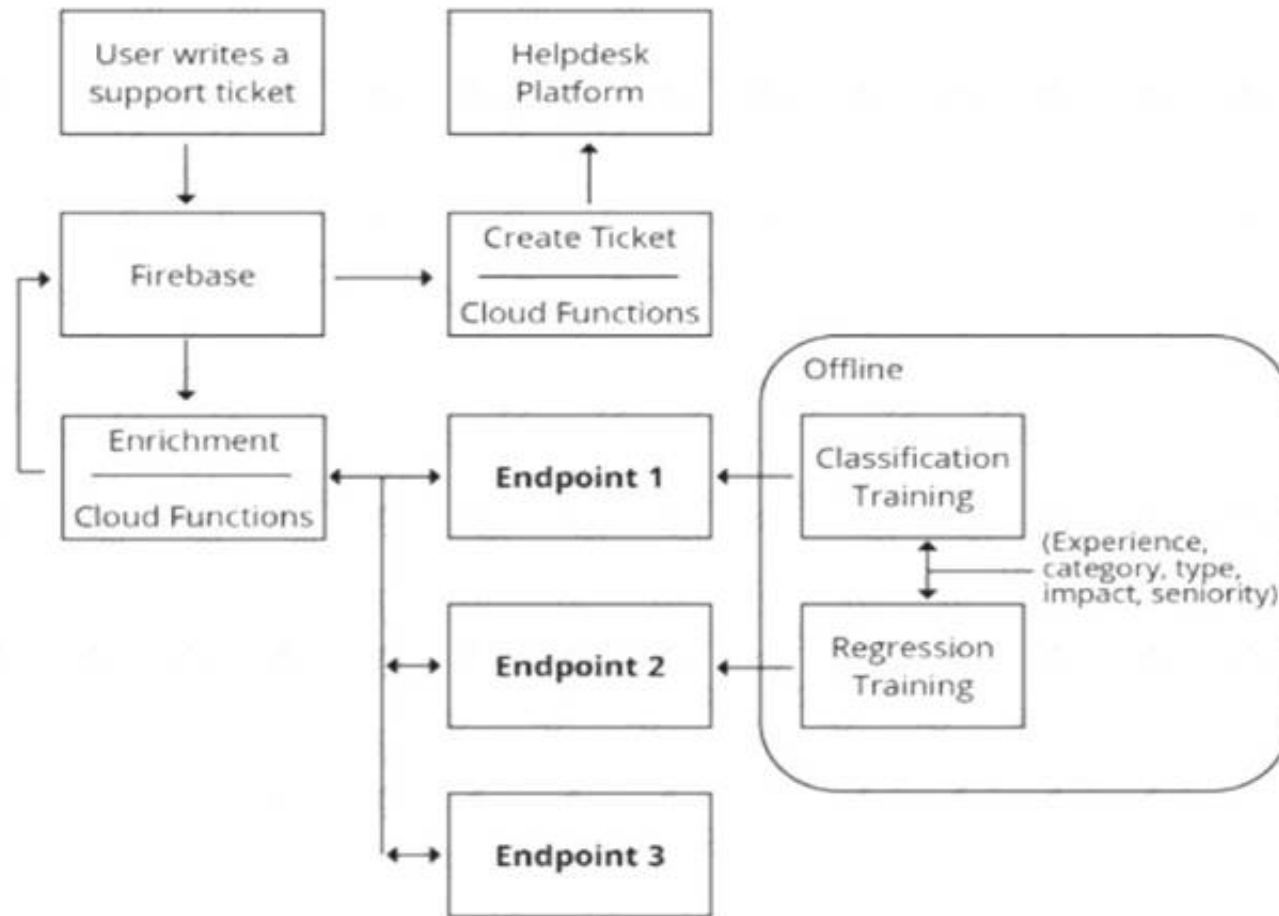
Which endpoints should the Enrichment Cloud Functions call?

A. 1 = AI Platform, 2 = AI Platform, 3 = AutoML Vision
B. 1 = AI Platform, 2 = AI Platform, 3 = AutoML Natural Language
C. 1 = AI Platform, 2 = AI Platform, 3 = Cloud Natural Language API
D. 1 = cloud Natural Language API, 2 = AI Platform, 3 = Cloud Vision API

Answer: B

NEW QUESTION 2
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: B

NEW QUESTION 3
You recently designed and built a custom neural network that uses critical dependencies specific to your organization's framework. You need to train the model using a managed training service on Google Cloud. However, the ML framework and related dependencies are not supported by AI Platform Training. Also, both your model and your data are too large to fit in memory on a single machine. Your ML framework of choice uses the scheduler, workers, and servers distribution structure. What should you do?

A. Use a built-in model available on AI Platform Training
B. Build your custom container to run jobs on AI Platform Training
C. Build your custom containers to run distributed training jobs on AI Platform Training
D. Reconfigure your code to a ML framework with dependencies that are supported by AI Platform Training

Answer: C

NEW QUESTION 4
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 model
B. AutoML Vision Edge mobile-versatile-1 model
C. AutoML Vision Edge mobile-low-latency-1 model

D. AutoML Vision Edge mobile-high-accuracy-1 model

Answer: A

NEW QUESTION 5
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: D

NEW QUESTION 6
You developed an ML model with AI Platform, and you want to move it to production. You serve a few thousand queries per second and are experiencing latency issues. Incoming requests are served by a load balancer that distributes them across multiple Kubeflow CPU-only pods running on Google Kubernetes Engine (GKE). Your goal is to improve the serving latency without changing the underlying infrastructure. What should you do?

A. Significantly increase the max_batch_size TensorFlow Serving parameter
B. Switch to the tensorflow-model-server-universal version of TensorFlow Serving
C. Significantly increase the max_enqueued_batches TensorFlow Serving parameter
D. Recompile TensorFlow Serving using the source to support CPU-specific optimizations Instruct GKE to choose an appropriate baseline minimum CPU platform for serving nodes

Answer: A

NEW QUESTION 7
You are building a real-time prediction engine that streams files which may contain Personally Identifiable Information (Pll) to Google Cloud. You want to use the Cloud Data Loss Prevention (DLP) API to scan the files. How should you ensure that the Pll is not accessible by unauthorized individuals?

A. Stream all files to Google CloudT 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.
E. Periodically conduct a bulk scan of that bucket using the DLP API, and move the data to either the Sensitive or Non-Sensitive bucket

Answer: A

NEW QUESTION 8
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?

A. Differential privacy
B. Federated learning
C. MD5 to encrypt data
D. Data Loss Prevention API

Answer: B

NEW QUESTION 9
You recently joined a machine learning team that will soon release a new project. As a lead on the project, you are asked to determine the production readiness of the ML components. The team has already tested features and data, model development, and infrastructure. Which additional readiness check should you recommend to the team?

A. Ensure that training is reproducible
B. Ensure that all hyperparameters are tuned
C. Ensure that model performance is monitored
D. Ensure that feature expectations are captured in the schema

Answer: B

NEW QUESTION 10
You are developing ML models with AI Platform for image segmentation on CT scans. You frequently update your model architectures based on the newest available research papers, and have to rerun training on the same dataset to benchmark their performance. You want to minimize computation costs and manual intervention while having version control for your code. What should you do?

A. Use Cloud Functions to identify changes to your code in Cloud Storage and trigger a retraining job
B. Use the gcloud command-line tool to submit training jobs on AI Platform when you update your code
C. Use Cloud Build linked with Cloud Source Repositories to trigger retraining when new code is pushed to the repository
D. Create an automated workflow in Cloud Composer that runs daily and looks for changes in code in Cloud Storage using a sensor.

Answer: A

NEW QUESTION 10
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: DE

NEW QUESTION 12
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.

Answer: D

NEW QUESTION 16
You are developing models to classify customer support emails. You created models with TensorFlow Estimators using small datasets on your on-premises system, but you now need to train the models using large datasets to ensure high performance. You will port your models to Google Cloud and want to minimize code refactoring and infrastructure overhead for easier migration from on-prem to cloud. What should you do?

A. Use AI Platform for distributed training
B. Create a cluster on Dataproc for training
C. Create a Managed Instance Group with autoscaling
D. Use Kubeflow Pipelines to train on a Google Kubernetes Engine cluster.

Answer: D

NEW QUESTION 20
You work for a social media company. You need to detect whether posted images contain cars. Each training example is a member of exactly one class. You have trained an object detection neural network and deployed the model version to AI Platform Prediction for evaluation. Before deployment, you created an evaluation job and attached it to the AI Platform Prediction model version. You notice that the precision is lower than your business requirements allow. How should you adjust the model's final layer softmax threshold to increase precision?

A. Increase the recall
B. Decrease the recall.
C. Increase the number of false positives
D. Decrease the number of false negatives

Answer: D

NEW QUESTION 23
You are training a TensorFlow model on a structured data set with 100 billion records stored in several CSV files. You need to improve the input/output execution performance. What should you do?

A. Load the data into BigQuery and read the data from BigQuery.
B. Load the data into Cloud Bigtable, and read the data from Bigtable
C. Convert the CSV files into shards of TFRecords, and store the data in Cloud Storage
D. Convert the CSV files into shards of TFRecords, and store the data in the Hadoop Distributed File System (HDFS)

Answer: B

NEW QUESTION 26
You are an ML engineer at a regulated insurance company. You are asked to develop an insurance approval model that accepts or rejects insurance applications from potential customers. What factors should you consider before building the model?

A. Redaction, reproducibility, and explainability
B. Traceability, reproducibility, and explainability
C. Federated learning, reproducibility, and explainability
D. Differential privacy federated learning, and explainability

Answer: B

NEW QUESTION 31
You need to build classification workflows over several structured datasets currently stored in BigQuery.
Because you will be performing the classification several times, you want to complete the following steps without writing code: exploratory data analysis, feature selection, model building, training, and hyperparameter tuning and serving. What should you do?

A. Configure AutoML Tables to perform the classification task
B. Run a BigQuery ML task to perform logistic regression for the classification
C. Use AI Platform Notebooks to run the classification model with pandas library
D. Use AI Platform to run the classification model job configured for hyperparameter tuning

Answer: C

NEW QUESTION 33
You have trained a model on a dataset that required computationally expensive preprocessing operations. You need to execute the same preprocessing at prediction time. You deployed the model on AI Platform for high-throughput online prediction. Which architecture should you use?

A. • Validate the accuracy of the model that you trained on preprocessed data• Create a new model that uses the raw data and is available in real time• Deploy the new model onto AI Platform for online prediction
B. • Send incoming prediction requests to a Pub/Sub topic• Transform the incoming data using a Dataflow job• Submit a prediction request to AI Platform using the transformed data• Write the predictions to an outbound Pub/Sub queue
C. • Stream incoming prediction request data into Cloud Spanner• Create a view to abstract your preprocessing logic.• Query the view every second for new records• Submit a prediction request to AI Platform using the transformed data• Write the predictions to an outbound Pub/Sub queue.
D. • Send incoming prediction requests to a Pub/Sub topic• Set up a Cloud Function that is triggered when messages are published to the Pub/Sub topic.• Implement your preprocessing logic in the Cloud Function• Submit a prediction request to AI Platform using the transformed data• Write the predictions to an outbound Pub/Sub queue

Answer: D

NEW QUESTION 36
......

Thank You for Trying Our Product

We offer two products:

1st - We have Practice Tests Software with Actual Exam Questions

2nd - Questons and Answers in PDF Format

Professional-Machine-Learning-Engineer Practice Exam Features:

* Professional-Machine-Learning-Engineer Questions and Answers Updated Frequently

* Professional-Machine-Learning-Engineer Practice Questions Verified by Expert Senior Certified Staff

* Professional-Machine-Learning-Engineer Most Realistic Questions that Guarantee you a Pass on Your FirstTry

* Professional-Machine-Learning-Engineer Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year

100% Actual & Verified — Instant Download, Please Click
[Order The Professional-Machine-Learning-Engineer Practice Test Here](https://www.surepassexam.com/Professional-Machine-Learning-Engineer-exam-dumps.html)