Free Questions for Professional-Machine-Learning-Engineer

Shared by Moran on 12-12-2023

For More Free Questions and Preparation Resources

Check the Links on Last Page

Question 1

Question Type: MultipleChoice

You work at an ecommerce startup. You need to create a customer churn prediction model Your company's recent sales records are stored in a BigQuery table You want to understand how your initial model is making predictions. You also want to iterate on the model as quickly as possible while minimizing cost How should you build your first model?

Options:

A- Export the data to a Cloud Storage Bucket Load the data into a pandas DataFrame on Vertex AI Workbench and train a logistic regression model with scikit-learn.
B- Create a tf.data.Dataset by using the TensorFlow BigQueryChent Implement a deep neural network in TensorFlow.
C- Prepare the data in BigQuery and associate the data with a Vertex AI dataset Create an AutoMLTabuiarTrainmgJob to train a classification model.
D- Export the data to a Cloud Storage Bucket Create tf. data. Dataset to read the data from Cloud Storage Implement a deep neural network in TensorFlow.

Answer:

C

Explanation:

BigQuery is a service that allows you to store and query large amounts of data in a scalable and cost-effective way. You can use BigQuery to prepare the data for your customer churn prediction model, such as filtering, aggregating, and transforming the data. You can then associate the data with a Vertex AI dataset, which is a service that allows you to store and manage your ML data on Google Cloud. By using a Vertex AI dataset, you can easily access the data from other Vertex AI services, such as AutoML. AutoML is a service that allows you to create and train ML models without writing code. You can use AutoML to create an AutoMLTabularTrainingJob, which is a type of job that trains a classification model for tabular data, such as customer churn. By using an AutoMLTabularTrainingJob, you can benefit from the automated feature engineering, model selection, and hyperparameter tuning that AutoML provides. You can also use Vertex Explainable AI to understand how your model is making predictions, such as which features are most important and how they affect the prediction outcome. By using BigQuery, Vertex AI dataset, and AutoMLTabularTrainingJob, you can build your first model as quickly as possible while minimizing cost and complexity.Reference:

BigQuery documentation

Vertex AI dataset documentation

AutoMLTabularTrainingJob documentation

Vertex Explainable AI documentation

Preparing for Google Cloud Certification: Machine Learning Engineer Professional Certificate

Question 2

Question Type: MultipleChoice

You are developing an ML model to identify your company s products in images. You have access to over one million images in a Cloud Storage bucket. You plan to experiment with different TensorFlow models by using Vertex AI Training You need to read images at scale during training while minimizing data I/O bottlenecks What should you do?

Options:

A- Load the images directly into the Vertex AI compute nodes by using Cloud Storage FUSE Read the images by using the tf .data.Dataset.from_tensor_slices function.
B- Create a Vertex AI managed dataset from your image data Access the aip_training_data_uri environment variable to read the images by using the tf. data. Dataset. list_flies function.
C- Convert the images to TFRecords and store them in a Cloud Storage bucket Read the TFRecords by using the tf. ciata.TFRecordDataset function.
D- Store the URLs of the images in a CSV file Read the file by using the tf.data.experomental.CsvDataset function.

Answer:

C

Explanation:

TFRecords are a binary file format that can store large amounts of data efficiently. By converting the images to TFRecords and storing them in a Cloud Storage bucket, you can reduce the data size and improve the data transfer speed. You can then read the TFRecords by using the tf.data.TFRecordDataset function, which creates a dataset of tensors from the TFRecord files. This way, you can read images at scale during training while minimizing data I/O bottlenecks.Reference:

TFRecord documentation

tf.data.TFRecordDataset documentation

Preparing for Google Cloud Certification: Machine Learning Engineer Professional Certificate

Question 3

Question Type: MultipleChoice

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?

Options:

A- 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.
B- 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 tnggered redeploy the pipeline.
C- 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.
D- 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.

Answer:

B

Explanation:

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.Reference:

Vertex AI documentation

Vertex AI endpoint documentation

Model monitoring documentation

Preparing for Google Cloud Certification: Machine Learning Engineer Professional Certificate

Question 4

Question Type: MultipleChoice

You need to use TensorFlow to train an image classification model. Your dataset is located in a Cloud Storage directory and contains millions of labeled images Before training the model, you need to prepare the dat

a. You want the data preprocessing and model training workflow to be as efficient scalable, and low maintenance as possible. What should you do?

Options:

A- 1 Create a Dataflow job that creates sharded TFRecord files in a Cloud Storage directory.
2 Reference tf .data.TFRecordDataset in the training script.
3. Train the model by using Vertex AI Training with a V100 GPU.
B- 1 Create a Dataflow job that moves the images into multiple Cloud Storage directories, where each directory is named according to the corresponding label.
2 Reference tfds.fclder_da-asst.imageFclder in the training script.
3. Train the model by using Vertex AI Training with a V100 GPU.
C- 1 Create a Jupyter notebook that uses an n1-standard-64, V100 GPU Vertex AI Workbench instance.
2 Write a Python script that creates sharded TFRecord files in a directory inside the instance
3. Reference tf. da-a.TFRecrrdDataset in the training script.
4. Train the model by using the Workbench instance.
D- 1 Create a Jupyter notebook that uses an n1-standard-64, V100 GPU Vertex AI Workbench instance.
2 Write a Python scnpt that copies the images into multiple Cloud Storage directories, where

each directory is named according to the corresponding label.
3 Reference tf ds. f older_dataset. imageFolder in the training script.
4. Train the model by using the Workbench instance.

Answer:

A

Explanation:

TFRecord is a binary file format that stores your data as a sequence of binary strings1.TFRecord files are efficient, scalable, and easy to process1.Sharding is a technique that splits a large file into smaller files, which can improve parallelism and performance2.Dataflow is a service that allows you to create and run data processing pipelines on Google Cloud3.Dataflow can create sharded TFRecord files from your images in a Cloud Storage directory4.

tf.data.TFRecordDataset is a class that allows you to read and parse TFRecord files in TensorFlow. You can use this class to create a tf.data.Dataset object that represents your input data for training. tf.data.Dataset is a high-level API that provides various methods to transform, batch, shuffle, and prefetch your data.

Vertex AI Training is a service that allows you to train your custom models on Google Cloud using various hardware accelerators, such as GPUs. Vertex AI Training supports TensorFlow models and can read data from Cloud Storage. You can use Vertex AI Training to train your image classification model by using a V100 GPU, which is a powerful and fast GPU for deep learning.

TFRecord and tf.Example | TensorFlow Core

Sharding | TensorFlow Core

Dataflow | Google Cloud

Creating sharded TFRecord files | Google Cloud

[tf.data.TFRecordDataset | TensorFlow Core v2.6.0]

[tf.data: Build TensorFlow input pipelines | TensorFlow Core]

[Vertex AI Training | Google Cloud]

[NVIDIA Tesla V100 GPU | NVIDIA]

Question 5

Question Type: MultipleChoice

You work for a semiconductor manufacturing company. You need to create a real-time application that automates the quality control process High-definition images of each semiconductor are taken at the end of the assembly line in real time. The photos are uploaded to a Cloud Storage bucket along with tabular data that includes each semiconductor's batch number serial number dimensions, and weight You need to configure model training and serving while maximizing model accuracy. What should you do?

Options:

A- Use Vertex Al Data Labeling Service to label the images and train an AutoML image classification model.
Deploy the model and configure Pub/Sub to publish a message when an image is categorized into the failing class.
B- Use Vertex Al Data Labeling Service to label the images and train an AutoML image classification model. Schedule a daily batch prediction job that publishes a Pub/Sub message when the job completes.
C- Convert the images into an embedding representation Import this data into BigQuery, and train a BigQuery. ML K-means clustenng model with two clusters Deploy the model and configure Pub/Sub to publish a message when a semiconductor's data is categorized into the failing cluster.
D- Import the tabular data into BigQuery use Vertex Al Data Labeling Service to label the data and train an AutoML tabular classification model Deploy the model and configure Pub/Sub to publish a message when a semiconductor's data is categorized into the failing class.

Answer:

A

Explanation:

Vertex AI is a unified platform for building and managing machine learning solutions on Google Cloud. It provides various services and tools for different stages of the machine learning lifecycle, such as data preparation, model training, deployment, monitoring, and experimentation. Vertex AI Data Labeling Service is a service that allows you to create and manage human-labeled datasets for machine learning. You can use Vertex AI Data Labeling Service to label the images of semiconductors with binary labels, such as ''pass'' or ''fail'', based on the quality criteria. You can also use Vertex AI AutoML Image Classification, which is a service that allows you to create and train custom image classification models without writing any code. You can use Vertex AI AutoML Image Classification to train an image classification model on the labeled images of semiconductors, and optimize the model for accuracy. You can also use Vertex AI to deploy the model to an endpoint, which is a service that allows you to serve online predictions from your model. You can configure Pub/Sub, which is a service that allows you to publish and subscribe to messages, to publish a message when an image is categorized into the failing class by the

model. You can use the message to trigger an action, such as alerting the quality control team or stopping the production line. This solution can help you create a real-time application that automates the quality control process of semiconductors, and maximizes the model accuracy.Reference: The answer can be verified from official Google Cloud documentation and resources related to Vertex AI, Vertex AI Data Labeling Service, Vertex AI AutoML Image Classification, and Pub/Sub.

Vertex AI | Google Cloud

Vertex AI Data Labeling Service | Google Cloud

Vertex AI AutoML Image Classification | Google Cloud

Pub/Sub | Google Cloud

P2P exams

Question 6

Question Type: MultipleChoice

Your team frequently creates new ML models and runs experiments. Your team pushes code to a single repository hosted on Cloud Source Repositories. You want to create a continuous integration pipeline that automatically retrains the models whenever there is any modification of the code. What should be your first step to set up the CI pipeline?

Options:

A- Configure a Cloud Build trigger with the event set as 'Pull Request'
B- Configure a Cloud Build trigger with the event set as 'Push to a branch'
C- Configure a Cloud Function that builds the repository each time there is a code change.
D- Configure a Cloud Function that builds the repository each time a new branch is created.

P2P exams

Answer:

B

Explanation:

According to the web search results, Cloud Build1is a service that executes your builds on Google Cloud Platform infrastructure.Cloud Build can import source code from Cloud Source Repositories2, Cloud Storage, GitHub, Bitbucket, or any publicly hosted Git repository. Cloud Build allows you to create and manage build triggers, which are automated workflows that run whenever a code change is pushed to your source repository. You can use Cloud Build triggers to

automatically retrain your ML models whenever there is any modification of the code. Therefore, option B is the best way to set up the CI pipeline for the given use case, as it allows you to configure a Cloud Build trigger with the event set as ''Push to a branch'', which means the trigger will run whenever a new commit is pushed to a specific branch of your source repository. The other options are not relevant or optimal for this scenario.Reference:

Cloud Build

Cloud Source Repositories

Google Professional Machine Learning Certification Exam 2023

Latest Google Professional Machine Learning Engineer Actual Free Exam Questions

Question 7

Question Type: MultipleChoice

You work for a company that sells corporate electronic products to thousands of businesses worldwide. Your company stores historical customer data in BigQuery. You need to build a model that predicts customer lifetime value over the next three years. You want to use the simplest approach to build the model and you want to have access to visualization tools. What should you do?

Options:

A- Create a Vertex AI Workbench notebook to perform exploratory data analysis. Use IPython magics to create a new BigQuery table with input features Use the BigQuery console to run the create model statement Validate the results by using the ml. evaluate and ml. predict statements.
B- Run the create model statement from the BigQuery console to create an AutoML model Validate the results by using the ml. evaluate and ml. predict statements.
C- Create a Vertex AI Workbench notebook to perform exploratory data analysis and create input features Save the features as a CSV file in Cloud Storage Import the CSV file as a new BigQuery table Use the BigQuery console to run the create model statement Validate the results by using the ml. evaluate and ml. predict statements.
D- Create a Vertex AI Workbench notebook to perform exploratory data analysis Use IPython magics to create a new BigQuery table with input features, create the model and validate the results by using the create model, ml. evaluates, and ml. predict statements.

Answer:

B

Explanation:

BigQuery is a service that allows you to store and query large amounts of data in a scalable and cost-effective way. You can use BigQuery to build a model that predicts customer lifetime value over the next three years, by using the create model statement. The create model statement is a SQL command that allows you to create and train an ML model using your data in BigQuery. You can use the create model statement to create an AutoML model, which is a type of model that automatically selects the best features and architecture for your data. By using an AutoML model, you can use the simplest approach to build the model, without writing any code or performing any feature engineering. You can also use the ml.evaluate and ml.predict statements to validate the results of your model. The ml.evaluate statement is a SQL command that allows you to evaluate the performance and quality of your model using various metrics. The ml.predict statement is a SQL command that allows you to make predictions using your model and new data. You can also use the BigQuery console to access visualization tools, such as charts and graphs, to explore and analyze your data and model results. By using the BigQuery console, the create model statement, and the ml.evaluate and ml.predict statements, you can build and validate a model that predicts customer lifetime value over the next three years, and have access to visualization tools.Reference:

BigQuery documentation

create model statement documentation

ml.evaluate statement documentation

ml.predict statement documentation

Preparing for Google Cloud Certification: Machine Learning Engineer Professional Certificate

Question 8

Question Type: MultipleChoice

You work for a delivery company. You need to design a system that stores and manages features such as parcels delivered and truck locations over time. The system must retrieve the features with low latency and feed those features into a model for online prediction. The data science team will retrieve historical data at a specific point in time for model training. You want to store the features with minimal effort. What should you do?

Options:

A- Store features in Bigtable as key/value data.
B- Store features in Vertex AI Feature Store.
C- Store features as a Vertex AI dataset and use those features to tram the models hosted in Vertex AI endpoints.
D- Store features in BigQuery timestamp partitioned tables, and use the BigQuery Storage Read API to serve the features.

Answer:

B

Explanation:

Vertex AI Feature Store is a service that allows you to store and manage your ML features on Google Cloud. You can use Vertex AI Feature Store to store features such as parcels delivered and truck locations over time, and retrieve them with low latency for online prediction. Online prediction is a type of prediction that provides low-latency responses to individual or small batches of input data. You can also use Vertex AI Feature Store to retrieve historical data at a specific point in time for model training. Model training is a process of learning the parameters of a ML model from data. By using Vertex AI Feature Store, you can store the features with minimal effort, and avoid the complexity of managing your own data storage and serving system.Reference:

Vertex AI Feature Store documentation

Preparing for Google Cloud Certification: Machine Learning Engineer Professional Certificate

Question 9

Question Type: MultipleChoice

You work for a bank You have been asked to develop an ML model that will support loan application decisions. You need to determine which Vertex AI services to include in the workflow You want to track the model's training parameters and the metrics per training epoch. You plan to compare the performance of each version of the model to determine the best model based on your chosen metrics. Which Vertex AI services should you use?

Options:

A- Vertex ML Metadata Vertex AI Feature Store, and Vertex AI Vizier
B- Vertex AI Pipelines. Vertex AI Experiments, and Vertex AI Vizier

C- Vertex ML Metadata Vertex Al Experiments, and Vertex Al TensorBoard
D- Vertex Al Pipelines. Vertex Al Feature Store, and Vertex Al TensorBoard

Answer:

C

Explanation:

According to the official exam guide1, one of the skills assessed in the exam is to ''track the lineage of pipeline artifacts''.Vertex ML Metadata2is a service that allows you to store, query, and visualize metadata associated with your ML workflows, such as datasets, models, metrics, and executions. Vertex ML Metadata helps you track the provenance and lineage of your ML artifacts and understand the relationships between them.Vertex AI Experiments3is a service that allows you to track and compare the results of your model training runs. Vertex AI Experiments automatically logs metadata such as hyperparameters, metrics, and artifacts for each training run. You can use Vertex AI Experiments to train your custom model using TensorFlow, PyTorch, XGBoost, or scikit-learn.Vertex AI TensorBoard4is a service that allows you to visualize and monitor your ML experiments using TensorBoard, an open source tool for ML visualization. Vertex AI TensorBoard helps you track the model's training parameters and the metrics per training epoch, and compare the performance of each version of the model. Therefore, option C is the best way to determine which Vertex AI services to include in the workflow for the given use case. The other options are not relevant or optimal for this scenario.Reference:

Professional ML Engineer Exam Guide

Vertex ML Metadata

Vertex AI Experiments

Vertex AI TensorBoard

Google Professional Machine Learning Certification Exam 2023

Latest Google Professional Machine Learning Engineer Actual Free Exam Questions

Question 10

Question Type: MultipleChoice

You have developed an AutoML tabular classification model that identifies high-value customers who interact with your organization's website.

You plan to deploy the model to a new Vertex Al endpoint that will integrate with your website

application. You expect higher traffic to the website during

nights and weekends. You need to configure the model endpoint's deployment settings to minimize latency and cost. What should you do?

Options:

A- Configure the model deployment settings to use an n1-standard-32 machine type.
B- Configure the model deployment settings to use an n1-standard-4 machine type. Set the minReplicaCount value to 1 and the maxReplicaCount value to 8.
C- Configure the model deployment settings to use an n1-standard-4 machine type and a GPU accelerator. Set the minReplicaCount value to 1 and the maxReplicaCount value to 4.
D- Configure the model deployment settings to use an n1-standard-8 machine type and a GPU accelerator.

Answer:

B

Explanation:

Deploying a model to an endpoint in Vertex AI associates physical resources with the model so it can serve online predictions with low latency1.By configuring the model deployment settings to use an n1-standard-4 machine type and setting the minReplicaCount value to 1 and the maxReplicaCount value to 8, you can ensure that the model scales according to the traffic, thereby minimizing latency and cost1.The n1-standard-4 machine type provides a balance between computing power and cost, and the dynamic scaling allows the model to handle higher traffic during nights and weekends without incurring unnecessary costs during off-peak times

To Get Premium Files for Professional-Machine-Learning-Engineer Visit
https://www.p2pexams.com/products/professional-machine-learning-engineer

For More Free Questions Visit
https://www.p2pexams.com/google/pdf/professional-machine-learning-engineer

