


Google Professional-Machine-Learning-Engineer: Google Professional Machine Learning Engineer braindumps PDF & Testking echter Test



P.S. Kostenlose 2026 Google Professional-Machine-Learning-Engineer Prüfungsfragen sind auf Google Drive freigegeben von Fast2test verfügbar: <https://drive.google.com/open?id=1IVKJaQ4utreAQ0blxqc5DK32WOpb0SbF>

Wir versprechen, dass Sie die Prüfung zum ersten Mal mit unseren Schulungsunterlagen zur Google Professional-Machine-Learning-Engineer Zertifizierungsprüfung bestehen können. Sonst erstatten wir Ihnen die gesamte Summe zurück.

Die Erlangung der Google Professional Machine Learning Engineer-Zertifizierung zeigt die Fähigkeit des Kandidaten, Machine-Learning-Modelle unter Verwendung von Google Cloud-Technologien zu entwerfen und umzusetzen und kann zu Karrierechancen und erhöhten Jobaussichten führen. Es ist eine hoch angesehene Zertifizierung im Bereich des maschinellen Lernens und wird von Industrieexperten weltweit anerkannt.

>> **Professional-Machine-Learning-Engineer Online Prüfung** <<

Professional-Machine-Learning-Engineer Übungstest: Google Professional Machine Learning Engineer & Professional-Machine-Learning-Engineer Braindumps Prüfung

Sorgen Sie noch um die Prüfungsunterlagen der Google Professional-Machine-Learning-Engineer? Jetzt brauchen Sie keine Sorgen! Weil uns zu finden bedeutet, dass Sie schon die Schlüssel zur Prüfungszertifizierung der Google Professional-Machine-Learning-Engineer gefunden haben. Wir Fast2test beschäftigen uns seit Jahren mit der Entwicklung der Software der IT-Zertifizierungsprüfung.

Jetzt genießen wir einen guten Ruf weltweit. Wir bieten Ihnen die effektivsten Hilfe bei der Vorbereitung der Google Professional-Machine-Learning-Engineer.

Google Professional Machine Learning Engineer Professional-Machine-Learning-Engineer Prüfungsfragen mit Lösungen (Q81-Q86):

81. Frage

You need to build an ML model for a social media application to predict whether a user's submitted profile photo meets the requirements. The application will inform the user if the picture meets the requirements. How should you build a model to ensure that the application does not falsely accept a non-compliant picture?

- A. Use AutoML to optimize the model's F1 score in order to balance the accuracy of false positives and false negatives.
- B. Use Vertex AI Workbench user-managed notebooks to build a custom model that has three times as many examples of pictures that do not meet the profile photo requirements.
- C. Use Vertex AI Workbench user-managed notebooks to build a custom model that has three times as many examples of pictures that meet the profile photo requirements.
- **D. Use AutoML to optimize the model's recall in order to minimize false negatives.**

Antwort: D

Begründung:

Recall is the ratio of true positives to the sum of true positives and false negatives. It measures how well the model can identify all the relevant cases. In this scenario, the relevant cases are the pictures that do not meet the profile photo requirements. Therefore, minimizing false negatives means minimizing the cases where the model incorrectly predicts that a non-compliant picture meets the requirements. By using AutoML to optimize the model's recall, the model will be more likely to reject a non-compliant picture and inform the user accordingly. References:

* [AutoML Vision] is a service that allows you to train custom ML models for image classification and object detection tasks. You can use AutoML to optimize your model for different metrics, such as recall, precision, or F1 score.

* [Recall] is one of the evaluation metrics for ML models. It is defined as $TP / (TP + FN)$, where TP is the number of true positives and FN is the number of false negatives. Recall measures how well the model can identify all the relevant cases. A high recall means that the model has a low rate of false negatives.

82. Frage

You work for a food product company. Your company's historical sales data is stored in BigQuery. You need to use Vertex AI's custom training service to train multiple TensorFlow models that read the data from BigQuery and predict future sales. You plan to implement a data preprocessing algorithm that performs min-max scaling and bucketing on a large number of features before you start experimenting with the models.

You want to minimize preprocessing time, cost and development effort. How should you configure this workflow?

- **A. Add the transformations as a preprocessing layer in the TensorFlow models.**
- B. Create a Dataflow pipeline that uses the BigQueryIO connector to ingest the data, process it, and write it back to BigQuery.
- C. Write the transformations into Spark that uses the spark-bigquery-connector and use Dataproc to preprocess the data.
- D. Write SQL queries to transform the data in-place in BigQuery.

Antwort: A

Begründung:

The best option for configuring the workflow is to add the transformations as a preprocessing layer in the TensorFlow models. This option allows you to leverage the power and simplicity of TensorFlow to preprocess and transform the data with simple Python code. TensorFlow is a framework for building and training machine learning models. TensorFlow provides various tools and libraries for data analysis and machine learning. A preprocessing layer is a type of layer in TensorFlow that can perform data preprocessing and feature engineering operations on the input data. A preprocessing layer can help you customize the data transformation and preprocessing logic, and handle complex or non-standard data formats. A preprocessing layer can also help you minimize the preprocessing time, cost, and development effort, as you only need to write a few lines of code to implement the preprocessing layer, and you do not need to create any intermediate data sources or pipelines. By adding the transformations as a preprocessing layer in the TensorFlow models, you can use Vertex AI's custom training service to train multiple TensorFlow models that read the data from BigQuery and predict future sales.

The other options are not as good as option A, for the following reasons:

* Option A: Writing the transformations into Spark that uses the spark-bigquery-connector and using Dataproc to preprocess the data would require more skills and steps than using a preprocessing layer in TensorFlow. Spark is a framework for distributed data

processing and machine learning. Spark can read and write data from BigQuery by using the spark-bigquery-connector, which is a library that allows Spark to communicate with BigQuery. Dataproc is a service that can create and manage Spark clusters on Google Cloud. Dataproc can help you run Spark jobs on Google Cloud, and scale the clusters according to the workload. However, writing the transformations into Spark that uses the spark-bigquery-connector and using Dataproc to preprocess the data would require more skills and steps than using a preprocessing layer in TensorFlow. You would need to write code, create and configure the Spark cluster, install and import the spark-bigquery-connector, load and preprocess the data, and write the data back to BigQuery. Moreover, this option would create an intermediate data source in BigQuery, which can increase the storage and computation costs².

* Option B: Writing SQL queries to transform the data in-place in BigQuery would not allow you to use Vertex AI's custom training service to train multiple TensorFlow models that read the data from BigQuery and predict future sales. BigQuery is a service that can perform data analysis and machine learning by using SQL queries. BigQuery can perform data transformation and preprocessing by using SQL functions and clauses, such as MIN, MAX, CASE, and TRANSFORM. BigQuery can also perform machine learning by using BigQuery ML, which is a feature that can create and train machine learning models by using SQL queries. However, writing SQL queries to transform the data in-place in BigQuery would not allow you to use Vertex AI's custom training service to train multiple TensorFlow models that read the data from BigQuery and predict future sales. Vertex AI's custom training service is a service that can run your custom machine learning code on Vertex AI. Vertex AI's custom training service can support various machine learning frameworks, such as TensorFlow, PyTorch, and scikit-learn. Vertex AI's custom training service cannot support SQL queries, as SQL is not a machine learning framework. Therefore, if you want to use Vertex AI's custom training service, you cannot use SQL queries to transform the data in-place in BigQuery³.

* Option D: Creating a Dataflow pipeline that uses the BigQueryIO connector to ingest the data, process it, and write it back to BigQuery would require more skills and steps than using a preprocessing layer in TensorFlow. Dataflow is a service that can create and run data processing and machine learning pipelines on Google Cloud. Dataflow can read and write data from BigQuery by using the BigQueryIO connector, which is a library that allows Dataflow to communicate with BigQuery. Dataflow can perform data transformation and preprocessing by using Apache Beam, which is a framework for distributed data processing and machine learning. However, creating a Dataflow pipeline that uses the BigQueryIO connector to ingest the data, process it, and write it back to BigQuery would require more skills and steps than using a preprocessing layer in TensorFlow. You would need to write code, create and configure the Dataflow pipeline, install and import the BigQueryIO connector, load and preprocess the data, and write the data back to BigQuery. Moreover, this option would create an intermediate data source in BigQuery, which can increase the storage and computation costs⁴.

References:

* Preparing for Google Cloud Certification: Machine Learning Engineer, Course 3: Production ML Systems, Week 2: Serving ML Predictions

* Google Cloud Professional Machine Learning Engineer Exam Guide, Section 2: Developing ML models, 2.1 Developing ML models by using TensorFlow

* Official Google Cloud Certified Professional Machine Learning Engineer Study Guide, Chapter 4: Developing ML Models, Section 4.1: Developing ML Models by Using TensorFlow

* TensorFlow Preprocessing Layers

* Spark and BigQuery

* Dataproc

* BigQuery ML

* Dataflow and BigQuery

* Apache Beam

83. Frage

While performing exploratory data analysis on a dataset, you find that an important categorical feature has 5% null values. You want to minimize the bias that could result from the missing values. How should you handle the missing values?

- A. Replace the missing values with the feature's mean.
- B. Remove the rows with missing values, and upsample your dataset by 5%.
- C. Replace the missing values with a placeholder category indicating a missing value.
- D. Move the rows with missing values to your validation dataset.

Antwort: C

Begründung:

The best option for handling missing values in a categorical feature is to replace them with a placeholder category indicating a missing value. This is a type of imputation, which is a method of estimating the missing values based on the observed data. Imputing the missing values with a placeholder category preserves the information that the data is missing, and avoids introducing bias or distortion in the feature distribution. It also allows the machine learning model to learn from the missingness pattern, and potentially use it as a predictor for the target variable. The other options are not suitable for handling missing values in a categorical feature, because:

* Removing the rows with missing values and upsampling the dataset by 5% would reduce the size of the dataset and potentially lose important information. It would also introduce sampling bias and overfitting, as the upsampling process would create duplicate or synthetic observations that do not reflect the true population.

* Replacing the missing values with the feature's mean would not make sense for a categorical feature, as the mean is a numerical measure that does not capture the mode or frequency of the categories. It would

* also create a new category that does not exist in the original data, and might confuse the machine learning model.

* Moving the rows with missing values to the validation dataset would compromise the validity and reliability of the model evaluation, as the validation dataset would not be representative of the test or production data. It would also reduce the amount of data available for training the model, and might introduce leakage or inconsistency between the training and validation datasets.

References:

* Imputation of missing values

* Effective Strategies to Handle Missing Values in Data Analysis

* How to Handle Missing Values of Categorical Variables?

* Google Cloud launches machine learning engineer certification

* Google Professional Machine Learning Engineer Certification

* Professional ML Engineer Exam Guide

* Preparing for Google Cloud Certification: Machine Learning Engineer Professional Certificate

84. Frage

You work at a bank. You have a custom tabular ML model that was provided by the bank's vendor. The training data is not available due to its sensitivity. The model is packaged as a Vertex AI Model serving container which accepts a string as input for each prediction instance. In each string the feature values are separated by commas. You want to deploy this model to production for online predictions, and monitor the feature distribution over time with minimal effort. What should you do?

- A. 1 Upload the model to Vertex AI Model Registry and deploy the model to a Vertex AI endpoint.
2 Create a Vertex AI Model Monitoring job with feature skew detection as the monitoring objective and provide an instance schema.
- B. 1 Refactor the serving container to accept key-value pairs as input format.
2. Upload the model to Vertex AI Model Registry and deploy the model to a Vertex AI endpoint.
3. Create a Vertex AI Model Monitoring job with feature drift detection as the monitoring objective.
- C. 1 Upload the model to Vertex AI Model Registry and deploy the model to a Vertex AI endpoint.
2. Create a Vertex AI Model Monitoring job with feature drift detection as the monitoring objective, and provide an instance schema.
- D. 1 Refactor the serving container to accept key-value pairs as input format.
2 Upload the model to Vertex AI Model Registry and deploy the model to a Vertex AI endpoint.
3. Create a Vertex AI Model Monitoring job with feature skew detection as the monitoring objective.

Antwort: C

Begründung:

The best option for deploying a custom tabular ML model to production for online predictions, and monitoring the feature distribution over time with minimal effort, using a model that was provided by the bank's vendor, the training data is not available due to its sensitivity, and the model is packaged as a Vertex AI Model serving container which accepts a string as input for each prediction instance, is to upload the model to Vertex AI Model Registry and deploy the model to a Vertex AI endpoint, create a Vertex AI Model Monitoring job with feature drift detection as the monitoring objective, and provide an instance schema. This option allows you to leverage the power and simplicity of Vertex AI to serve and monitor your model with minimal code and configuration. Vertex AI is a unified platform for building and deploying machine learning solutions on Google Cloud. Vertex AI can deploy a trained model to an online prediction endpoint, which can provide low-latency predictions for individual instances. Vertex AI can also provide various tools and services for data analysis, model development, model deployment, model monitoring, and model governance. A Vertex AI Model Registry is a resource that can store and manage your models on Vertex AI. A Vertex AI Model Registry can help you organize and track your models, and access various model information, such as model name, model description, and model labels. A Vertex AI Model serving container is a resource that can run your custom model code on Vertex AI. A Vertex AI Model serving container can help you package your model code and dependencies into a container image, and deploy the container image to an online prediction endpoint. A Vertex AI Model serving container can accept various input formats, such as JSON, CSV, or TFRecord. A string input format is a type of input format that accepts a string as input for each prediction instance. A string input format can help you encode your feature values into a single string, and separate them by commas. By uploading the model to Vertex AI Model Registry and deploying the model to a Vertex AI endpoint, you can serve your model for online predictions with minimal code and configuration. You can use the Vertex AI API or the gcloud command-line tool to upload the model to Vertex AI Model Registry, and provide the model name, model description, and model labels. You can also use the Vertex AI API or the gcloud command-line tool to deploy the model to a Vertex AI endpoint, and provide the endpoint name,

endpoint description, endpoint labels, and endpoint resources. A Vertex AI Model Monitoring job is a resource that can monitor the performance and quality of your deployed models on Vertex AI. A Vertex AI Model Monitoring job can help you detect and diagnose issues with your models, such as data drift, prediction drift, training/serving skew, or model staleness. Feature drift is a type of model monitoring metric that measures the difference between the distributions of the features used to train the model and the features used to serve the model over time. Feature drift can indicate that the online data is changing over time, and the model performance is degrading. By creating a Vertex AI Model Monitoring job with feature drift detection as the monitoring objective, and providing an instance schema, you can monitor the feature distribution over time with minimal effort. You can use the Vertex AI API or the `gcloud` command-line tool to create a Vertex AI Model Monitoring job, and provide the monitoring objective, the monitoring frequency, the alerting threshold, and the notification channel. You can also provide an instance schema, which is a JSON file that describes the features and their types in the prediction input data. An instance schema can help Vertex AI Model Monitoring parse and analyze the string input format, and calculate the feature distributions and distance scores¹.

The other options are not as good as option A, for the following reasons:

Option B: Uploading the model to Vertex AI Model Registry and deploying the model to a Vertex AI endpoint, creating a Vertex AI Model Monitoring job with feature skew detection as the monitoring objective, and providing an instance schema would not help you monitor the changes in the online data over time, and could cause errors or poor performance. Feature skew is a type of model monitoring metric that measures the difference between the distributions of the features used to train the model and the features used to serve the model at a given point in time. Feature skew can indicate that the model is not trained on the representative data, or that the data is changing over time. By creating a Vertex AI Model Monitoring job with feature skew detection as the monitoring objective, and providing an instance schema, you can monitor the feature distribution at a given point in time with minimal effort. However, uploading the model to Vertex AI Model Registry and deploying the model to a Vertex AI endpoint, creating a Vertex AI Model Monitoring job with feature skew detection as the monitoring objective, and providing an instance schema would not help you monitor the changes in the online data over time, and could cause errors or poor performance. You would need to use the Vertex AI API or the `gcloud` command-line tool to upload the model to Vertex AI Model Registry, deploy the model to a Vertex AI endpoint, create a Vertex AI Model Monitoring job, and provide an instance schema. Moreover, this option would not monitor the feature drift, which is a more direct and relevant metric for measuring the changes in the online data over time, and the model performance and quality¹.

Option C: Refactoring the serving container to accept key-value pairs as input format, uploading the model to Vertex AI Model Registry and deploying the model to a Vertex AI endpoint, creating a Vertex AI Model Monitoring job with feature drift detection as the monitoring objective would require more skills and steps than uploading the model to Vertex AI Model Registry and deploying the model to a Vertex AI endpoint, creating a Vertex AI Model Monitoring job with feature drift detection as the monitoring objective, and providing an instance schema. A key-value pair input format is a type of input format that accepts a key-value pair as input for each prediction instance. A key-value pair input format can help you specify the feature names and values in a JSON object, and separate them by colons. By refactoring the serving container to accept key-value pairs as input format, uploading the model to Vertex AI Model Registry and deploying the model to a Vertex AI endpoint, creating a Vertex AI Model Monitoring job with feature drift detection as the monitoring objective, you can serve and monitor your model with minimal code and configuration. You can write code to refactor the serving container to accept key-value pairs as input format, and use the Vertex AI API or the `gcloud` command-line tool to upload the model to Vertex AI Model Registry, deploy the model to a Vertex AI endpoint, and create a Vertex AI Model Monitoring job. However, refactoring the serving container to accept key-value pairs as input format, uploading the model to Vertex AI Model Registry and deploying the model to a Vertex AI endpoint, creating a Vertex AI Model Monitoring job with feature drift detection as the monitoring objective would require more skills and steps than uploading the model to Vertex AI Model Registry and deploying the model to a Vertex AI endpoint, creating a Vertex AI Model Monitoring job with feature drift detection as the monitoring objective, and providing an instance schema. You would need to write code, refactor the serving container, upload the model to Vertex AI Model Registry, deploy the model to a Vertex AI endpoint, and create a Vertex AI Model Monitoring job. Moreover, this option would not use the instance schema, which is a JSON file that can help Vertex AI Model Monitoring parse and analyze the string input format, and calculate the feature distributions and distance scores¹.

Option D: Refactoring the serving container to accept key-value pairs as input format, uploading the model to Vertex AI Model Registry and deploying the model to a Vertex AI endpoint, creating a Vertex AI Model Monitoring job with feature skew detection as the monitoring objective would require more skills and steps than uploading the model to Vertex AI Model Registry and deploying the model to a Vertex AI endpoint, creating a Vertex AI Model Monitoring job with feature drift detection as the monitoring objective, and providing an instance schema, and would not help you monitor the changes in the online data over time, and could cause errors or poor performance. Feature skew is a type of model monitoring metric that measures the difference between the distributions of the features used to train the model and the features used to serve the model at a given point in time. Feature skew can indicate that the model is not trained on the representative data, or that the data is changing over time. By creating a Vertex AI Model Monitoring job with feature skew detection as the monitoring objective, you can monitor the feature distribution at a given point in time with minimal effort. However, refactoring the serving container to accept key-value pairs as input format, uploading the model to Vertex AI Model Registry and deploying the model to a Vertex AI endpoint, creating a Vertex AI Model Monitoring job with feature skew detection as the monitoring objective would require more skills and steps than uploading the model to Vertex AI Model Registry and deploying the model to a Vertex AI endpoint, creating a Vertex AI Model Monitoring job with feature drift detection as the monitoring objective, and providing an instance schema, and would not help you monitor the changes in the online data over time, and could cause errors or poor performance. You would need to write code, refactor the serving container, upload the model to Vertex AI Model Registry, deploy the model to a Vertex AI endpoint, and create a Vertex AI

Model Monitoring job. Moreover, this option would not monitor the feature drift, which is a more direct and relevant metric for measuring the changes in the online data over time, and the model performance and quality.

Reference:

Using Model Monitoring | Vertex AI | Google Cloud

85. Frage

You are developing an ML model intended to classify whether X-Ray images indicate bone fracture risk. You have trained on Api Resnet architecture on Vertex AI using a TPU as an accelerator, however you are unsatisfied with the training time and use memory usage. You want to quickly iterate your training code but make minimal changes to the code. You also want to minimize impact on the models accuracy. What should you do?

- A. Reduce the number of layers in the model architecture
- **B. Reduce the global batch size from 1024 to 256**
- C. Reduce the dimensions of the images used un the model
- D. Configure your model to use bfloat16 instead float32

Antwort: B

86. Frage

.....

Wir bemühen uns nun darum, den Kandidaten rechtzeitigen und effizienten Service zu bieten, um Ihre wertvolle Zeit zu ersparen. Fast2test bietet Ihnen zahlreiche Lerntipps, Fragen und Antworten zur Google Professional-Machine-Learning-Engineer Zertifizierungsprüfung. Einige Websites bieten Ihnen auch Lernmaterialien zur Professional-Machine-Learning-Engineer Zertifizierungsprüfung, die von guter Qualität ist und mit dem Zeit Schritt halten. Aber Fast2test ist die einzige Website, die beste Schulungsunterlagen zur Professional-Machine-Learning-Engineer Zertifizierungsprüfung bietet. Mit Hilfe der Lernmaterialien und der Anleitung von Fast2test können Sie die Google Professional-Machine-Learning-Engineer Zertifizierungsprüfung einmalig bestehen.

Professional-Machine-Learning-Engineer Pruefungssimulationen: <https://de.fast2test.com/Professional-Machine-Learning-Engineer-premium-file.html>

- Sie können so einfach wie möglich - Professional-Machine-Learning-Engineer bestehen! Suchen Sie jetzt auf (www.zertsoft.com) nach Professional-Machine-Learning-Engineer und laden Sie es kostenlos herunter
➔ Professional-Machine-Learning-Engineer Prüfungsaufgaben
- Google Professional-Machine-Learning-Engineer Prüfung Übungen und Antworten Suchen Sie einfach auf www.itzert.com nach kostenloser Download von { Professional-Machine-Learning-Engineer } Professional-Machine-Learning-Engineer Testking
- Professional-Machine-Learning-Engineer Buch Professional-Machine-Learning-Engineer Prüfung Professional-Machine-Learning-Engineer Deutsche Suchen Sie einfach auf [de.fast2test.com] nach kostenloser Download von [Professional-Machine-Learning-Engineer] Professional-Machine-Learning-Engineer Deutsche
- Google Professional-Machine-Learning-Engineer Prüfung Übungen und Antworten Suchen Sie einfach auf { www.itzert.com } nach kostenloser Download von ➤ Professional-Machine-Learning-Engineer Professional-Machine-Learning-Engineer Examengine
- Professional-Machine-Learning-Engineer zu bestehen mit allseitigen Garantien Öffnen Sie die Webseite ➡ www.zertpruefung.de und suchen Sie nach kostenloser Download von ▷ Professional-Machine-Learning-Engineer ◁ Professional-Machine-Learning-Engineer Fragenkatalog
- Professional-Machine-Learning-Engineer zu bestehen mit allseitigen Garantien Sie müssen nur zu [www.itzert.com] gehen um nach kostenloser Download von ▷ Professional-Machine-Learning-Engineer ◁ zu suchen Professional-Machine-Learning-Engineer Echte Fragen
- Sie können so einfach wie möglich - Professional-Machine-Learning-Engineer bestehen! Öffnen Sie die Website ⇒ www.pass4test.de ⇐ Suchen Sie ➡ Professional-Machine-Learning-Engineer Kostenloser Download Professional-Machine-Learning-Engineer Prüfungen
- Professional-Machine-Learning-Engineer: Google Professional Machine Learning Engineer Dumps - PassGuide Professional-Machine-Learning-Engineer Examen Suchen Sie jetzt auf ▶ www.itzert.com ◀ nach Professional-Machine-Learning-Engineer und laden Sie es kostenlos herunter Professional-Machine-Learning-Engineer Prüfungen
- Professional-Machine-Learning-Engineer Testking Professional-Machine-Learning-Engineer Prüfungsfragen Professional-Machine-Learning-Engineer Fragen&Antworten Öffnen Sie 《 www.it-pruefung.com 》 geben Sie (Professional-Machine-Learning-Engineer) ein und erhalten Sie den kostenlosen Download Professional-Machine-Learning-Engineer Buch

