


Google Professional-Machine-Learning-Engineer資料勉強、Professional-Machine-Learning-Engineer模擬試験最新版



さらに、Tech4Exam Professional-Machine-Learning-Engineerダンプの一部が現在無料で提供されています：https://drive.google.com/open?id=1AJ8TjN7MoAMQwXTtFbxhOarxV23QN_B3

Tech4ExamはIT認定試験を受験した多くの人々を助きました。また、受験生からいろいろな良い評価を得ています。Tech4ExamのProfessional-Machine-Learning-Engineer問題集の合格率が100%に達することも数え切れない受験生に証明された事実です。もし試験の準備をするために大変を感じているとしたら、ぜひTech4ExamのProfessional-Machine-Learning-Engineer問題集を見逃さないでください。これは試験の準備をするために非常に効率的なツールです。この問題集はあなたが少ない労力で最高の結果を取得することができます。

Google Professional Machine Learning Engineer認定を取得するには、候補者は、複数選択およびシナリオベースの質問で構成される2時間の試験に合格する必要があります。この試験では、Google Cloudプラットフォームを使用して、スケーラブルで効率的で安全な機械学習モデルを設計および開発する能力について候補者を評価します。この認定は世界的に認識されており、専門家の専門知識と機械学習の分野で最先端のテクノロジーを扱う能力を実証するため、雇用主は高く評価されています。この認定はまた、Googleのリソースと機械学習の専門家コミュニティへのアクセスを提供し、この分野でのキャリアを進歩させたい人にとって貴重な資産となっています。

Google Professional-Machine-Learning-Engineer（Google Professional Machine Learning Engineer）認定試験は、Googleが提供するGoogleが提供するプロのレベルの認定試験であり、Google Cloudプラットフォームに機械学習モデルの構築と展開をテストします。この試験は、機械学習の概念を確実に理解し、Google Cloudプラットフォームに機械学習モデルの構築と展開の経験がある個人向けに設計されています。

Professional-Machine-Learning-Engineer試験の準備方法 | 完璧な Professional-Machine-Learning-Engineer資料勉強試験 | 正確的な Google Professional Machine Learning Engineer模擬試験最新版

Tech4ExamのGoogleのProfessional-Machine-Learning-Engineer試験トレーニング資料はPDF形式とソフトウェアの形式で提供します。それはTech4ExamのGoogleのProfessional-Machine-Learning-Engineer試験の問題と解答を含めます。そして、その学習教材の内容はカバー率が高くて、正確率も高いです。それはきっと君のGoogleのProfessional-Machine-Learning-Engineer試験に合格することの良い参考資料です。もし不合格になる場合は、ご心配なく、私たちは資料の費用を全部返金します。

Google Professional Machine Learning Engineer 認定 Professional-Machine-Learning-Engineer 試験問題 (Q163-Q168):

質問 # 163

You work for a pet food company that manages an online forum. Customers upload photos of their pets on the forum to share with others. About 20 photos are uploaded daily. You want to automatically and in near real time detect whether each uploaded photo has an animal. You want to prioritize time and minimize cost of your application development and deployment. What should you do?

- A. Manually label previously submitted images with bounding boxes around any animals. Build an AutoML object detection model by using Vertex AI. Deploy the model to a Vertex AI endpoint. Send new user-submitted images to your model endpoint to detect whether each photo has an animal.
- B. Manually label previously submitted images as having animals or not. Create an image dataset on Vertex AI. Train a classification model by using Vertex AutoML to distinguish the two classes. Deploy the model to a Vertex AI endpoint. Send new user-submitted images to your model endpoint to classify whether each photo has an animal.
- C. Download an object detection model from TensorFlow Hub. Deploy the model to a Vertex AI endpoint. Send new user-submitted images to the model endpoint to classify whether each photo has an animal.
- **D. Send user-submitted images to the Cloud Vision API. Use object localization to identify all objects in the image and compare the results against a list of animals.**

正解: D

解説:

Cloud Vision API is a service that allows you to analyze images using pre-trained machine learning models¹. You can use Cloud Vision API to perform various tasks, such as face detection, text extraction, logo recognition, and object localization¹. Object localization is a feature that allows you to detect multiple objects in an image and draw bounding boxes around them². You can also get the labels and confidence scores for each detected object².

By sending user-submitted images to the Cloud Vision API, you can use object localization to identify all objects in the image and compare the results against a list of animals. You can use the OBJECT_LOCALIZATION feature type in the AnnotateImageRequest to request object localization³. You can then use the localizedObjectAnnotations field in the AnnotateImageResponse to get the list of detected objects, their labels, and their confidence scores. You can compare the labels with a predefined list of animals, such as dogs, cats, birds, etc., and determine whether the image has an animal or not.

This option is the best for your scenario, because it allows you to automatically and in near real time detect whether each uploaded photo has an animal, without requiring any manual labeling, model training, or model deployment. You can also prioritize time and minimize cost of your application development and deployment, as you can use the Cloud Vision API as a ready-to-use service, without needing any machine learning expertise or infrastructure.

The other options are not suitable for your scenario, because they either require manual labeling, model training, or model deployment, which would increase the time and cost of your application development and deployment, or they use object detection models, which are more complex and computationally expensive than object localization models, and are not necessary for your simple task of detecting whether an image has an animal or not.

Reference:

Cloud Vision API | Google Cloud

Object localization | Cloud Vision API | Google Cloud

AnnotateImageRequest | Cloud Vision API | Google Cloud

[AnnotateImageResponse | Cloud Vision API | Google Cloud]

質問 # 164

You have created a Vertex AI pipeline that includes two steps. The first step preprocesses 10 TB data completes in about 1 hour, and saves the result in a Cloud Storage bucket. The second step uses the processed data to train a model. You need to update the model's code to allow you to test different algorithms. You want to reduce pipeline execution time and cost, while also minimizing pipeline changes. What should you do?

- A. Add a pipeline parameter and an additional pipeline step. Depending on the parameter value, the pipeline step conducts or skips data preprocessing and starts model training.
- **B. Enable caching for the pipeline job, and disable caching for the model training step.**
- C. Configure a machine with more CPU and RAM from the compute-optimized machine family for the data preprocessing step.
- D. Create another pipeline without the preprocessing step, and hardcode the preprocessed Cloud Storage file location for model training.

正解: B

解説:

The best option for reducing pipeline execution time and cost, while also minimizing pipeline changes, is to enable caching for the pipeline job, and disable caching for the model training step. This option allows you to leverage the power and simplicity of Vertex AI Pipelines to reuse the output of the data preprocessing step, and avoid unnecessary recomputation. Vertex AI Pipelines is a service that can orchestrate machine learning workflows using Vertex AI. Vertex AI Pipelines can run preprocessing and training steps on custom Docker images, and evaluate, deploy, and monitor the machine learning model. Caching is a feature of Vertex AI Pipelines that can store and reuse the output of a pipeline step, and skip the execution of the step if the input parameters and the code have not changed. Caching can help you reduce the pipeline execution time and cost, as you do not need to re-run the same step with the same input and code. Caching can also help you minimize the pipeline changes, as you do not need to add or remove any pipeline steps or parameters. By enabling caching for the pipeline job, and disabling caching for the model training step, you can create a Vertex AI pipeline that includes two steps. The first step preprocesses 10 TB data, completes in about 1 hour, and saves the result in a Cloud Storage bucket. The second step uses the processed data to train a model. You can update the model's code to allow you to test different algorithms, and run the pipeline job with caching enabled. The pipeline job will reuse the output of the data preprocessing step from the cache, and skip the execution of the step. The pipeline job will run the model training step with the updated code, and disable the caching for the step. This way, you can reduce the pipeline execution time and cost, while also minimizing pipeline changes.

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

Option A: Adding a pipeline parameter and an additional pipeline step, depending on the parameter value, the pipeline step conducts or skips data preprocessing and starts model training, would require more skills and steps than enabling caching for the pipeline job, and disabling caching for the model training step. A pipeline parameter is a variable that can be used to control the input or output of a pipeline step. A pipeline parameter can help you customize the pipeline logic and behavior, and experiment with different values. An additional pipeline step is a new instance of a pipeline component that can perform a part of the pipeline workflow, such as data preprocessing or model training. An additional pipeline step can help you extend the pipeline functionality and complexity, and handle different scenarios. However, adding a pipeline parameter and an additional pipeline step, depending on the parameter value, the pipeline step conducts or skips data preprocessing and starts model training, would require more skills and steps than enabling caching for the pipeline job, and disabling caching for the model training step. You would need to write code, define the pipeline parameter, create the additional pipeline step, implement the conditional logic, and compile and run the pipeline. Moreover, this option would not reuse the output of the data preprocessing step from the cache, but rather from the Cloud Storage bucket, which can increase the data transfer and access costs.

Option B: Creating another pipeline without the preprocessing step, and hardcoding the preprocessed Cloud Storage file location for model training, would require more skills and steps than enabling caching for the pipeline job, and disabling caching for the model training step. A pipeline without the preprocessing step is a pipeline that only includes the model training step, and uses the preprocessed data from the Cloud Storage bucket as the input. A pipeline without the preprocessing step can help you avoid running the data preprocessing step every time, and reduce the pipeline execution time and cost. However, creating another pipeline without the preprocessing step, and hardcoding the preprocessed Cloud Storage file location for model training, would require more skills and steps than enabling caching for the pipeline job, and disabling caching for the model training step. You would need to write code, create a new pipeline, remove the preprocessing step, hardcode the Cloud Storage file location, and compile and run the pipeline. Moreover, this option would not reuse the output of the data preprocessing step from the cache, but rather from the Cloud Storage bucket, which can increase the data transfer and access costs. Furthermore, this option would create another pipeline, which can increase the maintenance and management costs.

Option C: Configuring a machine with more CPU and RAM from the compute-optimized machine family for the data preprocessing step, would not reduce the pipeline execution time and cost, while also minimizing pipeline changes, but rather increase the pipeline execution cost and complexity. A machine with more CPU and RAM from the compute-optimized machine family is a virtual machine that has a high ratio of CPU cores to memory, and can provide high performance and scalability for compute-intensive workloads. A machine with more CPU and RAM from the compute-optimized machine family can help you optimize the data

preprocessing step, and reduce the pipeline execution time. However, configuring a machine with more CPU and RAM from the compute-optimized machine family for the data preprocessing step, would not reduce the pipeline execution time and cost, while also minimizing pipeline changes, but rather increase the pipeline execution cost and complexity. You would need to write code, configure the machine type parameters for the data preprocessing step, and compile and run the pipeline. Moreover, this option would increase the pipeline execution cost, as machines with more CPU and RAM from the compute-optimized machine family are more expensive than machines with less CPU and RAM from other machine families. Furthermore, this option would not reuse the output of the data preprocessing step from the cache, but rather re-run the data preprocessing step every time, which can increase the pipeline execution time and cost.

Reference:

Preparing for Google Cloud Certification: Machine Learning Engineer, Course 3: Production ML Systems, Week 3: MLOps Google Cloud Professional Machine Learning Engineer Exam Guide, Section 3: Scaling ML models in production, 3.2 Automating ML workflows Official Google Cloud Certified Professional Machine Learning Engineer Study Guide, Chapter 6: Production ML Systems, Section 6.4: Automating ML Workflows Vertex AI Pipelines Caching Pipeline parameters Machine types

質問 # 165

A Machine Learning Specialist is developing a custom video recommendation model for an application. The dataset used to train this model is very large with millions of data points and is hosted in an Amazon S3 bucket.

The Specialist wants to avoid loading all of this data onto an Amazon SageMaker notebook instance because it would take hours to move and will exceed the attached 5 GB Amazon EBS volume on the notebook instance.

Which approach allows the Specialist to use all the data to train the model?

- A. Launch an Amazon EC2 instance with an AWS Deep Learning AMI and attach the S3 bucket to the instance. Train on a small amount of the data to verify the training code and hyperparameters. Go back to Amazon SageMaker and train using the full dataset
- **B. Load a smaller subset of the data into the SageMaker notebook and train locally. Confirm that the training code is executing and the model parameters seem reasonable. Initiate a SageMaker training job using the full dataset from the S3 bucket using Pipe input mode.**
- C. Use AWS Glue to train a model using a small subset of the data to confirm that the data will be compatible with Amazon SageMaker. Initiate a SageMaker training job using the full dataset from the S3 bucket using Pipe input mode.
- D. Load a smaller subset of the data into the SageMaker notebook and train locally. Confirm that the training code is executing and the model parameters seem reasonable. Launch an Amazon EC2 instance with an AWS Deep Learning AMI and attach the S3 bucket to train the full dataset.

正解: B

質問 # 166

You are developing an ML model that predicts the cost of used automobiles based on data such as location, condition model type color, and engine-battery efficiency. The data is updated every night Car dealerships will use the model to determine appropriate car prices. You created a Vertex AI pipeline that reads the data splits the data into training/evaluation/test sets performs feature engineering trains the model by using the training dataset and validates the model by using the evaluation dataset. You need to configure a retraining workflow that minimizes cost What should you do?

- **A. Compare the training and evaluation losses of the current run If the losses are similar deploy the model to a Vertex AI endpoint with training/serving skew threshold model monitoring When the model monitoring threshold is triggered redeploy the pipeline.**
- B. Compare the training and evaluation losses of the current run If the losses are similar, deploy the model to a Vertex AI endpoint Configure a cron job to redeploy the pipeline every night.
- C. Compare the results to the evaluation results from a previous run If the performance improved deploy the model to a Vertex AI endpoint with training/serving skew threshold model monitoring. When the model monitoring threshold is triggered, redeploy the pipeline.
- D. Compare the results to the evaluation results from a previous run If the performance improved deploy the model to a Vertex AI endpoint Configure a cron job to redeploy the pipeline every night.

正解: A

解説:

Comparing the training and evaluation losses of the current run is a good way to check if the model is overfitting or underfitting. If the losses are similar, it means that the model is generalizing well and can be deployed to a Vertex AI endpoint. Vertex AI endpoint is a service that allows you to serve your ML models online and scale them automatically. By using a training/serving skew threshold

model monitoring, you can detect if there is a significant difference between the data used for training and the data used for serving. This can indicate that the model is becoming stale or inaccurate over time. When the model monitoring threshold is triggered, it means that the model needs to be retrained with the latest data. By redeploying the pipeline, you can automate the retraining process and update the model with the new data. This way, you can minimize the cost of retraining and ensure that your model is always up-to-date and accurate. References:

- * Vertex AI documentation
- * Vertex AI endpoint documentation
- * Model monitoring documentation
- * Preparing for Google Cloud Certification: Machine Learning Engineer Professional Certificate

質問 # 167

You need to execute a batch prediction on 100 million records in a BigQuery table with a custom TensorFlow DNN regressor model, and then store the predicted results in a BigQuery table. You want to minimize the effort required to build this inference pipeline. What should you do?

- A. Create a Dataflow pipeline to convert the data in BigQuery to TFRecords. Run a batch inference on Vertex AI Prediction, and write the results to BigQuery.
- B. Use the TensorFlow BigQuery reader to load the data, and use the BigQuery API to write the results to BigQuery.
- C. Load the TensorFlow SavedModel in a Dataflow pipeline. Use the BigQuery I/O connector with a custom function to perform the inference within the pipeline, and write the results to BigQuery.
- **D. Import the TensorFlow model with BigQuery ML, and run the ml.predict function.**

正解: D

解説:

* Option A is correct because importing the TensorFlow model with BigQuery ML, and running the ml.predict function is the easiest way to execute a batch prediction on a large BigQuery table with a custom TensorFlow model, and store the predicted results in another BigQuery table. BigQuery ML

* allows you to import TensorFlow models that are stored in Cloud Storage, and use them for prediction with SQL queries¹. The ml.predict function returns a table with the predicted values, which can be saved to another BigQuery table².

* Option B is incorrect because using the TensorFlow BigQuery reader to load the data, and using the BigQuery API to write the results to BigQuery requires more effort to build the inference pipeline than option A. The TensorFlow BigQuery reader is a way to read data from BigQuery into TensorFlow datasets, which can be used for training or prediction³. However, this option also requires writing code to load the TensorFlow model, run the prediction, and use the BigQuery API to write the results back to BigQuery⁴.

* Option C is incorrect because creating a Dataflow pipeline to convert the data in BigQuery to TFRecords, running a batch inference on Vertex AI Prediction, and writing the results to BigQuery requires more effort to build the inference pipeline than option A. Dataflow is a service for creating and running data processing pipelines, such as ETL (extract, transform, load) or batch processing⁵. Vertex AI Prediction is a service for deploying and serving ML models for online or batch prediction. However, this option also requires writing code to create the Dataflow pipeline, convert the data to TFRecords, run the batch inference, and write the results to BigQuery.

* Option D is incorrect because loading the TensorFlow SavedModel in a Dataflow pipeline, using the BigQuery I/O connector with a custom function to perform the inference within the pipeline, and writing the results to BigQuery requires more effort to build the inference pipeline than option A. The BigQuery I/O connector is a way to read and write data from BigQuery within a Dataflow pipeline. However, this option also requires writing code to load the TensorFlow SavedModel, create the custom function for inference, and write the results to BigQuery.

References:

- * Importing models into BigQuery ML
- * Using imported models for prediction
- * TensorFlow BigQuery reader
- * BigQuery API
- * Dataflow overview
- * [Vertex AI Prediction overview]
- * [Batch prediction with Dataflow]
- * [BigQuery I/O connector]
- * [Using TensorFlow models in Dataflow]

質問 # 168

.....

Google Professional-Machine-Learning-Engineer試験参考書を利用すれば、あなたは多くの時間を節約するだけでなく、いろいろな知識を身につけます。最も重要なのは、Professional-Machine-Learning-Engineer認定試験資格証明書を取得できるということです。また、Professional-Machine-Learning-Engineer試験参考書の合格率は高いので、Professional-Machine-Learning-Engineer試験に落ちる必要がないです。

Professional-Machine-Learning-Engineer模擬試験最新版: <https://www.tech4exam.com/Professional-Machine-Learning-Engineer-pass-shiken.html>

- Professional-Machine-Learning-Engineer試験関連赤本 □ Professional-Machine-Learning-Engineer問題例 □ Professional-Machine-Learning-Engineer試験対策 □ ➡ www.passtest.jp □ を開いて ✨ Professional-Machine-Learning-Engineer □ ✨ □ を検索し、試験資料を無料でダウンロードしてください Professional-Machine-Learning-Engineer日本語版トレーニング
- 最新のProfessional-Machine-Learning-Engineer資料勉強 - 合格スムーズ Professional-Machine-Learning-Engineer模擬試験最新版 | 一番優秀なProfessional-Machine-Learning-Engineer問題と解答 Google Professional Machine Learning Engineer □ ▷ www.goshiken.com ◁ サイトにて最新 ➡ Professional-Machine-Learning-Engineer □ 問題集をダウンロード Professional-Machine-Learning-Engineer試験解説問題
- 最新のProfessional-Machine-Learning-Engineer資料勉強 - 合格スムーズ Professional-Machine-Learning-Engineer模擬試験最新版 | 一番優秀なProfessional-Machine-Learning-Engineer問題と解答 Google Professional Machine Learning Engineer □ ▷ www.passtest.jp ◁ を入力して “Professional-Machine-Learning-Engineer” を検索し、無料でダウンロードしてください Professional-Machine-Learning-Engineer問題数
- Professional-Machine-Learning-Engineer復習過去問 □ Professional-Machine-Learning-Engineer練習問題 □ Professional-Machine-Learning-Engineer問題例 □ 最新 ✨ Professional-Machine-Learning-Engineer □ ✨ □ 問題集ファイルは (www.goshiken.com) にて検索 Professional-Machine-Learning-Engineer日本語版問題集
- 最新のProfessional-Machine-Learning-Engineer資料勉強 - 合格スムーズ Professional-Machine-Learning-Engineer模擬試験最新版 | 一番優秀なProfessional-Machine-Learning-Engineer問題と解答 Google Professional Machine Learning Engineer □ ⇒ www.mogixam.com ⇐ サイトにて ✨ Professional-Machine-Learning-Engineer □ ✨ □ 問題集を無料で使おう Professional-Machine-Learning-Engineer試験問題
- Professional-Machine-Learning-Engineer合格記 □ Professional-Machine-Learning-Engineer試験関連赤本 □ Professional-Machine-Learning-Engineer練習問題 □ 《 www.goshiken.com 》 にて限定無料の □ Professional-Machine-Learning-Engineer □ 問題集をダウンロードせよ Professional-Machine-Learning-Engineer PDF問題サンプル
- 最新のProfessional-Machine-Learning-Engineer資料勉強 - 合格スムーズ Professional-Machine-Learning-Engineer模擬試験最新版 | 一番優秀なProfessional-Machine-Learning-Engineer問題と解答 Google Professional Machine Learning Engineer □ ✨ Professional-Machine-Learning-Engineer □ ✨ □ を無料でダウンロード 【 www.xhs1991.com 】 ウェブサイトを入力するだけ Professional-Machine-Learning-Engineer受験内容
- Professional-Machine-Learning-Engineer PDF □ Professional-Machine-Learning-Engineer試験対策 □ Professional-Machine-Learning-Engineer試験攻略 □ ✨ www.goshiken.com □ ✨ □ サイトにて ➡ Professional-Machine-Learning-Engineer □ 問題集を無料で使おう Professional-Machine-Learning-Engineer練習問題
- Professional-Machine-Learning-Engineer日本語学習内容 □ Professional-Machine-Learning-Engineer受験準備 □ Professional-Machine-Learning-Engineer試験関連赤本 □ “ www.passtest.jp ” で ➡ Professional-Machine-Learning-Engineer □ を検索して、無料でダウンロードしてください Professional-Machine-Learning-Engineer PDF
- 初段のProfessional-Machine-Learning-Engineer資料勉強 - 有効的なGoogle認定トレーニング練習 Google Google Professional Machine Learning Engineer □ ウェブサイト { www.goshiken.com } を開き、 ➡ Professional-Machine-Learning-Engineer □ を検索して無料でダウンロードしてください Professional-Machine-Learning-Engineer日本語版トレーニング
- 素敵なProfessional-Machine-Learning-Engineer資料勉強 - 合格スムーズ Professional-Machine-Learning-Engineer模擬試験最新版 | 便利なProfessional-Machine-Learning-Engineer問題と解答 Google Professional Machine Learning Engineer □ ▷ www.xhs1991.com ◁ は、 ➡ Professional-Machine-Learning-Engineer □ □ □ を無料でダウンロードするのに最適なサイトです Professional-Machine-Learning-Engineer受験準備
- yourbookmarklist.com, xanderhurn993875.blogofchange.com, neleikai293915.wikilowdown.com, ebiz-directory.com, nicolehslly580015.blogproducer.com, bookmarklethq.com, www.jkkfg.com, halemaikai706512.blogunteer.com, socialnetworkadsinfo.com, honeymho402176.onzeblog.com, Disposable vapes

無料でクラウドストレージから最新のTech4Exam Professional-Machine-Learning-Engineer PDFダンプをダウンロードする: https://drive.google.com/open?id=1AJ8TjN7MoAMQwXTtFbXhOarxV23QN_B3