Google Cloud Machine Learning - Beginner to Intermediate

Course Overview

A Machine Learning Engineer designs, builds, productionizes, optimizes, operates, and maintains ML systems. This learning path guides you through a curated collection of concepts and labs that provide you with real-world, hands-on experience using Google Cloud technologies essential to the ML Engineer role.

Duration: 05 days / 40 hours

Level: Professional (Foundation to Intermediate)

Prerequisites: There is no prerequisite for this learning path. Basic knowledge of Python and SQL query is helpful but not mandatory.

Course Outcome: Learner can take **Google Cloud Certified Professional Machine Learning Engineer** exam

Table of Content

Foundations of ML on Google Cloud

- Introduction to AI & ML on Google Cloud
- Generative AI Overview
- ML Development Workflow and Tools
- Vertex AI Overview and Development Options
- Low-code ML Solutions using BigQuery ML and AutoML

Focus: Understanding the GCP ML ecosystem and tools for both traditional and low-code/no-code approaches.

Data Exploration, AutoML & BigQuery ML

- Exploratory Data Analysis using GCP Notebooks
- AutoML with Vertex AI
- BigQuery ML: Build Models in SQL
- Looker Studio: Visualizing Data and ML Outputs
- ML APIs (Vision, NLP, Speech, Translation) and Model Garden

Focus: Fast ML prototyping with minimal code and effective use of APIs.

Building and Training Models with TensorFlow/Keras

- TensorFlow Ecosystem and Keras API
- Designing Input Pipelines
- Building & Training Neural Networks

- Training at Scale on Vertex AI
- Responsible AI (fairness, privacy, interpretability)

Focus: Custom model development and scalable training workflows.

Feature Engineering & Transformation

- Vertex AI Feature Store Overview
- Feature Engineering & Preprocessing
- TensorFlow Transform Use Cases
- Data to Feature Pipeline Demo
- Looker Studio: Feature Visualization
- Privacy-compliant design (DLP API, anonymization)

Focus: Efficient and responsible feature engineering practices.

Deployment, Monitoring & Communication

- Model Deployment with Vertex AI Endpoints
- Setting Up Model Monitoring
- ML Pipelines Overview
- Looker Studio: Communicating Results
- Best Practices for ML Ops
- Continuous Evaluation & Training-Serving Skew Detection

Focus: Deploying production-ready models and monitoring accuracy post-deployment.