

Google

GCP-PMLE

Google Professional Machine Learning Engineer

For More Information – Visit link below:

<https://www.examsempire.com/>

Product Version

1. Up to Date products, reliable and verified.
2. Questions and Answers in PDF Format.



<https://examsempire.com/>

Latest Version: 6.0

Question: 1

You need to build an object detection model for a small startup company to identify if and where the company's logo appears in an image. You were given a large repository of images, some with logos and some without.

These images are not yet labelled. You need to label these pictures, and then train and deploy the model. What should you do?

- A. Use Google Cloud's Data Labelling Service to label your data. Use AutoML Object Detection to train and deploy the model.
- B. Use Vision API to detect and identify logos in pictures and use it as a label. Use AI Platform to build and train a convolutional neural network.
- C. Create two folders: one where the logo appears and one where it doesn't. Manually place images in each folder. Use AI Platform to build and train a convolutional neural network.
- D. Create two folders: one where the logo appears and one where it doesn't. Manually place images in each folder. Use AI Platform to build and train a real time object detection model.

Answer: A

Question: 2

You work for a gaming company that develops and manages a popular massively multiplayer online (MMO) game.

The game's environment is open-ended, and a large number of positions and moves can be taken by a player. Your team has developed an ML model with TensorFlow that predicts the next move of each player.

Edge deployment is not possible, but low-latency serving is required. How should you configure the deployment?

- A. Use a Cloud TPU to optimize model training speed.
- B. Use AI Platform Prediction with a NVIDIA GPU to make real-time predictions.
- C. Use AI Platform Prediction with a high-CPU machine type to get a batch prediction for the players.
- D. Use AI Platform Prediction with a high-memory machine type to get a batch prediction for the players.

Answer: B

Question: 3

You work for a textile manufacturer and have been asked to build a model to detect and classify fabric defects.

You trained a machine learning model with high recall based on high resolution images taken at the end of the production line. You want quality control inspectors to gain trust in your model.

Which technique should you use to understand the rationale of your classifier?

- A. Use K-fold cross validation to understand how the model performs on different test datasets.
- B. Use the Integrated Gradients method to efficiently compute feature attributions for each predicted image.
- C. Use PCA (Principal Component Analysis) to reduce the original feature set to a smaller set of easily understood features.
- D. Use k-means clustering to group similar images together, and calculate the Davies-Bouldin index to evaluate the separation between clusters.

Answer: B

Question: 4

You work for a large financial institution that is planning to use Dialogflow to create a chatbot for the company's mobile app.

You have reviewed old chat logs and tagged each conversation for intent based on each customer's stated intention for contacting customer service.

About 70% of customer inquiries are simple requests that are solved within 10 intents. The remaining 30% of inquiries require much longer and more complicated requests.

Which intents should you automate first?

- A. Automate a blend of the shortest and longest intents to be representative of all intents.
- B. Automate the more complicated requests first because those require more of the agents' time.
- C. Automate the 10 intents that cover 70% of the requests so that live agents can handle the more complicated requests.
- D. Automate intents in places where common words such as "payment" only appear once to avoid confusing the software.

Answer: C

Question: 5

You work for a manufacturing company that owns a high-value machine which has several machine settings and multiple sensors.

A history of the machine's hourly sensor readings and known failure event data are stored in BigQuery. You need to predict if the machine will fail within the next 3 days in order to schedule maintenance before the machine fails.

Which data preparation and model training steps should you take?

- A. Data preparation: Daily max value feature engineering; Model training: AutoML classification with BQML
- B. Data preparation: Daily min value feature engineering; Model training: Logistic regression with BQML and AUTO_CLASS_WEIGHTS set to True
- C. Data preparation: Rolling average feature engineering; Model training: Logistic regression with BQML and AUTO_CLASS_WEIGHTS set to False
- D. Data preparation: Rolling average feature engineering; Model training: Logistic regression with BQML and AUTO_CLASS_WEIGHTS set to True

Answer: D

Question: 6

You work on a team where the process for deploying a model into production starts with data scientists training different versions of models in a Kubeflow pipeline.

The workflow then stores the new model artifact into the corresponding Cloud Storage bucket. You need to build the next steps of the pipeline after the submitted model is ready to be tested and deployed in production on AI Platform.

How should you configure the architecture before deploying the model to production?

- A. Deploy model in test environment -> Evaluate and test model -> Create a new AI Platform model version
- B. Validate model -> Deploy model in test environment -> Create a new AI Platform model version
- C. Create a new AI Platform model version -> Evaluate and test model -> Deploy model in test environment
- D. Create a new AI Platform model version -> Deploy model in test environment -> Validate model

Answer: A

Question: 7

You are an ML engineer at a media company. You want to use machine learning to analyze video content, identify objects, and alert users if there is inappropriate content.

Which Google Cloud products should you use to build this project?

- A. Pub/Sub, Cloud Function, Cloud Vision API
- B. Pub/Sub, Cloud IoT, Dataflow, Cloud Vision API, Cloud Logging
- C. Pub/Sub, Cloud Function, Video Intelligence API, Cloud Logging
- D. Pub/Sub, Cloud Function, AutoML Video Intelligence, Cloud Logging

Answer: C

Question: 8

Your team is using a TensorFlow Inception-v3 CNN model pretrained on ImageNet for an image classification prediction challenge on 10,000 images. You will use AI Platform to perform the model training.

What TensorFlow distribution strategy and AI Platform training job configuration should you use to train the model and optimize for wall-clock time?

- A. Default Strategy; Custom tier with a single master node and four v100 GPUs.
- B. One Device Strategy; Custom tier with a single master node and four v100 GPUs.
- C. One Device Strategy; Custom tier with a single master node and eight v100 GPUs.
- D. MirroredStrategy; Custom tier with a single master node and four v100 GPUs.

Answer: D

Question: 9

You work for a large retailer. You want to use ML to forecast future sales leveraging 10 years of historical sales data.

The historical data is stored in Cloud Storage in Avro format. You want to rapidly experiment with all the available data.

How should you build and train your model for the sales forecast?

- A. Load data into BigQuery and use the ARIMA model type on BigQuery ML.
- B. Convert the data into CSV format and create a regression model on AutoML Tables.
- C. Convert the data into TFRecords and create an RNN model on TensorFlow on AI Platform Notebooks.
- D. Convert and refactor the data into CSV format and use the built-in XGBoost algorithm on AI Platform Training.

Answer: A

Question: 10

You need to write a generic test to verify whether Dense Neural Network (DNN) models automatically released by your team have a sufficient number of parameters to learn the task for which they were built.

What should you do?

- A. Train the model for a few iterations, and check for NaN values.
- B. Train the model for a few iterations, and verify that the loss is constant.
- C. Train a simple linear model, and determine if the DNN model outperforms it.
- D. Train the model with no regularization, and verify that the loss function is close to zero.

Answer: D

Thank You for Trying Our Product

Special 16 USD Discount Coupon: **NSZUBG3X**

Email: support@examsempire.com

**Check our Customer Testimonials and ratings
available on every product page.**

Visit our website.

<https://examsempire.com/>