Machine learning on Google Cloud

Course 4329 – 40 Hours

Overview

This course teaches you how to build Vertex AI AutoML models without writing a single line of code; build BigQuery ML models knowing basic SQL; create Vertex AI custom training jobs you deploy using containers (with little knowledge of Docker0; use Feature Store for data management and governance; use feature engineering for model improvement; determine the appropriate data preprocessing options for your use case; write distributed ML models that scale in TensorFlow; and leverage best practices to implement machine learning on Google Cloud. Learn all this and more!

On Completion, Delegates will be able to

- Build, train, and deploy a machine learning model without writing a single line of code using Vertex AI AutoML.
- Understand when to use AutoML and Big Query ML.
- Create Vertex AI managed datasets.
- Add features to a Feature Store.
- Describe Analytics Hub, Dataplex, and Data Catalog.
- Describe hyperparameter tuning using Vertex Vizier and how it can be used to improve model performance.
- Create a Vertex AI Workbench User-Managed Notebook, build a custom training job, and then deploy it using a Docker container.
- Describe batch and online predictions and model monitoring.
- Describe how to improve data quality.
- Perform exploratory data analysis.
- Build and train supervised learning models.
- Optimize and evaluate models using loss functions and performance metrics.
- Create repeatable and scalable train, eval, and test datasets.
- Implement ML models using TensorFlow/Keras.
- Describe how to represent and transform features.
- Understand the benefits of using feature engineering.
- Explain Vertex AI Pipelines.

Who Should Attend

- Aspiring machine learning data analysts, data scientists and data engineers
- Learners who want exposure to ML using Vertex AI AutoML, BQML, Feature Store, Workbench, Dataflow, Vizier for hyperparameter tuning, and TensorFlow/Keras

Prerequisites

To get the most out of this course, participants should have:

- Some familiarity with basic machine learning concepts
- Basic proficiency with a scripting language, preferably Python

Course Contents

Module 1: How Google Does Machine Learning

- Describe the Vertex AI Platform and how it is used to quickly build, train, and deploy AutoML machine learning models without writing a single line of code.
- Describe best practices for implementing machine learning on Google Cloud.
- Develop a data strategy around machine learning.
- Examine use cases that are then reimagined through an ML lens.
- Leverage Google Cloud Platform tools and environment to do ML.
- Learn from Google's experience to avoid common pitfalls.
- Carry out data science tasks in online collaborative notebooks.

Module 2: Launching into Machine Learning

- Describe Vertex AI AutoML and how to build, train, and deploy an ML model without writing a single line of code.
- Describe Big Query ML and its benefits.
- Describe how to improve data quality.
- Perform exploratory data analysis.
- Build and train supervised learning models.
- Optimize and evaluate models using loss functions and performance metrics.
- Mitigate common problems that arise in machine learning.
- Create repeatable and scalable training, evaluation, and test datasets.

Module 3: TensorFlow on Google Cloud

- Create TensorFlow and Keras machine learning models.
- Describe TensorFlow key components.
- Use the tf.data library to manipulate data and large datasets.
- Build a ML model using tf.keras preprocessing layers.
- Use the Keras Sequential and Functional APIs for simple and advanced model creation. Understand how model subclassing can be used for more customized models.
- Use tf.keras.preprocessing utilities for working with image data, text data, and sequence data.
- Train, deploy, and productionalize ML models at scale with Cloud AI Platform.

Module 4: Feature Engineering

- Describe Vertex AI Feature Store.
- Compare the key required aspects of a good feature.
- Combine and create new feature combinations through feature crosses.
- Perform feature engineering using BigQuery ML, Keras, and TensorFlow.
- Understand how to preprocess and explore features with Dataflow and Dataprep by Trifacta.
- Understand and apply how TensorFlow transforms features.

Module 5: Machine Learning in the Enterprise

- Understand the tools required for data management and governance.
- Describe the best approach for data preprocessing: from providing an overview of
- Dataflow and Dataprep to using SQL for preprocessing tasks.
- Explain how AutoML, BigQuery ML, and custom training differ and when to use a particular framework.
- Describe hyperparameter tuning using Vertex Vizier and how it can be used to improve model performance.
- Explain prediction and model monitoring and how Vertex AI can be used to manage
- ML models.
- Describe the benefits of Vertex AI Pipelines.