Reliable Professional-Machine-Learning-Engineer Exam Papers, Exam Professional-Machine-Learning-Engineer Questions

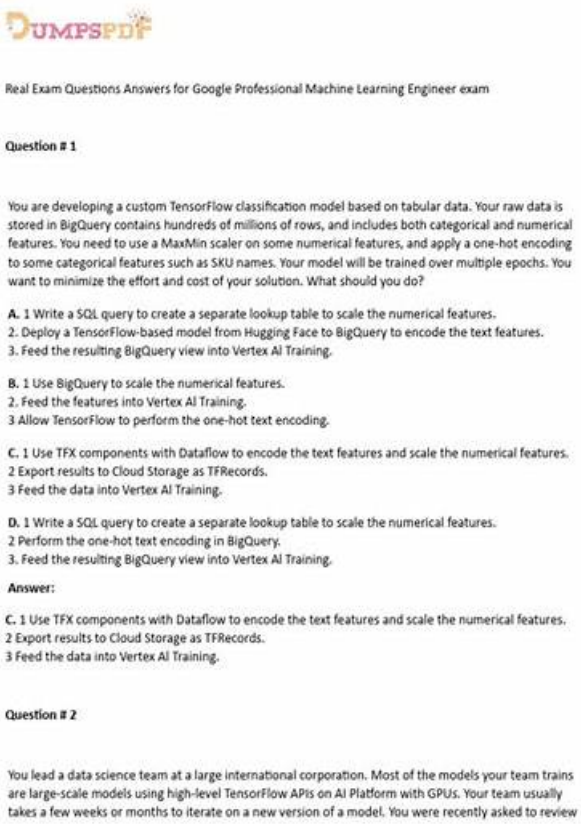

DUMPSPDF

Real Exam Questions Answers for Google Professional Machine Learning Engineer exam

Question # 1

You are developing a custom TensorFlow classification model based on tabular data. Your raw data is stored in BigQuery contains hundreds of millions of rows, and includes both categorical and numerical features. You need to use a MaxMin scaler on some numerical features, and apply a one-hot encoding to some categorical features such as SKU names. Your model will be trained over multiple epochs. You want to minimize the effort and cost of your solution. What should you do?

A. 1 Write a SQL query to create a separate lookup table to scale the numerical features.
2. Deploy a TensorFlow-based model from Hugging Face to BigQuery to encode the text features.
3. Feed the resulting BigQuery view into Vertex AI Training.

B. 1 Use BigQuery to scale the numerical features.
2. Feed the features into Vertex AI Training.
3 Allow TensorFlow to perform the one-hot text encoding.

C. 1 Use TFX components with Dataflow to encode the text features and scale the numerical features.
2 Export results to Cloud Storage as TFRecords.
3 Feed the data into Vertex AI Training.

D. 1 Write a SQL query to create a separate lookup table to scale the numerical features.
2 Perform the one-hot text encoding in BigQuery.
3. Feed the resulting BigQuery view into Vertex AI Training.

Answer:

C. 1 Use TFX components with Dataflow to encode the text features and scale the numerical features.
2 Export results to Cloud Storage as TFRecords.
3 Feed the data into Vertex AI Training.

Question # 2

You lead a data science team at a large international corporation. Most of the models your team trains are large-scale models using high-level TensorFlow APIs on AI Platform with GPUs. Your team usually takes a few weeks or months to iterate on a new version of a model. You were recently asked to review

P.S. Free & New Professional-Machine-Learning-Engineer dumps are available on Google Drive shared by PassTorrent: https://drive.google.com/open?id=1JkBstlUBBUr28wXvBt1F0F9J5ckWyBpo

The Google Professional Machine Learning Engineer (Professional-Machine-Learning-Engineer) questions is currently in use by many customers, and they are preparing for the test effectively. The applicants who used it previously to prepare for the Professional-Machine-Learning-Engineer certification exam have rated our Professional-Machine-Learning-Engineer Dumps as one of the best. Our customers receive Google Professional Machine Learning Engineer (Professional-Machine-Learning-Engineer) questions updates for up to 365 days after their purchase.

We don't just want to make profitable deals, but also to help our users pass the exams with the least amount of time to get Professional-Machine-Learning-Engineer certificate. Choosing our Professional-Machine-Learning-Engineer exam practice, you only need to spend 20-30 hours to prepare for the exam. Maybe you will ask whether such a short time can finish all the content, we want to tell you that you can rest assured ,because our Professional-Machine-Learning-Engineer Learning Materials are closely related to the exam outline and the questions of our Professional-Machine-Learning-Engineer guide questions are related to the latest and basic knowledge. You will pass the Professional-Machine-Learning-Engineer exam only with our Professional-Machine-Learning-Engineer exam questions.

>> Reliable Professional-Machine-Learning-Engineer Exam Papers <<

Exam Professional-Machine-Learning-Engineer Questions - Professional-Machine-Learning-Engineer Actual Test Answers

The above formats of PassTorrent are made to help customers prepare as per their unique styles and crack the Professional-Machine-Learning-Engineer exam certification on the very first attempt. Our Google Professional Machine Learning Engineer (Professional-Machine-Learning-Engineer) questions product is getting updated regularly as per the original Google Professional Machine Learning Engineer (Professional-Machine-Learning-Engineer) practice test's content. So that customers can prepare according to the latest Professional-Machine-Learning-Engineer exam content and pass it with ease.

Google Professional Machine Learning Engineer certification is highly respected in the industry and is recognized as a benchmark for excellence in machine learning. Achieving this certification demonstrates to employers and peers that a candidate has the skills and knowledge required to design, build, and deploy machine learning models on Google Cloud Platform. Google Professional Machine Learning Engineer certification is ideal for data scientists, machine learning engineers, software engineers, and other professionals who are looking to enhance their skills in machine learning and advance their career in this field.

Google Professional Machine Learning Engineer Sample Questions (Q60-Q65):

NEW QUESTION # 60
You are going to train a DNN regression model with Keras APIs using this code:
How many trainable weights does your model have? (The arithmetic below is correct.)

- A. 501*256+257*128+2 = 161154
- B. 500*256*0 25+256*128*0 25+128*2 = 40448
- C. 500*256+256*128+128*2 = 161024
- D. 501*256+257*128+128*2=161408

Answer: C

Explanation:
The number of trainable weights in a DNN regression model with Keras APIs can be calculated by multiplying the number of input units by the number of output units for each layer, and adding the number of bias units for each layer. The bias units are usually equal to the number of output units,except for the last layer, which does not have bias units if the activation function is softmax1. In this code, the model has three layers: a dense layer with 256 units and relu activation, a dropout layer with 0.25 rate, and a dense layer with 2 units and softmax activation. The input shape is 500. Therefore, the number of trainable weights is:
* For the first layer: 500 input units * 256 output units + 256 bias units = 128256
* For the second layer: The dropout layer does not have any trainable weights, as it only randomly sets some of the input units to zero to prevent overfitting2.
* For the third layer: 256 input units * 2 output units + 0 bias units = 512 The total number of trainable weights is 128256 + 512 = 161024. Therefore, the correct answer is B.
References:
* How to calculate the number of parameters for a Convolutional Neural Network?
* Dropout (keras.io)

NEW QUESTION # 61
Your data science team has requested a system that supports scheduled model retraining, Docker containers, and a service that supports autoscaling and monitoring for online prediction requests. Which platform components should you choose for this system?

- A. Cloud Composer, Vertex AI Training with custom containers, and App Engine
- B. Vertex AI Pipelines and App Engine
- C. Vertex AI Pipelines, Vertex AI Prediction, and Vertex AI Model Monitoring
- D. Cloud Composer, BigQuery ML, and Vertex AI Prediction

Answer: C

Explanation:
Option A is incorrect because Vertex AI Pipelines and App Engine do not meet all the requirements of the system. Vertex AI Pipelines is a service that allows you to create, run, and manage ML workflows using TensorFlow Extended (TFX) components or custom components1. App Engine is a service that allows you to build and deploy scalable web applications using standard or

flexible environments2. However, App Engine does not support Docker containers in the standard environment, and does not provide a dedicated service for online prediction and monitoring of ML models3.
Option B is correct because Vertex AI Pipelines, Vertex AI Prediction, and Vertex AI Model Monitoring meet all the requirements of the system. Vertex AI Prediction is a service that allows you to deploy and serve ML models for online or batch prediction, with support for autoscaling and custom containers4. Vertex AI Model Monitoring is a service that allows you to monitor the performance and fairness of your deployed models, and get alerts for any issues or anomalies5.
Option C is incorrect because Cloud Composer, BigQuery ML, and Vertex AI Prediction do not meet all the requirements of the system. Cloud Composer is a service that allows you to create, schedule, and manage workflows using Apache Airflow. BigQuery ML is a service that allows you to create and use ML models within BigQuery using SQL queries. However, BigQuery ML does not support custom containers, and Vertex AI Prediction does not support scheduled model retraining or model monitoring.
Option D is incorrect because Cloud Composer, Vertex AI Training with custom containers, and App Engine do not meet all the requirements of the system. Vertex AI Training is a service that allows you to train ML models using built-in algorithms or custom containers. However, Vertex AI Training does not support online prediction or model monitoring, and App Engine does not support Docker containers in the standard environment or online prediction and monitoring of ML models3.
Reference:
Vertex AI Pipelines overview
App Engine overview
Choosing an App Engine environment
Vertex AI Prediction overview
Vertex AI Model Monitoring overview
[Cloud Composer overview]
[BigQuery ML overview]
[BigQuery ML limitations]
[Vertex AI Training overview]

NEW QUESTION # 62
You developed an ML model with Al Platform, and you want to move it to production. You serve a few thousand queries per second and are experiencing latency issues. Incoming requests are served by a load balancer that distributes them across multiple Kubeflow CPU-only pods running on Google Kubernetes Engine (GKE). Your goal is to improve the serving latency without changing the underlying infrastructure. What should you do?

- A. Significantly increase the max_batch_size TensorFlow Serving parameter
- B. Switch to the tensorflow-model-server-universal version of TensorFlow Serving
- C. Recompile TensorFlow Serving using the source to support CPU-specific optimizations Instruct GKE to choose an appropriate baseline minimum CPU platform for serving nodes
- D. Significantly increase the max_enqueued_batches TensorFlow Serving parameter

Answer: C

Explanation:
TensorFlow Serving is a service that allows you to deploy and serve TensorFlow models in a scalable and efficient way. TensorFlow Serving supports various platforms and hardware, such as CPU, GPU, and TPU. However, the default TensorFlow Serving binaries are built with generic CPU instructions, which may not leverage the full potential of the CPU architecture. To improve the serving latency and performance, you can recompile TensorFlow Serving using the source code and enable CPU-specific optimizations, such as AVX, AVX2, and FMA1. These optimizations can speed up the computation and inference of the TensorFlow models, especially for deep neural networks.
Google Kubernetes Engine (GKE) is a service that allows you to run and manage containerized applications on Google Cloud using Kubernetes. GKE supports various types and sizes of nodes, which are the virtual machines that run the containers. GKE also supports different CPU platforms, which are the generations and models of the CPUs that power the nodes. GKE allows you to choose a baseline minimum CPU platform for your node pool, which is a group of nodes with the same configuration. By choosing a baseline minimum CPU platform, you can ensure that your nodes have the CPU features and capabilities that match your workload requirements2.
For the use case of serving a few thousand queries per second and experiencing latency issues, the best option is to recompile TensorFlow Serving using the source to support CPU-specific optimizations, and instruct GKE to choose an appropriate baseline minimum CPU platform for serving nodes. This option can improve the serving latency and performance without changing the underlying infrastructure, as it only involves rebuilding the TensorFlow Serving binary and selecting the CPU platform for the GKE nodes. This option can also take advantage of the CPU-only pods that are running on GKE, as it can optimize the CPU utilization and efficiency. Therefore, recompiling TensorFlow Serving using the source to support CPU-specific optimizations and instructing GKE to choose an appropriate baseline minimum CPU platform for serving nodes is the best option for this use case.
Reference:

Building TensorFlow Serving from source
Specifying a minimum CPU platform for a node pool

NEW QUESTION # 63
Your team is training a large number of ML models that use different algorithms, parameters and datasets.
Some models are trained in Vertex Ai Pipelines, and some are trained on Vertex Al Workbench notebook instances. Your team wants to compare the performance of the models across both services. You want to minimize the effort required to store the parameters and metrics What should you do?

- A. Implement all models in Vertex Al Pipelines Create a Vertex Al experiment, and associate all pipeline runs with that experiment.
- B. Create a Vertex Al experiment Submit all the pipelines as experiment runs. For models trained on notebooks log parameters and metrics by using the Vertex Al SDK.
- C. Implement an additional step for all the models running in pipelines and notebooks to export parameters and metrics to BigQuery.
- D. Store all model parameters and metrics as mode! metadata by using the Vertex Al Metadata API.

Answer: B

Explanation:
Vertex AI Experiments is a service that allows you to track, compare, and manage experiments with Vertex AI. You can use Vertex AI Experiments to record the parameters, metrics, and artifacts of each model training run, and compare them in a graphical interface. Vertex AI Experiments supports models trained in Vertex AI Pipelines, Vertex AI Custom Training, and Vertex AI Workbench notebooks. To use Vertex AI Experiments, you need to create an experiment and submit your pipeline runs or custom training jobs as experiment runs.
For models trained on notebooks, you need to use the Vertex AI SDK to log the parameters and metrics to the experiment. This way, you can minimize the effort required to store and compare the model performance across different services. References: Track, compare, manage experiments with Vertex AI Experiments, Vertex AI Pipelines: Metrics visualization and run comparison using the KFP SDK, [Vertex AI SDK for Python]

NEW QUESTION # 64
You work for a bank and are building a random forest model for fraud detection. You have a dataset that includes transactions, of which 1% are identified as fraudulent. Which data transformation strategy would likely improve the performance of your classifier?

- A. Write your data in TFRecords.
- B. Use one-hot encoding on all categorical features.
- C. Z-normalize all the numeric features.
- D. Oversample the fraudulent transaction 10 times.

Answer: D

NEW QUESTION # 65
.....

With our Professional-Machine-Learning-Engineer practice test software, you can simply assess yourself by going through the Professional-Machine-Learning-Engineer practice tests. We highly recommend going through the Professional-Machine-Learning-Engineer answers multiple times so you can assess your preparation for the Professional-Machine-Learning-Engineer exam. Make sure that you are preparing yourself for the Professional-Machine-Learning-Engineer test with our practice test software as it will help you get a clear idea of the real Professional-Machine-Learning-Engineer exam scenario. By passing the exams multiple times on practice test software, you will be able to pass the real Professional-Machine-Learning-Engineer test in the first attempt.

Exam Professional-Machine-Learning-Engineer Questions: https://www.passtorrent.com/Professional-Machine-Learning-Engineer-latest-torrent.html

- Reliable Professional-Machine-Learning-Engineer Exam Papers - High Pass Rate Guaranteed. □ Download □ Professional-Machine-Learning-Engineer □ for free by simply entering ➽ www.testkingpass.com □ website □ □Professional-Machine-Learning-Engineer Reliable Exam Question
- Reliable Professional-Machine-Learning-Engineer Exam Papers - High Pass Rate Guaranteed. □ Easily obtain 「

Professional-Machine-Learning-Engineer 」 for free download through ➠ www.pdfvce.com □ □Valid Professional-Machine-Learning-Engineer Exam Fee

- How to Crack Google Professional-Machine-Learning-Engineer Certification Exam Easily? □ Copy URL 《 www.pdfdumps.com 》 open and search for □ Professional-Machine-Learning-Engineer □ to download for free □Valid Professional-Machine-Learning-Engineer Exam Fee
- Professional-Machine-Learning-Engineer Reliable Dump □ Professional-Machine-Learning-Engineer Dumps Free □ Professional-Machine-Learning-Engineer Reliable Exam Question □ Go to website ➠ www.pdfvce.com □ open and search for 【 Professional-Machine-Learning-Engineer 】 to download for free □Professional-Machine-Learning-Engineer New Learning Materials
- Professional-Machine-Learning-Engineer Dumps Free □ Professional-Machine-Learning-Engineer Exam Sample □ Professional-Machine-Learning-Engineer Dumps Free □ Download ▹ Professional-Machine-Learning-Engineer ◃ for free by simply entering □ www.pdfdumps.com □ website □Professional-Machine-Learning-Engineer Pdf Free
- Valid Professional-Machine-Learning-Engineer Exam Fee □ Professional-Machine-Learning-Engineer Pdf Free □ Professional-Machine-Learning-Engineer Test Voucher □ Immediately open ✔ www.pdfvce.com □✔□ and search for ✔ Professional-Machine-Learning-Engineer □✔□ to obtain a free download □Professional-Machine-Learning-Engineer Latest Exam Experience
- Valid Professional-Machine-Learning-Engineer Exam Fee □ Professional-Machine-Learning-Engineer Vce Exam □ Professional-Machine-Learning-Engineer Reliable Exam Question □ Immediately open ✔ www.practicevce.com □✔□ and search for 「 Professional-Machine-Learning-Engineer 」 to obtain a free download ✶ Professional-Machine-Learning-Engineer Reliable Exam Question
- Professional-Machine-Learning-Engineer Free Pdf - Professional-Machine-Learning-Engineer Pass4sure Vce - Professional-Machine-Learning-Engineer Practice Torrent □ Download [Professional-Machine-Learning-Engineer] for free by simply searching on ✔ www.pdfvce.com □✔□ □Professional-Machine-Learning-Engineer Valid Test Fee
- Certification Professional-Machine-Learning-Engineer Questions ➡ Professional-Machine-Learning-Engineer Study Group □ Professional-Machine-Learning-Engineer Exam Sample □ Easily obtain free download of ⇛ Professional-Machine-Learning-Engineer ⇚ by searching on 《 www.pass4test.com 》 □Professional-Machine-Learning-Engineer Vce Exam
- Professional-Machine-Learning-Engineer Dumps Free □ Professional-Machine-Learning-Engineer Latest Exam Experience □ Professional-Machine-Learning-Engineer Study Group □ Download ➠ Professional-Machine-Learning-Engineer □ for free by simply searching on □ www.pdfvce.com □ □Professional-Machine-Learning-Engineer Exam Labs
- How to Crack Google Professional-Machine-Learning-Engineer Certification Exam Easily? □ Search for □ Professional-Machine-Learning-Engineer □ and download it for free on { www.practicevce.com } website □Professional-Machine-Learning-Engineer Test Voucher
- www.stes.tyc.edu.tw, dkpacademy.in, www.stes.tyc.edu.tw, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, www.stes.tyc.edu.tw, www.stes.tyc.edu.tw, www.stes.tyc.edu.tw, www.stes.tyc.edu.tw, www.stes.tyc.edu.tw, www.stes.tyc.edu.tw, Disposable vapes

2026 Latest PassTorrent Professional-Machine-Learning-Engineer PDF Dumps and Professional-Machine-Learning-Engineer Exam Engine Free Share: https://drive.google.com/open?id=1JkBstlUBBUr28wXvBt1F0F9J5ckWyBpo