



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.geekcert.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

GeekCert.com

<https://www.geekcert.com/professional-machine-learning-engineer.html>

2024 Latest geekcert 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

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 mode?

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

Correct Answer: B

With federated learning, all the data is collected, and the model is trained with algorithms across multiple decentralized edge devices such as cell phones or websites, without exchanging them.

QUESTION 2

You are building a model to predict daily temperatures. You split the data randomly and then transformed the training and test datasets. Temperature data for model training is uploaded hourly. During testing, your model performed with 97% accuracy; however, after deploying to production, the model's accuracy dropped to 66%. How can you make your production model more accurate?

- A. Normalize the data for the training, and test datasets as two separate steps.
- B. Split the training and test data based on time rather than a random split to avoid leakage.
- C. Add more data to your test set to ensure that you have a fair distribution and sample for testing.
- D. Apply data transformations before splitting, and cross-validate to make sure that the transformations are applied to both the training and test sets.

Correct Answer: B

QUESTION 3

Your company manages an application that aggregates news articles from many different online sources and sends them to users. You need to build a recommendation model that will suggest articles to readers that are similar to the articles they are currently reading. Which approach should you use?

- A. Create a collaborative filtering system that recommends articles to a user based on the user's past behavior.
- B. Encode all articles into vectors using word2vec, and build a model that returns articles based on vector similarity.
- C. Build a logistic regression model for each user that predicts whether an article should be recommended to a user.
- D. Manually label a few hundred articles, and then train an SVM classifier based on the manually classified articles that



categorizes additional articles into their respective categories.

Correct Answer: B

<https://cloud.google.com/blog/topics/developers-practitioners/meet-ais-multitool-vector-embeddings>

QUESTION 4

Your data science team is training a PyTorch model for image classification based on a pre-trained ResNet model. You need to perform hyperparameter tuning to optimize for several parameters. What should you do?

- A. Convert the model to a Keras model, and run a Keras Tuner job.
- B. Run a hyperparameter tuning job on AI Platform using custom containers.
- C. Create a Kuberflow Pipelines instance, and run a hyperparameter tuning job on Katib.
- D. Convert the model to a TensorFlow model, and run a hyperparameter tuning job on AI Platform.

Correct Answer: B

<https://cloud.google.com/blog/topics/developers-practitioners/pytorch-google-cloud-how-train-and-tune-pytorch-models-vertex-ai> <https://cloud.google.com/blog/topics/developers-practitioners/pytorch-google-cloud-how-deploy-pytorch-modelsvertex-ai>

QUESTION 5

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.

Correct Answer: A

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

[PROFESSIONAL-MACHINE-LEARNING-ENGINEER Study Guide](#)

[PROFESSIONAL-MACHINE-LEARNING-ENGINEER Braindumps](#)