

Machine Learning Engineer

Course Outline

Module 1: Google Cloud Big Data and Machine Learning Fundamentals

Topics:

- Big Data and Machine Learning on Google Cloud
- Data Engineering for Streaming Data
- Big Data with BigQuery
- Machine Learning Options on Google Cloud
- The Machine Learning Workflow with Vertex AI

Hands-On:

- AI Platform: Qwik Start
- Dataprep: Qwik Start
- Dataflow: Qwik Start - Templates
- Dataflow: Qwik Start - Python
- Dataproc: Qwik Start - Console
- Dataproc: Qwik Start - Command Line

Module 2: How Google Does Machine Learning

Topics:

- What It Means to be AI-First
- How Google Does ML
- Machine Learning Development with Vertex AI
- Machine Learning Development with Vertex Notebooks
- Best Practices for Implementing Machine Learning on Vertex AI
- Responsible AI Development

Hands-On:

- Vertex AI: Qwik Start
- Using an Image Dataset to Train an AutoML Model

Module 3: Launching into Machine Learning

Topics:

- Introduction
- Get to Know Your Data: Improve Data through Exploratory Data Analysis
- Machine Learning in Practice
- Training AutoML Models Using Vertex AI
- BigQuery Machine Learning: Develop ML Models Where Your Data Lives
- Optimization
- Generalization and Sampling

Hands-On:

- Exploratory Data Analysis Using Python and BigQuery
- Using BigQuery ML to Predict Penguin Weight

Module 4: TensorFlow on Google Cloud

Topics:

- Introduction to the TensorFlow ecosystem

- Design and Build an Input Data Pipeline
- Building Neural Networks with the TensorFlow and Keras API
- Training at Scale with Vertex AI

Hands-On:

- Classifying Structured Data using Keras Preprocessing Layers
- Build a DNN using the Keras Functional API

Module 5: Feature Engineering

Topics:

- Introduction to Vertex AI Feature Store
- Raw Data to Features
- Feature Engineering
- Preprocessing and Feature Creation
- Feature Crosses - TensorFlow Playground
- Introduction to TensorFlow Transform

Hands-On:

- Using Feature Store
- Performing Basic Feature Engineering in BQML
- Performing Basic Feature Engineering in Keras

Module 6: Machine Learning in the Enterprise

Topics:

- Introduction
- Understanding the ML Enterprise Workflow
- Data in the Enterprise
- Science of Machine Learning and Custom Training
- Vertex Vizier Hyperparameter Tuning
- Prediction and Model Monitoring Using Vertex AI
- Vertex AI Pipelines
- Best Practices for ML Development

Hands-On:

- Vertex Pipelines: Qwik Start
- Cloud Natural Language API: Qwik Start
- Google Cloud Speech API: Qwik Start
- Video Intelligence: Qwik Start

Module 7: End-to-End Machine Learning with TensorFlow on Google Cloud

Topics:

- Machine Learning (ML) on Google Cloud Platform (GCP)
- Explore the Data
- Create the dataset
- Build the Model
- Operationalize the model

Hands-On:

- Identify Damaged Car Parts with Vertex AutoML Vision
- Deploy a BigQuery ML Customer Churn Classifier to Vertex AI for Online Predictions

Module 8: Production Machine Learning Systems

Topics:

- Introduction to Advanced Machine Learning on Google Cloud
- Architecting Production ML Systems
- Designing Adaptable ML Systems
- Designing High-Performance ML Systems
- Building Hybrid ML Systems

Hands-On:

- Structured data prediction using Vertex AI Platform
- Serving ML Predictions in Batch and Real Time
- Distributed Training with Keras
- Using Kubeflow Pipelines with AI Platform

Module 9: Computer Vision Fundamentals with Google Cloud

Topics:

- Introduction to Computer Vision and Pre-built ML Models for Image Classification
- Vertex AI and AutoML Vision on Vertex AI
- Custom Training with Linear, Neural Network and Deep Neural Network models
- Convolutional Neural Networks
- Dealing with Image Data

Hands-On:

- Using the What-If Tool with Image Recognition Models
- Identifying Bias in Mortgage Data using Cloud AI Platform and the What-if Tool
- Compare Cloud AI Platform Models using the What-If Tool to Identify potential bias

Module 10: Sequence Models for Time Series and Natural Language Processing on Google Cloud

Topics:

- Working with Sequences
- Recurrent Neural Networks
- Dealing with Longer Sequences
- Text Classification
- Reusable Embeddings
- Encoder-Decoder Models

Hands-On:

- Time Series Prediction with a DNN Model
- Time Series Prediction with a Two-Layer RNN Model
- Text Classification using TensorFlow/Keras on AI Platform
- Text generation using tensor2tensor on Cloud AI Platform

Module 11: Recommendation Systems on Google Cloud

Topics:

- Recommendation Systems Overview
- Content-Based Recommendation Systems
- Collaborative Filtering Recommendations Systems
- Neural Networks for Recommendation Systems

- Reinforcement Learning

Hands-On:

- Using Neural Networks for Content-Based Filtering
- Collaborative Filtering on Google Analytics data
- ML on GCP: Hybrid Recommendations with the MovieLens Dataset
- Applying Contextual Bandits for Recommendations with Tensorflow and TFAgents

Module 12: MLOps (Machine Learning Operations) Fundamentals

Topics:

- Why and When do we Need MLOps
- Understanding the Main Kubernetes Components (Optional)
- Introduction to AI Platform Pipelines
- Training, Tuning and Serving on AI Platform
- Kubeflow Pipelines on AI Platform
- CI/CD for Kubeflow Pipelines on AI Platform

Hands-On:

- Working with Cloud Build
- Creating Google Kubernetes Engine Deployments
- Using custom containers with AI Platform Training
- Continuous Training Pipeline with Kubeflow Pipeline and Cloud AI Platform
- CI/CD for a Kubeflow pipeline

Module 13: ML Pipelines on Google Cloud

Topics:

- Introduction to TFX Pipelines
- Pipeline orchestration with TFX
- Custom components and CI/CD for TFX pipelines
- ML Metadata with TFX
- Continuous Training with multiple SDKs, KubeFlow & AI Platform Pipelines
- Continuous Training with Cloud Composer
- ML Pipelines with MLflow

Hands-On:

- TFX Standard Components Walkthrough
- TFX on Cloud AI Platform Pipelines
- CI/CD for a TFX pipeline
- Continuous Training Pipelines with Cloud Composer