

Google Professional-Machine-Learning-Engineer Exam Quizzes: Google Professional Machine Learning Engineer & Leader in Qualification Exams



2026 Latest PassExamDumps Professional-Machine-Learning-Engineer PDF Dumps and Professional-Machine-Learning-Engineer Exam Engine Free Share: <https://drive.google.com/open?id=1P17XepqjeY9QSTu2J6eRnVZ6yl7dmRGa>

All questions in our Professional-Machine-Learning-Engineer pass guide are at here to help you prepare for the certification exam. We have developed our learning materials with accurate Professional-Machine-Learning-Engineer exam answers and detailed explanations to ensure you pass test in your first try. Our PDF files are printable that you can share your Professional-Machine-Learning-Engineer free demo with your friends and classmates. You can practice Professional-Machine-Learning-Engineer real questions and review our study guide anywhere and anytime.

The Google Professional-Machine-Learning-Engineer exam covers a range of topics, including data preparation and feature engineering, model training, optimization, and deployment. Candidates must demonstrate their ability to use machine learning tools and technologies to create scalable, efficient, and accurate models. Professional-Machine-Learning-Engineer Exam also tests their knowledge of best practices for machine learning, data processing, and model evaluation.

>> **Professional-Machine-Learning-Engineer Exam Quizzes** <<

Professional-Machine-Learning-Engineer Test Braindumps: Google Professional Machine Learning Engineer & Professional-Machine-Learning-Engineer VCE Dumps

Get benefits from PassExamDumps exam questions update offer and prepare well with the assistance of Google Professional-Machine-Learning-Engineer updated exam questions. The Google Professional-Machine-Learning-Engineer exam dumps are being offered at affordable charges. We guarantee you that the Professional-Machine-Learning-Engineer Exam Dumps prices are entirely affordable for every Professional-Machine-Learning-Engineer exam candidate.

Google Professional Machine Learning Engineer Sample Questions (Q270-Q275):

NEW QUESTION # 270

You developed a custom model by using Vertex AI to forecast the sales of your company's products based on historical transactional data. You anticipate changes in the feature distributions and the correlations between the features in the near future. You also expect to receive a large volume of prediction requests. You plan to use Vertex AI Model Monitoring for drift detection and you want to minimize the cost. What should you do?

- A. Use the features and the feature attributions for monitoring. Set a prediction-sampling-rate value that is closer to 0 than 1.
- B. Use the features and the feature attributions for monitoring. Set a monitoring-frequency value that is lower than the default.
- C. Use the features for monitoring. Set a prediction-sampling-rare value that is closer to 1 than 0.

- D. Use the features for monitoring Set a monitoring- frequency value that is higher than the default.

Answer: A

Explanation:

The best option for using Vertex AI Model Monitoring for drift detection and minimizing the cost is to use the features and the feature attributions for monitoring, and set a prediction-sampling-rate value that is closer to 0 than 1. This option allows you to leverage the power and flexibility of Google Cloud to detect feature drift in the input predict requests for custom models, and reduce the storage and computation costs of the model monitoring job. Vertex AI Model Monitoring is a service that can track and compare the results of multiple machine learning runs. Vertex AI Model Monitoring can monitor the model's prediction input data for feature skew and drift. Feature drift occurs when the feature data distribution in production changes over time. If the original training data is not available, you can enable drift detection to monitor your models for feature drift.

Vertex AI Model Monitoring uses TensorFlow Data Validation (TFDV) to calculate the distributions and distance scores for each feature, and compares them with a baseline distribution. The baseline distribution is the statistical distribution of the feature's values in the training data. If the training data is not available, the baseline distribution is calculated from the first 1000 prediction requests that the model receives. If the distance score for a feature exceeds an alerting threshold that you set, Vertex AI Model Monitoring sends you an email alert. However, if you use a custom model, you can also enable feature attribution monitoring, which can provide more insights into the feature drift. Feature attribution monitoring analyzes the feature attributions, which are the contributions of each feature to the prediction output. Feature attribution monitoring can help you identify the features that have the most impact on the model performance, and the features that have the most significant drift over time. Feature attribution monitoring can also help you understand the relationship between the features and the prediction output, and the correlation between the features¹. The prediction-sampling-rate is a parameter that determines the percentage of prediction requests that are logged and analyzed by the model monitoring job. Using a lower prediction-sampling-rate can reduce the storage and computation costs of the model monitoring job, but also the quality and validity of the data. Using a lower prediction-sampling-rate can introduce sampling bias and noise into the data, and make the model monitoring job miss some important features or patterns of the data. However, using a higher prediction-sampling-rate can increase the storage and computation costs of the model monitoring job, and also the amount of data that needs to be processed and analyzed. Therefore, there is a trade-off between the prediction-sampling-rate and the cost and accuracy of the model monitoring job, and the optimal prediction-sampling-rate depends on the business objective and the data characteristics². By using the features and the feature attributions for monitoring, and setting a prediction-sampling-rate value that is closer to 0 than 1, you can use Vertex AI Model Monitoring for drift detection and minimize the cost.

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

* Option A: Using the features for monitoring and setting a monitoring-frequency value that is higher than the default would not enable feature attribution monitoring, and could increase the cost of the model monitoring job. The monitoring-frequency is a parameter that determines how often the model monitoring job analyzes the logged prediction requests and calculates the distributions and distance scores for each feature. Using a higher monitoring-frequency can increase the frequency and timeliness of the model monitoring job, but also the computation costs of the model monitoring job. Moreover, using the features for monitoring would not enable feature attribution monitoring, which can provide more insights into the feature drift and the model performance¹.

* Option B: Using the features for monitoring and setting a prediction-sampling-rate value that is closer to 1 than 0 would not enable feature attribution monitoring, and could increase the cost of the model monitoring job. The prediction-sampling-rate is a parameter that determines the percentage of prediction requests that are logged and analyzed by the model monitoring job. Using a higher prediction-sampling-rate can increase the quality and validity of the data, but also the storage and computation costs of the model monitoring job. Moreover, using the features for monitoring would not enable feature attribution monitoring, which can provide more insights into the feature drift and the model performance^{1,2}.

* Option C: Using the features and the feature attributions for monitoring and setting a monitoring-frequency value that is lower than the default would enable feature attribution monitoring, but could reduce the frequency and timeliness of the model monitoring job. The monitoring-frequency is a parameter that determines how often the model monitoring job analyzes the logged prediction requests and calculates the distributions and distance scores for each feature. Using a lower monitoring-frequency can reduce the computation costs of the model monitoring job, but also the frequency and timeliness of the model monitoring job. This can make the model monitoring job less responsive and effective in detecting and alerting the feature drift¹.

References:

- * Preparing for Google Cloud Certification: Machine Learning Engineer, Course 3: Production ML Systems, Week 4: Evaluation
- * Google Cloud Professional Machine Learning Engineer Exam Guide, Section 3: Scaling ML models in production, 3.3 Monitoring ML models in production
- * Official Google Cloud Certified Professional Machine Learning Engineer Study Guide, Chapter 6: Production ML Systems, Section 6.3: Monitoring ML Models
- * Using Model Monitoring
- * Understanding the score threshold slider

NEW QUESTION # 271

You have recently trained a scikit-learn model that you plan to deploy on Vertex AI. This model will support both online and batch

prediction. You need to preprocess input data for model inference. You want to package the model for deployment while minimizing additional code. What should you do?

- A. 1 Wrap your model in a custom prediction routine (CPR), and build a container image from the CPR local model
2 Upload your sci-kit learn model container to Vertex AI Model Registry
3 Deploy your model to Vertex AI Endpoints, and create a Vertex AI batch prediction job
- B. 1. Create a custom container for your sci-kit learn model,
2 Define a custom serving function for your model
3 Upload your model and custom container to Vertex AI Model Registry
4 Deploy your model to Vertex AI Endpoints, and create a Vertex AI batch prediction job
- C. 1 Create a custom container for your sci-kit learn model.
2 Upload your model and custom container to Vertex AI Model Registry
3 Deploy your model to Vertex AI Endpoints, and create a Vertex AI batch prediction job that uses the `instanceConfig.instanceType` setting to transform your input data
- D. 1 Upload your model to the Vertex AI Model Registry by using a prebuilt scikit-learn prediction container
2 Deploy your model to Vertex AI Endpoints, and create a Vertex AI batch prediction job that uses the `instanceConfig.instanceType` setting to transform your input data

Answer: A

Explanation:

The best option for deploying a scikit-learn model on Vertex AI with minimal additional code is to wrap the model in a custom prediction routine (CPR) and build a container image from the CPR local model. Upload your scikit-learn model container to Vertex AI Model Registry. Deploy your model to Vertex AI Endpoints, and create a Vertex AI batch prediction job. This option allows you to leverage the power and simplicity of Google Cloud to deploy and serve a scikit-learn model that supports both online and batch prediction. Vertex AI is a unified platform for building and deploying machine learning solutions on Google Cloud. Vertex AI can deploy a trained scikit-learn model to an online prediction endpoint, which can provide low-latency predictions for individual instances. Vertex AI can also create a batch prediction job, which can provide high-throughput predictions for a large batch of instances. A custom prediction routine (CPR) is a Python script that defines the logic for preprocessing the input data, running the prediction, and postprocessing the output data. A CPR can help you customize the prediction behavior of your model, and handle complex or non-standard data formats. A CPR can also help you minimize the additional code, as you only need to write a few functions to implement the prediction logic. A container image is a package that contains the model, the CPR, and the dependencies. A container image can help you standardize and simplify the deployment process, as you only need to upload the container image to Vertex AI Model Registry, and deploy it to Vertex AI Endpoints. By wrapping the model in a CPR and building a container image from the CPR local model, uploading the scikit-learn model container to Vertex AI Model Registry, deploying the model to Vertex AI Endpoints, and creating a Vertex AI batch prediction job, you can deploy a scikit-learn model on Vertex AI with minimal additional code¹.

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

Option A: Uploading your model to the Vertex AI Model Registry by using a prebuilt scikit-learn prediction container, deploying your model to Vertex AI Endpoints, and creating a Vertex AI batch prediction job that uses the `instanceConfig.instanceType` setting to transform your input data would not allow you to preprocess the input data for model inference, and could cause errors or poor performance. A prebuilt scikit-learn prediction container is a container image that is provided by Google Cloud, and contains the scikit-learn framework and the dependencies. A prebuilt scikit-learn prediction container can help you deploy a scikit-learn model without writing any code, but it also limits your customization options. A prebuilt scikit-learn prediction container can only handle standard data formats, such as JSON or CSV, and cannot perform any preprocessing or postprocessing on the input or output data. If your input data requires any transformation or normalization before running the prediction, you cannot use a prebuilt scikit-learn prediction container. The `instanceConfig.instanceType` setting is a parameter that determines the machine type and the accelerator type for the batch prediction job. The `instanceConfig.instanceType` setting can help you optimize the performance and the cost of the batch prediction job, but it cannot help you transform your input data².

Option C: Creating a custom container for your scikit-learn model, defining a custom serving function for your model, uploading your model and custom container to Vertex AI Model Registry, and deploying your model to Vertex AI Endpoints, and creating a Vertex AI batch prediction job would require more skills and steps than using a CPR and a container image. A custom container is a container image that contains the model, the dependencies, and a web server. A custom container can help you customize the prediction behavior of your model, and handle complex or non-standard data formats. A custom serving function is a Python function that defines the logic for running the prediction on the model. A custom serving function can help you implement the prediction logic of your model, and handle complex or non-standard data formats. However, creating a custom container and defining a custom serving function would require more skills and steps than using a CPR and a container image. You would need to write code, build and test the container image, configure the web server, and implement the prediction logic. Moreover, creating a custom container and defining a custom serving function would not allow you to preprocess the input data for model inference, as the custom serving function only runs the prediction on the model³.

Option D: Creating a custom container for your scikit-learn model, uploading your model and custom container to Vertex AI Model

Registry, deploying your model to Vertex AI Endpoints, and creating a Vertex AI batch prediction job that uses the `instanceConfig.instanceType` setting to transform your input data would not allow you to preprocess the input data for model inference, and could cause errors or poor performance. A custom container is a container image that contains the model, the dependencies, and a web server. A custom container can help you customize the prediction behavior of your model, and handle complex or non-standard data formats. However, creating a custom container would require more skills and steps than using a CPR and a container image. You would need to write code, build and test the container image, and configure the web server. The `instanceConfig.instanceType` setting is a parameter that determines the machine type and the accelerator type for the batch prediction job. The `instanceConfig.instanceType` setting can help you optimize the performance and the cost of the batch prediction job, but it cannot help you transform your input data²³.

Reference:

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 3: Scaling ML models in production, 3.1 Deploying ML models to production Official Google Cloud Certified Professional Machine Learning Engineer Study Guide, Chapter 6: Production ML Systems, Section 6.2: Serving ML Predictions Custom prediction routines Using pre-built containers for prediction Using custom containers for prediction

NEW QUESTION # 272

A trucking company is collecting live image data from its fleet of trucks across the globe. The data is growing rapidly and approximately 100 GB of new data is generated every day. The company wants to explore machine learning uses cases while ensuring the data is only accessible to specific IAM users.

Which storage option provides the most processing flexibility and will allow access control with IAM?

- A. Use an Amazon S3-backed data lake to store the raw images, and set up the permissions using bucket policies.
- B. Configure Amazon EFS with IAM policies to make the data available to Amazon EC2 instances owned by the IAM users.
- **C. Setup up Amazon EMR with Hadoop Distributed File System (HDFS) to store the files, and restrict access to the EMR instances using IAM policies.**
- D. Use a database, such as Amazon DynamoDB, to store the images, and set the IAM policies to restrict access to only the desired IAM users.

Answer: C

Explanation:

Explanation

NEW QUESTION # 273

You need to develop an image classification model by using a large dataset that contains labeled images in a Cloud Storage Bucket. What should you do?

- A. Use Vertex AI Pipelines with TensorFlow Extended (TFX) to create a pipeline that reads the images from Cloud Storage and trams the model.
- **B. Import the labeled images as a managed dataset in Vertex AI: and use AutoML to tram the model.**
- C. Use Vertex AI Pipelines with the Kubeflow Pipelines SDK to create a pipeline that reads the images from Cloud Storage and trains the model.
- D. Convert the image dataset to a tabular format using Dataflow Load the data into BigQuery and use BigQuery ML to tram the model.

Answer: B

Explanation:

The best option for developing an image classification model by using a large dataset that contains labeled images in a Cloud Storage bucket is to import the labeled images as a managed dataset in Vertex AI and use AutoML to train the model. This option allows you to leverage the power and simplicity of Google Cloud to create and deploy a high-quality image classification 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 create a managed dataset from a Cloud Storage bucket that contains labeled images, which can be used to train an AutoML model. AutoML is a service that can automatically build and optimize machine learning models for various tasks, such as image classification, object detection, natural language processing, and tabular data analysis. AutoML can handle the complex aspects of machine learning, such as feature engineering, model architecture, hyperparameter tuning, and model evaluation. AutoML can also evaluate, deploy, and monitor the image classification model, and provide online or batch predictions. By using Vertex AI and AutoML, users can develop an image classification model by using a large dataset with ease and efficiency.

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

* Option A: Using Vertex AI Pipelines with the Kubeflow Pipelines SDK to create a pipeline that reads the images from Cloud Storage and trains the model would require more skills and steps than using Vertex AI and AutoML. 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. Kubeflow Pipelines SDK is a Python library that can create and run pipelines on Vertex AI Pipelines or on Kubeflow, an open-source platform for machine learning on Kubernetes. However, using Vertex AI Pipelines and Kubeflow Pipelines SDK would require writing code, building Docker images, defining pipeline components and steps, and managing the pipeline execution and artifacts. Moreover, Vertex AI Pipelines and Kubeflow Pipelines SDK are not specialized for image classification, and users would need to use other libraries or frameworks, such as TensorFlow or PyTorch, to build and train the image classification model.

* Option B: Using Vertex AI Pipelines with TensorFlow Extended (TFX) to create a pipeline that reads the images from Cloud Storage and trains the model would require more skills and steps than using Vertex AI and AutoML. TensorFlow Extended (TFX) is a framework that can create and run end-to-end machine learning pipelines on TensorFlow, a popular library for building and training deep learning models. TFX can preprocess the data, train and evaluate the model, validate and push the model, and serve the model for online or batch predictions. However, using Vertex AI Pipelines and TFX would

* require writing code, building Docker images, defining pipeline components and steps, and managing the pipeline execution and artifacts. Moreover, TFX is not optimized for image classification, and users would need to use other libraries or tools, such as TensorFlow Data Validation, TensorFlow Transform, and TensorFlow Hub, to handle the image data and the model architecture.

* Option D: Converting the image dataset to a tabular format using Dataflow, loading the data into BigQuery, and using BigQuery ML to train the model would not handle the image data properly and could result in a poor model performance. Dataflow is a service that can create scalable and reliable pipelines to process large volumes of data from various sources. Dataflow can preprocess the data by using Apache Beam, a programming model for defining and executing data processing workflows. BigQuery is a serverless, scalable, and cost-effective data warehouse that can perform fast and interactive queries on large datasets. BigQuery ML is a service that can create and train machine learning models by using SQL queries on BigQuery. However, converting the image data to a tabular format would lose the spatial and semantic information of the images, which are essential for image classification. Moreover, BigQuery ML is not specialized for image classification, and users would need to use other tools or techniques, such as feature hashing, embedding, or one-hot encoding, to handle the categorical features.

NEW QUESTION # 274

You are creating a model training pipeline to predict sentiment scores from text-based product reviews. You want to have control over how the model parameters are tuned, and you will deploy the model to an endpoint after it has been trained. You will use Vertex AI Pipelines to run the pipeline. You need to decide which Google Cloud pipeline components to use. What components should you choose?

- A.
- B.
- C.
- D.

Answer: C

Explanation:

According to the web search results, Vertex AI Pipelines is a serverless orchestrator for running ML pipelines, using either the KFP SDK or TFX1. Vertex AI Pipelines provides a set of prebuilt components that can be used to perform common ML tasks, such as training, evaluation, deployment, and more2. Vertex AI ModelEvaluationOp and ModelDeployOp are two such components that can be used to evaluate and deploy a model to an endpoint for online inference3. However, Vertex AI Pipelines does not provide a prebuilt component for hyperparameter tuning. Therefore, to have control over how the model parameters are tuned, you need to use a custom component that calls the Vertex AI HyperparameterTuningJob service4. Therefore, option A is the best way to decide which Google Cloud pipeline components to use for the given use case, as it includes a custom component for hyperparameter tuning, and prebuilt components for model evaluation and deployment. The other options are not relevant or optimal for this scenario. References:

* Vertex AI Pipelines

* Google Cloud Pipeline Components

* Vertex AI ModelEvaluationOp and ModelDeployOp

* Vertex AI HyperparameterTuningJob

* Google Professional Machine Learning Certification Exam 2023

* Latest Google Professional Machine Learning Engineer Actual Free Exam Questions

