



PROFESSIONAL-MACHINE- LEARNING-ENGINEER^{Q&As}

Professional Machine Learning Engineer

**Pass Google PROFESSIONAL-MACHINE-LEARNING-
ENGINEER Exam with 100% Guarantee**

Free Download Real Questions & Answers **PDF** and **VCE** file from:

<https://www.pass4itsure.com/professional-machine-learning-engineer.html>

100% Passing Guarantee
100% Money Back Assurance

Following Questions and Answers are all new published by Google
Official Exam Center



VCE & PDF

Pass4itSure.com

<https://www.pass4itsure.com/professional-machine-learning-engineer.html>

2024 Latest pass4itsure PROFESSIONAL-MACHINE-LEARNING-ENGINEER

PDF and VCE dumps Download

- ⚙️ **Instant Download** After Purchase
- ⚙️ **100% Money Back** Guarantee
- ⚙️ **365 Days** Free Update
- ⚙️ **800,000+** Satisfied Customers



**QUESTION 1**

You need to build classification workflows over several structured datasets currently stored in BigQuery. Because you will be performing the classification several times, you want to complete the following steps without writing code: exploratory data analysis, feature selection, model building, training, and hyperparameter tuning and serving.

What should you do?

- A. Configure AutoML Tables to perform the classification task.
- B. Run a BigQuery ML task to perform logistic regression for the classification.
- C. Use AI Platform Notebooks to run the classification model with pandas library.
- D. Use AI Platform to run the classification model job configured for hyperparameter tuning.

Correct Answer: A

QUESTION 2

You are an ML engineer at a global car manufacture. You need to build an ML model to predict car sales in different cities around the world. Which features or feature crosses should you use to train city-specific relationships between car type and number of sales?

- A. Three individual features: binned latitude, binned longitude, and one-hot encoded car type.
- B. One feature obtained as an element-wise product between latitude, longitude, and car type.
- C. One feature obtained as an element-wise product between binned latitude, binned longitude, and one-hot encoded car type.
- D. Two feature crosses as an element-wise product: the first between binned latitude and one-hot encoded car type, and the second between binned longitude and one-hot encoded car type.

Correct Answer: C

<https://developers.google.com/machine-learning/crash-course/feature-crosses/check-your-understanding>

QUESTION 3

You built a deep learning-based image classification model by using on-premises data. You want to use Vertex AI to deploy the model to production. Due to security concerns, you cannot move your data to the cloud. You are aware that the input data distribution might change over time. You need to detect model performance changes in production. What should you do?

- A. Use Vertex Explainable AI for model explainability. Configure feature-based explanations.
- B. Use Vertex Explainable AI for model explainability. Configure example-based explanations.
- C. Create a Vertex AI Model Monitoring job. Enable training-serving skew detection for your model.



D. Create a Vertex AI Model Monitoring job. Enable feature attribution skew and drift detection for your model.

Correct Answer: D

QUESTION 4

Your team is building a convolutional neural network (CNN)-based architecture from scratch. The preliminary experiments running on your on-premises CPU-only infrastructure were encouraging, but have slow convergence. You have been asked to speed up model training to reduce time-to-market. You want to experiment with virtual machines (VMs) on Google Cloud to leverage more powerful hardware. Your code does not include any manual device placement and has not been wrapped in Estimator model-level abstraction. Which environment should you train your model on?

- A. AVM on Compute Engine and 1 TPU with all dependencies installed manually.
- B. AVM on Compute Engine and 8 GPUs with all dependencies installed manually.
- C. A Deep Learning VM with an n1-standard-2 machine and 1 GPU with all libraries pre-installed.
- D. A Deep Learning VM with more powerful CPU e2-highcpu-16 machines with all libraries pre-installed.

Correct Answer: C

https://cloud.google.com/deep-learning-vm/docs/cli#creating_an_instance_with_one_or_more_gpus

https://cloud.google.com/deep-learning-vm/docs/introduction#pre-installed_packages

QUESTION 5

You recently joined an enterprise-scale company that has thousands of datasets. You know that there are accurate descriptions for each table in BigQuery, and you are searching for the proper BigQuery table to use for a model you are building on AI Platform. How should you find the data that you need?

- A. Use Data Catalog to search the BigQuery datasets by using keywords in the table description.
- B. Tag each of your model and version resources on AI Platform with the name of the BigQuery table that was used for training.
- C. Maintain a lookup table in BigQuery that maps the table descriptions to the table ID. Query the lookup table to find the correct table ID for the data that you need.
- D. Execute a query in BigQuery to retrieve all the existing table names in your project using the INFORMATION_SCHEMA metadata tables that are native to BigQuery. Use the result to find the table that you need.

Correct Answer: A

<https://cloud.google.com/data-catalog/docs/concepts/overview>

QUESTION 6

You are an ML engineer at a travel company. You have been researching customers' travel behavior for many years,



and you have deployed models that predict customers\' vacation patterns. You have observed that customers\' vacation destinations vary based on seasonality and holidays; however, these seasonal variations are similar across years. You want to quickly and easily store and compare the model versions and performance statistics across years. What should you do?

- A. Store the performance statistics in Cloud SQL. Query that database to compare the performance statistics across the model versions.
- B. Create versions of your models for each season per year in Vertex AI. Compare the performance statistics across the models in the Evaluate tab of the Vertex AI UI.
- C. Store the performance statistics of each pipeline run in Kubeflow under an experiment for each season per year. Compare the results across the experiments in the Kubeflow UI.
- D. Store the performance statistics of each version of your models using seasons and years as events in Vertex ML Metadata. Compare the results across the slices.

Correct Answer: D

<https://cloud.google.com/vertex-ai/docs/ml-metadata/analyzing#filtering>

QUESTION 7

You are an ML engineer in the contact center of a large enterprise. You need to build a sentiment analysis tool that predicts customer sentiment from recorded phone conversations. You need to identify the best approach to building a model while ensuring that the gender, age, and cultural differences of the customers who called the contact center do not impact any stage of the model development pipeline and results. What should you do?

- A. Convert the speech to text and extract sentiments based on the sentences.
- B. Convert the speech to text and build a model based on the words.
- C. Extract sentiment directly from the voice recordings.
- D. Convert the speech to text and extract sentiment using syntactical analysis.

Correct Answer: A

QUESTION 8

You are deploying a new version of a model to a production Vertex AI endpoint that is serving traffic. You plan to direct all user traffic to the new model. You need to deploy the model with minimal disruption to your application. What should you do?

- A. 1. Create a new endpoint
- 2.

Create a new model. Set it as the default version. Upload the model to Vertex AI Model Registry

- 3.

Deploy the new model to the new endpoint



4.

Update Cloud DNS to point to the new endpoint

B. 1. Create a new endpoint

2.

Create a new model. Set the parentModel parameter to the model ID of the currently deployed model and set it as the default version. Upload the model to Vertex AI Model Registry

3.

Deploy the new model to the new endpoint, and set the new model to 100% of the traffic.

C. 1. Create a new model. Set the parentModel parameter to the model ID of the currently deployed model. Upload the model to Vertex AI Model Registry.

2. Deploy the new model to the existing endpoint, and set the new model to 100% of the traffic

D. 1. Create a new model. Set it as the default version. Upload the model to Vertex AI Model Registry

2. Deploy the new model to the existing endpoint

Correct Answer: C

QUESTION 9

You work as an analyst at a large banking firm. You are developing a robust scalable ML pipeline to train several regression and classification models. Your primary focus for the pipeline is model interpretability. You want to productionize the pipeline as quickly as possible. What should you do?

A. Use Tabular Workflow for Wide and Deep through Vertex AI Pipelines to jointly train wide linear models and deep neural networks

B. Use Google Kubernetes Engine to build a custom training pipeline for XGBoost-based models

C. Use Tabular Workflow for TabNet through Vertex AI Pipelines to train attention-based models

D. Use Cloud Composer to build the training pipelines for custom deep learning-based models

Correct Answer: C

QUESTION 10

You work at a mobile gaming startup that creates online multiplayer games. Recently, your company observed an increase in players cheating in the games, leading to a loss of revenue and a poor user experience. You built a binary classification model to determine whether a player cheated after a completed game session, and then send a message to other downstream systems to ban the player that cheated. Your model has performed well during testing, and you now need to deploy the model to production. You want your serving solution to provide immediate classifications after a completed game session to avoid further loss of revenue. What should you do?



- A. Import the model into Vertex AI Model Registry. Use the Vertex Batch Prediction service to run batch inference jobs.
- B. Save the model files in a Cloud Storage bucket. Create a Cloud Function to read the model files and make online inference requests on the Cloud Function.
- C. Save the model files in a VM. Load the model files each time there is a prediction request, and run an inference job on the VM
- D. Import the model into Vertex AI Model Registry. Create a Vertex AI endpoint that hosts the model, and make online inference requests.

Correct Answer: D

QUESTION 11

You have developed a BigQuery ML model that predicts customer churn, and deployed the model to Vertex AI Endpoints. You want to automate the retraining of your model by using minimal additional code when model feature values change. You also want to minimize the number of times that your model is retrained to reduce training costs. What should you do?

- A. 1. Enable request-response logging on Vertex AI Endpoints
2.
Schedule a TensorFlow Data Validation job to monitor prediction drift
3.
Execute model retraining if there is significant distance between the distributions
- B. 1. Enable request-response logging on Vertex AI Endpoints
2.
Schedule a TensorFlow Data Validation job to monitor training/serving skew
3.
Execute model retraining if there is significant distance between the distributions
- C. 1. Create a Vertex AI Model Monitoring job configured to monitor prediction drift
2.
Configure alert monitoring to publish a message to a Pub/Sub queue when a monitoring alert is detected
3.
Use a Cloud Function to monitor the Pub/Sub queue, and trigger retraining in BigQuery
- D. 1. Create a Vertex AI Model Monitoring job configured to monitor training/serving skew
2.



Configure alert monitoring to publish a message to a Pub/Sub queue when a monitoring alert is detected

3.

Use a Cloud Function to monitor the Pub/Sub queue, and trigger retraining in BigQuery

Correct Answer: D

QUESTION 12

You recently deployed a model to a Vertex AI endpoint. Your data drifts frequently, so you have enabled request-response logging and created a Vertex AI Model Monitoring job. You have observed that your model is receiving higher traffic than expected. You need to reduce the model monitoring cost while continuing to quickly detect drift. What should you do?

- A. Replace the monitoring job with a DataFlow pipeline that uses TensorFlow Data Validation (TFDV)
- B. Replace the monitoring job with a custom SQL script to calculate statistics on the features and predictions in BigQuery
- C. Decrease the sample_rate parameter in the RandomSampleConfig of the monitoring job
- D. Increase the monitor_interval parameter in the ScheduleConfig of the monitoring job

Correct Answer: C

QUESTION 13

You are an ML engineer responsible for designing and implementing training pipelines for ML models. You need to create an end-to-end training pipeline for a TensorFlow model. The TensorFlow model will be trained on several terabytes of structured data. You need the pipeline to include data quality checks before training and model quality checks after training but prior to deployment. You want to minimize development time and the need for infrastructure maintenance. How should you build and orchestrate your training pipeline?

- A. Create the pipeline using Kubeflow Pipelines domain-specific language (DSL) and predefined Google Cloud components. Orchestrate the pipeline using Vertex AI Pipelines.
- B. Create the pipeline using TensorFlow Extended (TFX) and standard TFX components. Orchestrate the pipeline using Vertex AI Pipelines.
- C. Create the pipeline using Kubeflow Pipelines domain-specific language (DSL) and predefined Google Cloud components. Orchestrate the pipeline using Kubeflow Pipelines deployed on Google Kubernetes Engine.
- D. Create the pipeline using TensorFlow Extended (TFX) and standard TFX components. Orchestrate the pipeline using Kubeflow Pipelines deployed on Google Kubernetes Engine.

Correct Answer: B

<https://cloud.google.com/vertex-ai/docs/pipelines/build-pipeline#sdk>

**QUESTION 14**

You have recently used TensorFlow to train a classification model on tabular data. You have created a Dataflow pipeline that can transform several terabytes of data into training or prediction datasets consisting of TFRecords. You now need to productionize the model, and you want the predictions to be automatically uploaded to a BigQuery table on a weekly schedule. What should you do?

- A. Import the model into Vertex AI and deploy it to a Vertex AI endpoint. On Vertex AI Pipelines, create a pipeline that uses the DataflowPythonJobOp and the ModelBatchPredictOp components.
- B. Import the model into Vertex AI and deploy it to a Vertex AI endpoint. Create a Dataflow pipeline that reuses the data processing logic sends requests to the endpoint, and then uploads predictions to a BigQuery table.
- C. Import the model into Vertex AI. On Vertex AI Pipelines, create a pipeline that uses the DataflowPythonJobOp and the ModelBatchPredictOp components.
- D. Import the model into BigQuery. Implement the data processing logic in a SQL query. On Vertex AI Pipelines create a pipeline that uses the BigQueryQueryJobOp and the BigQueryPredictModelJobOp components.

Correct Answer: C

QUESTION 15

You are the Director of Data Science at a large company, and your Data Science team has recently begun using the Kubeflow Pipelines SDK to orchestrate their training pipelines. Your team is struggling to integrate their custom Python code into the Kubeflow Pipelines SDK. How should you instruct them to proceed in order to quickly integrate their code with the Kubeflow Pipelines SDK?

- A. Use the `func_to_container_op` function to create custom components from the Python code.
- B. Use the predefined components available in the Kubeflow Pipelines SDK to access Dataproc, and run the custom code there.
- C. Package the custom Python code into Docker containers, and use the `load_component_from_file` function to import the containers into the pipeline.
- D. Deploy the custom Python code to Cloud Functions, and use Kubeflow Pipelines to trigger the Cloud Function.

Correct Answer: A

https://kubeflow-pipelines.readthedocs.io/en/stable/source/kfp.components.html?highlight=func_to_container_op%20#kfp.components.func_to_container_op

[Latest PROFESSIONAL-MACHINE-LEARNING-ENGINEER Dumps](#)

[PROFESSIONAL-MACHINE-LEARNING-ENGINEER Practice Test](#)

[PROFESSIONAL-MACHINE-LEARNING-ENGINEER Braindumps](#)