



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



VCE & PDF

Pass4itSure.com

<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
- ⚙️ **800,000+** Satisfied Customers



**QUESTION 1**

Your data science team needs to rapidly experiment with various features, model architectures, and hyperparameters. They need to track the accuracy metrics for various experiments and use an API to query the metrics over time. What should they use to track and report their experiments while minimizing manual effort?

- A. Use Vertex AI Pipelines to execute the experiments. Query the results stored in MetadataStore using the Vertex AI API.
- B. Use Vertex AI Training to execute the experiments. Write the accuracy metrics to BigQuery, and query the results using the BigQuery API.
- C. Use Vertex AI Training to execute the experiments. Write the accuracy metrics to Cloud Monitoring, and query the results using the Monitoring API.
- D. Use Vertex AI Workbench user-managed notebooks to execute the experiments. Collect the results in a shared Google Sheets file, and query the results using the Google Sheets API.

Correct Answer: A

The Vertex AI Pipelines provide a powerful tool for automating machine learning workflows, including data preparation, training, and deployment. MetadataStore can be used to track the performance of different models by logging accuracy metrics and other important information. The Vertex AI API can then be used to query the metadata store and retrieve the results of different experiments.

QUESTION 2

You have deployed multiple versions of an image classification model on AI Platform. You want to monitor the performance of the model versions over time. How should you perform this comparison?

- A. Compare the loss performance for each model on a held-out dataset.
- B. Compare the loss performance for each model on the validation data.
- C. Compare the receiver operating characteristic (ROC) curve for each model using the What-If Tool.
- D. Compare the mean average precision across the models using the Continuous Evaluation feature.

Correct Answer: D

Compare the mean average precision across the models using the Continuous Evaluation feature

<https://cloud.google.com/vertex-ai/docs/evaluation/introduction>

Vertex AI provides model evaluation metrics, such as precision and recall, to help you determine the performance of your models...

Vertex AI supports evaluation of the following model types:

AuPRC: The area under the precision-recall (PR) curve, also referred to as average precision. This value ranges from zero to one, where a higher value indicates a higher-quality model.

**QUESTION 3**

You are an ML engineer at a large grocery retailer with stores in multiple regions. You have been asked to create an inventory prediction model. Your model's features include region, location, historical demand, and seasonal popularity. You want the algorithm to learn from new inventory data on a daily basis. Which algorithms should you use to build the model?

- A. Classification
- B. Reinforcement Learning
- C. Recurrent Neural Networks (RNN)
- D. Convolutional Neural Networks (CNN)

Correct Answer: C

QUESTION 4

You work for a credit card company and have been asked to create a custom fraud detection model based on historical data using AutoML Tables. You need to prioritize detection of fraudulent transactions while minimizing false positives. Which optimization objective should you use when training the model?

- A. An optimization objective that minimizes Log loss
- B. An optimization objective that maximizes the Precision at a Recall value of 0.50
- C. An optimization objective that maximizes the area under the precision-recall curve (AUC PR) value
- D. An optimization objective that maximizes the area under the receiver operating characteristic curve (AUC ROC) value

Correct Answer: C

<https://stats.stackexchange.com/questions/262616/roc-vs-precision-recall-curves-on-imbalanced-dataset>

<https://neptune.ai/blog/f1-score-accuracy-roc-auc-pr-auc>

QUESTION 5

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 Edge mobile-high-accuracy-1 model
- B. AutoML Vision Edge mobile-low-latency-1 model
- C. AutoML Vision model
- D. AutoML Vision Edge mobile-versatile-1 model



Correct Answer: B

Hence faster defect detection is a priority, AutoML Vision Edge mobile-low-latency-1 model should be the choice. This model is designed to run efficiently on mobile devices and prioritize low latency, which means that it can provide fast defect detection without requiring a connection to the cloud. <https://cloud.google.com/vision/automl/docs/train-edge>

[Latest PROFESSIONAL-MACHINE-LEARNING-ENGINEER Dumps](#)

[PROFESSIONAL-MACHINE-LEARNING-ENGINEER VCE Dumps](#)

[PROFESSIONAL-MACHINE-LEARNING-ENGINEER Exam Questions](#)