

PROFESSIONAL-MACHINE-LEARNING-ENGINEER^{Q&As}

Professional Machine Learning Engineer

Pass Google PROFESSIONAL-MACHINE-LEARNING-ENGINEER Exam with 100% Guarantee

Free Download Real Questions & Answers **PDF** and **VCE** file from:

https://www.pass4itsure.com/professional-machine-learning-engineer.html

100% Passing Guarantee 100% Money Back Assurance

Following Questions and Answers are all new published by Google Official Exam Center https://www.pass4itsure.com/professional-machine-learning-engineer.html 2024 Latest pass4itsure PROFESSIONAL-MACHINE-LEARNING-ENGINEER PDF and VCE dumps Download

- Instant Download After Purchase
- 100% Money Back Guarantee
- 🙆 365 Days Free Update

VCE & PDF

Pass4itSure.com

800,000+ Satisfied Customers





QUESTION 1

You are an ML engineer at a manufacturing company. You need to build a model that identifies defects in products based on images of the product taken at the end of the assembly line. You want your model to preprocess the images with lower computation to quickly extract features of defects in products. Which approach should you use to build the model?

- A. Reinforcement learning
- B. Recommender system
- C. Recurrent Neural Networks (RNN)
- D. Convolutional Neural Networks (CNN)

Correct Answer: D

https://developers.google.com/machine-learning/practica/image-classification/convolutional-neural-networks

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?

- 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.

Correct Answer: D

Cite from Google Pag: to construct a Dataset from data in memory, use tf.data.Dataset.from_tensors() or tf.data.Dataset.from_tensor_slices(). When input data is stored in a file (not in memory), the recommended TFRecord format, you can use tf.data.TFRecordDataset().

tf.data.Dataset is for data in memory. tf.data.TFRecordDataset is for data in non-memory storage.

QUESTION 3

You recently used BigQuery ML to train an AutoML regression model. You shared results with your team and received positive feedback. You need to deploy your model for online prediction as quickly as possible. What should you do?

A. Retrain the model by using BigQuery ML, and specify Vertex AI as the model registry. Deploy the model from Vertex AI Model Registry to a Vertex AI endpoint,

B. Retrain the model by using Vertex AI Deploy the model from Vertex AI Model. Registry to a Vertex AI endpoint.



C. Alter the model by using BigQuery ML, and specify Vertex AI as the model registry. Deploy the model from Vertex AI Model Registry to a Vertex AI endpoint.

D. Export the model from BigQuery ML to Cloud Storage. Import the model into Vertex Al Model Registry. Deploy the model to a Vertex Al endpoint.

Correct Answer: C

QUESTION 4

You have trained a text classification model in TensorFlow using AI Platform. You want to use the trained model for batch predictions on text data stored in BigQuery while minimizing computational overhead. What should you do?

A. Export the model to BigQuery ML.

- B. Deploy and version the model on AI Platform.
- C. Use Dataflow with the SavedModel to read the data from BigQuery.
- D. Submit a batch prediction job on AI Platform that points to the model location in Cloud Storage.

Correct Answer: D

https://cloud.google.com/vertex-ai/docs/tabular-data/classification-regression/get-batch-predictions

QUESTION 5

You are an ML engineer at a manufacturing company. You are creating a classification model for a predictive maintenance use case. You need to predict whether a crucial machine will fail in the next three days so that the repair crew has enough time to fix the machine before it breaks. Regular maintenance of the machine is relatively inexpensive, but a failure would be very costly. You have trained several binary classifiers to predict whether the machine will fail, where a prediction of 1 means that the ML model predicts a failure.

You are now evaluating each model on an evaluation dataset. You want to choose a model that prioritizes detection while ensuring that more than 50% of the maintenance jobs triggered by your model address an imminent machine failure. Which model should you choose?

A. The model with the highest area under the receiver operating characteristic curve (AUC ROC) and precision greater than 0.5

B. The model with the lowest root mean squared error (RMSE) and recall greater than 0.5.

- C. The model with the highest recall where precision is greater than 0.5.
- D. The model with the highest precision where recall is greater than 0.5.

Correct Answer: C

Priority is to detect(Pointing to Recall) and correctly detect (more that 50% - pointing to Precision)

Latest PROFESSIONAL-MA PROFESSIONAL-MACHIN PROFESSIONAL-MACHIN



https://www.pass4itsure.com/professional-machine-learning-engineer.html 2024 Latest pass4itsure PROFESSIONAL-MACHINE-LEARNING-ENGINEER PDF and VCE dumps Download

CHINE-LEARNING-ENGINEER Dumps E-LEARNING-ENGINEER E-LEARNING-ENGINEER Exam Questions Braindumps