



BrainDumps
Collection

Google

PROFESSIONAL-MACHINE-LEARNING-ENGINEER Exam

Google Professional Machine Learning Engineer

Thank you for Downloading PROFESSIONAL-MACHINE-LEARNING-ENGINEER exam PDF Demo

You can also try our PROFESSIONAL-MACHINE-LEARNING-ENGINEER practice exam software

Download Free Demo

<https://www.braindumpscollection.com/PROFESSIONAL-MACHINE-LEARNING-ENGINEER.html>

DEMO
VERSION

(LIMITED CONTENT)

Questions
& Answers

Version: 10.3

Question: 1

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: A

Explanation:

Batch prediction is the process of using an ML model to make predictions on a large set of data points. Batch prediction is suitable for scenarios where the predictions are not time-sensitive and can be done in batches, such as digitizing scanned customer forms at the end of each day. Batch prediction can also handle large volumes of data and scale up or down the resources as needed. AI Platform provides a batch prediction service that allows users to submit a job with their TensorFlow model and input data stored in Cloud Storage, and receive the output predictions in Cloud Storage as well. This service requires minimal manual intervention and can be automated with Cloud Scheduler or Cloud Functions. Therefore, using the batch prediction functionality of AI Platform is the best option for this use case.

Reference:

[Batch prediction overview](#)

[Using batch prediction](#)

Question: 2

You work for a global footwear retailer and need to predict when an item will be out of stock based on historical inventory data.

a. Customer behavior is highly dynamic since footwear demand is influenced by many different factors. You want to serve models that are trained on all available data, but track your performance on specific subsets of data before pushing to production. What is the most streamlined and reliable way to perform this validation?

- A. Use the TFX ModelValidator tools to specify performance metrics for production readiness
- B. Use k-fold cross-validation as a validation strategy to ensure that your model is ready for production.
- C. Use the last relevant week of data as a validation set to ensure that your model is performing

accurately on current data

D. Use the entire dataset and treat the area under the receiver operating characteristics curve (AUC ROC) as the main metric.

Answer: A

Explanation:

TFX ModelValidator is a tool that allows you to compare new models against a baseline model and evaluate their performance on different metrics and data slices¹. You can use this tool to validate your models before deploying them to production and ensure that they meet your expectations and requirements.

k-fold cross-validation is a technique that splits the data into k subsets and trains the model on k-1 subsets while testing it on the remaining subset. This is repeated k times and the average performance is reported². This technique is useful for estimating the generalization error of a model, but it does not account for the dynamic nature of customer behavior or the potential changes in data distribution over time.

Using the last relevant week of data as a validation set is a simple way to check the model's performance on recent data, but it may not be representative of the entire data or capture the long-term trends and patterns. It also does not allow you to compare the model with a baseline or evaluate it on different data slices.

Using the entire dataset and treating the AUC ROC as the main metric is not a good practice because it does not leave any data for validation or testing. It also assumes that the AUC ROC is the only metric that matters, which may not be true for your business problem. You may want to consider other metrics such as precision, recall, or revenue.

Question: 3

You work on a growing team of more than 50 data scientists who all use AI Platform. You are designing a strategy to organize your jobs, models, and versions in a clean and scalable way. Which strategy should you choose?

A. Set up restrictive IAM permissions on the AI Platform notebooks so that only a single user or group can access a given instance.

B. Separate each data scientist's work into a different project to ensure that the jobs, models, and versions created by each data scientist are accessible only to that user.

C. Use labels to organize resources into descriptive categories. Apply a label to each created resource so that users can filter the results by label when viewing or monitoring the resources

D. Set up a BigQuery sink for Cloud Logging logs that is appropriately filtered to capture information about AI Platform resource usage In BigQuery create a SQL view that maps users to the resources they are using.

Answer: C

Explanation:

Labels are key-value pairs that can be attached to any AI Platform resource, such as jobs, models, versions, or endpoints¹. Labels can help you organize your resources into descriptive categories, such as project, team, environment, or purpose. You can use labels to filter the results when you list or

[monitor your resources, or to group them for billing or quota purposes²](#). Using labels is a simple and scalable way to manage your AI Platform resources without creating unnecessary complexity or overhead. Therefore, using labels to organize resources is the best strategy for this use case.

Reference:

[Using labels](#)

[Filtering and grouping by labels](#)

Question: 4

During batch training of a neural network, you notice that there is an oscillation in the loss. How should you adjust your model to ensure that it converges?

- A. Increase the size of the training batch
- B. Decrease the size of the training batch
- C. Increase the learning rate hyperparameter
- D. Decrease the learning rate hyperparameter

Answer: D

Explanation:

Oscillation in the loss during batch training of a neural network means that the model is overshooting the optimal point of the loss function and bouncing back and forth. This can prevent the model from converging to the minimum loss value. One of the main reasons for this phenomenon is that the learning rate hyperparameter, which controls the size of the steps that the model takes along the gradient, is too high. Therefore, decreasing the learning rate hyperparameter can help the model take smaller and more precise steps and avoid oscillation. [This is a common technique to improve the stability and performance of neural network training¹²](#).

Reference:

[Interpreting Loss Curves](#)

[Is learning rate the only reason for training loss oscillation after few epochs?](#)

Question: 5

You are building a linear model with over 100 input features, all with values between -1 and 1. You suspect that many features are non-informative. You want to remove the non-informative features from your model while keeping the informative ones in their original form. Which technique should you use?

- A. Use Principal Component Analysis to eliminate the least informative features.
- B. Use L1 regularization to reduce the coefficients of uninformative features to 0.
- C. After building your model, use Shapley values to determine which features are the most informative.
- D. Use an iterative dropout technique to identify which features do not degrade the model when removed.

Answer: B

Explanation:

L1 regularization, also known as Lasso regularization, adds the sum of the absolute values of the model's coefficients to the loss function¹. It encourages sparsity in the model by shrinking some coefficients to precisely zero². This way, L1 regularization can perform feature selection and remove the non-informative features from the model while keeping the informative ones in their original form. Therefore, using L1 regularization is the best technique for this use case.

Reference:

Regularization in Machine Learning - GeeksforGeeks

Regularization in Machine Learning (with Code Examples) - Dataquest

L1 And L2 Regularization Explained & Practical How To Examples

L1 and L2 as Regularization for a Linear Model

Thank You for trying PROFESSIONAL-MACHINE- LEARNING-ENGINEER PDF Demo

To try our PROFESSIONAL-MACHINE-LEARNING-ENGINEER
practice exam software visit link below

[https://www.braindumpscollection.com/PROFESSIONAL-MACHINE-
LEARNING-ENGINEER.html](https://www.braindumpscollection.com/PROFESSIONAL-MACHINE-LEARNING-ENGINEER.html)

Start Your PROFESSIONAL- MACHINE-LEARNING- ENGINEER Preparation

Use Coupon “20OFF” for extra 20% discount on the purchase of
Practice Test Software. Test your PROFESSIONAL-MACHINE-
LEARNING-ENGINEER preparation with actual exam questions.