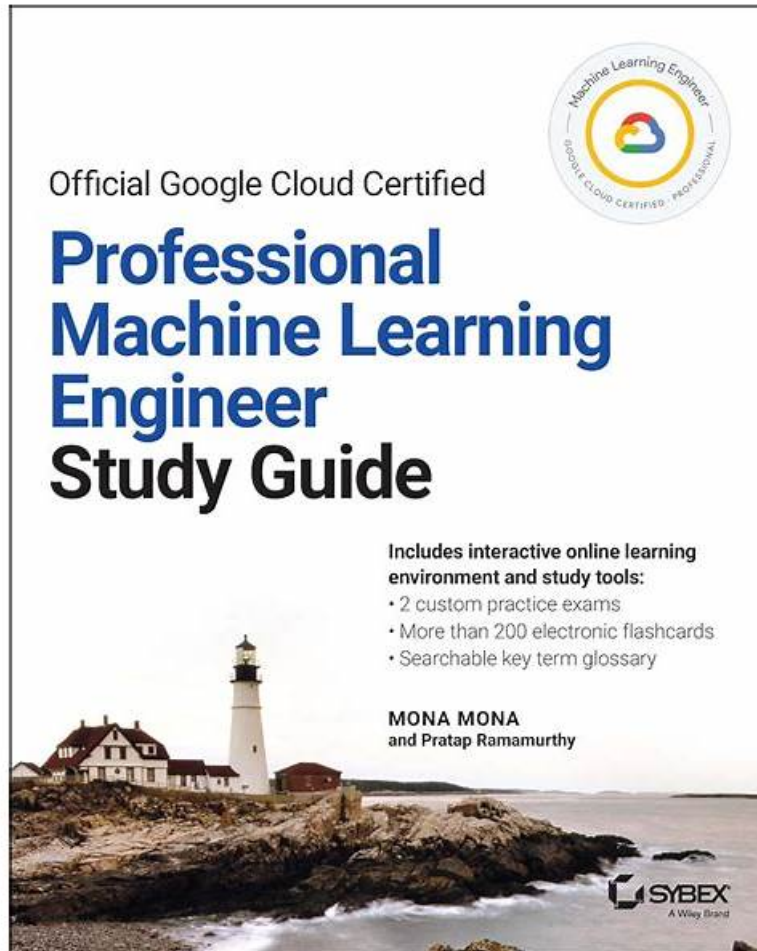


# 素敵な Professional-Machine-Learning-Engineer クラム メディア | 最初の試行で簡単に勉強して試験に合格する & 最高の Google Google Professional Machine Learning Engineer



P.S.Pass4TestがGoogle Driveで共有している無料の2026 Google Professional-Machine-Learning-Engineerダウンロード: [https://drive.google.com/open?id=1h1TK-BN9bdqsBSaD5J5\\_ozdNH1enDHtf](https://drive.google.com/open?id=1h1TK-BN9bdqsBSaD5J5_ozdNH1enDHtf)

GoogleのProfessional-Machine-Learning-Engineerは専門知識と情報技術の検査として認証試験で、Pass4Testはあなたに一日早くGoogleの認証試験に合格させて、多くの人が大量の時間とエネルギーを費やしても無駄になりました。Pass4Testにその問題が心配でなく、わずか20時間と少ないお金をを使って楽に試験に合格することができます。Pass4Testは君に対して特別の訓練を提供しています。

Google Professional Machine Learning Engineer Certification Examは、Google Cloud Platform上で機械学習モデルの設計、構築、展開のスキルを証明するための非常に評価の高い認定試験です。この認定試験は、機械学習アルゴリズム、データ分析、クラウドコンピューティングの深い理解が必要です。

Google認定プログラムは、対話型のコースと経験豊富な講師を提供し、知識とスキルセットの技術を改善します。認定後に認定プロフェッショナルのコミュニティとネットワークは、個人が利益を得るための重要な利点です。この試験は、実践的な経験を提供するリアルライフのケーススタディやプロジェクトを使用した実践的なトレーニングと経験を強調しています。Google Professional Machine Learning Engineer認定は、既に習得したスキルの信頼性と妥当性を向上させたい個人に適した投資であり、機械学習の分野での主要な認定資格となっています。

この認定を取得するには、候補者は機械学習とクラウドコンピューティングに関連する幅広いトピックをカバーする厳格な試験に合格する必要があります。試験は、複数選択とシナリオベースの質問で構成されてお

り、候補者は試験を完了するために2時間半が与えられます。試験はオンラインで管理されており、世界中のどこからでも取得できます。試験に合格すると、候補者はLinkedInプロフィール、履歴書、またはWebサイトに表示できるデジタルバッジを受け取り、機械学習の分野とGoogle Cloudプラットフォームの習熟度を示していることを示します。この認定は、業界の専門家によって認識され、個人が機械学習とクラウドコンピューティングの分野でキャリアを前進させるのに役立ちます。

>> Professional-Machine-Learning-Engineer クラムメディア <<

## 100%合格率のProfessional-Machine-Learning-Engineer クラムメディア & 合格スムーズ Professional-Machine-Learning-Engineer テストサンプル問題 | 更新する Professional-Machine-Learning-Engineer 試験対策書

この競争の激しい社会では、良い仕事をするためには、自分の能力を向上させ、可能性を常に探求し、関連するProfessional-Machine-Learning-Engineer認定を取得することが最善の方法です。しかし、私たちの専門的な能力は、試験を解読するのが難しいことであり、試験に関連するProfessional-Machine-Learning-Engineer準備質問が非常に多いため、試験に必要なすべてのキーポイントを体系化することは不可能です。

### Google Professional Machine Learning Engineer 認定 Professional-Machine-Learning-Engineer 試験問題 (Q88-Q93):

#### 質問 # 88

You work for an organization that operates a streaming music service. You have a custom production model that is serving a "next song" recommendation based on a user's recent listening history. Your model is deployed on a Vertex AI endpoint. You recently retrained the same model by using fresh data. The model received positive test results offline. You now want to test the new model in production while minimizing complexity. What should you do?

- A. Configure a model monitoring job for the existing Vertex AI endpoint. Configure the monitoring job to detect prediction drift, and set a threshold for alerts Update the model on the endpoint from the previous model to the new model If you receive an alert of prediction drift, revert to the previous model.
- B. Deploy the new model to the existing Vertex AI endpoint Use traffic splitting to send 5% of production traffic to the new model Monitor end-user metrics, such as listening time If end-user metrics improve between models over time, gradually increase the percentage of production traffic sent to the new model.
- C. Capture incoming prediction requests in BigQuery Create an experiment in Vertex AI Experiments Run batch predictions for both models using the captured data Use the user's selected song to compare the models performance side by side If the new models performance metrics are better than the previous model deploy the new model to production.
- D. Create a new Vertex AI endpoint for the new model and deploy the new model to that new endpoint Build a service to randomly send 5% of production traffic to the new endpoint Monitor end-user metrics such as listening time If end-user metrics improve between models over time gradually increase the percentage of production traffic sent to the new endpoint.

正解: B

解説:

Traffic splitting is a feature of Vertex AI that allows you to distribute the prediction requests among multiple models or model versions within the same endpoint. You can specify the percentage of traffic that each model or model version receives, and change it at any time. Traffic splitting can help you test the new model in production without creating a new endpoint or a separate service. You can deploy the new model to the existing Vertex AI endpoint, and use traffic splitting to send 5% of production traffic to the new model. You can monitor the end-user metrics, such as listening time, to compare the performance of the new model and the previous model. If the end-user metrics improve between models over time, you can gradually increase the percentage of production traffic sent to the new model. This solution can help you test the new model in production while minimizing complexity and cost.

References:

\* Traffic splitting | Vertex AI

\* Deploying models to endpoints | Vertex AI

#### 質問 # 89

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. Move the rows with missing values to your validation dataset.
- 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. Replace the missing values with the feature's mean.

正解: C

解説:

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

#### 質問 #90

You are training an object detection model using a Cloud TPU v2. Training time is taking longer than expected. Based on this simplified trace obtained with a Cloud TPU profile, what action should you take to decrease training time in a cost-efficient way?

- A. Rewrite your input function using parallel reads, parallel processing, and prefetch.
- B. Move from Cloud TPU v2 to Cloud TPU v3 and increase batch size.
- C. Rewrite your input function to resize and reshape the input images.
- D. Move from Cloud TPU v2 to 8 NVIDIA V100 GPUs and increase batch size.

正解: B

#### 質問 #91

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 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 do not meet the profile photo requirements.
- D. Use AutoML to optimize the model's recall in order to minimize false negatives.

正解: D

解説:

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.

## 質問 # 92

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 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.
- **B. 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.**
- 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 skew 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 drift detection as the monitoring objective.

正解: B

解説:

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 scores [1].

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 quality [1].

\* 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 scores [1].

\* 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 [1].

References:

Using Model Monitoring | Vertex AI | Google Cloud



