

Machine Learning on Google Cloud

5journs / 35h

What you'll learn

- Describe the technologies, products, and tools to build an ML model, an ML pipeline, and a Generative AI project.
- Understand when to use AutoML and BigQuery ML.
- Create Vertex AI-managed datasets.
- Add features to the Vertex AI Feature Store.
- Describe Analytics Hub, Dataplex, and Data Catalog.
- Describe how to improve model performance.
- Create Vertex AI Workbench user-managed notebook, build a custom training job, and deploy it by using a Docker container.
- Describe batch and online predictions and model monitoring.
- Describe how to improve data quality and explore your data.
- Build and train supervised learning models.
- Optimize and evaluate models by using loss functions and performance metrics.
- Create repeatable and scalable train, eval, and test datasets.
- Implement ML models by using TensorFlow or Keras.
- Understand the benefits of using feature engineering.
- Explain Vertex AI Model Monitoring and Vertex AI Pipelines.

Who this course is for

This class is primarily intended for the following participants:

- Aspiring machine learning data analysts, data scientists, and data engineers

- Learners who want exposure to ML and use Vertex AI, AutoML, BigQuery ML, Vertex AI Feature Store, Vertex AI Workbench, Dataflow, Vertex AI Vizier for hyperparameter tuning, and TensorFlow/Keras

Prerequisite

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

Programme

Module 1: Introduction to AI and Machine Learning on Google Cloud

Objectives

- Recognize the AI/ML framework on Google Cloud.
- Identify the major components of Google Cloud infrastructure.
- Define the data and ML products on Google Cloud and how they support the data-to-AI lifecycle.
- Build an ML model with BigQueryML to bring data to AI.
- Define different options to build an ML model on Google Cloud.
- Recognize the primary features and applicable situations of pre-trained APIs, AutoML, and custom training.
- Use the Natural Language API to analyze text.
- Define the workflow of building an ML model.
- Describe MLOps and workflow automation on Google Cloud.
- Build an ML model from end-to-end by using AutoML on Vertex AI.
- Define generative AI and large language models.
- Use generative AI capabilities in AI development.
- Recognize the AI solutions and the embedded generative AI features.

Activities

- Hands-On Labs
- Module Quizzes
- Module Readings

Module 2: Launching into Machine Learning

Objectives

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

Activities

- Hands-On Labs
- Module Quizzes
- Module Readings

Module 3: TensorFlow on Google Cloud

Objectives

- Create TensorFlow and Keras machine learning models.
- Describe the TensorFlow main components
- Use the tf.data library to manipulate data and large datasets.
- Build a ML model that uses tf.keras preprocessing layers.
- Use the Keras Sequential and Functional APIs for simple and advanced model creation.
- Train, deploy, and productionalize ML models at scale with the Vertex AI Training Service.

Activities

- Hands-On Labs
- Module Quizzes
- Module Readings

Module 4: Feature Engineering

Objectives

- Describe Vertex AI Feature Store.
- Compare the key required aspects of a good feature.

- Use `tf.keras.preprocessing` utilities for working with image data, text data, and sequence data.
- Perform feature engineering by using BigQuery ML, Keras, and TensorFlow.

Activities

- Hands-On Labs
- Module Quizzes
- Module Readings

Module 5: Machine Learning in the Enterprise

Objectives

- 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 by using Vertex AI Vizier 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.
- Describe best practices for model deployment and serving, model monitoring, Vertex AI Pipelines, and artifact organization.

Activities

- Hands-On Labs
- Module Quizzes
- Module Readings