

Google Cloud Machine Learning Engineer

Course Duration: 80 Hours

Course code: GCMLE

1. Course Overview

This course is designed to equip learners with the skills required to become a Google Cloud Machine Learning Engineer. It focuses on designing, building, deploying, and managing machine learning models using Google Cloud tools such as Vertex AI, BigQuery ML, and TensorFlow. The course emphasizes end-to-end ML lifecycle management, MLOps practices, and scalable AI solutions for real-world applications.

2. What you'll learn?

By the end of the course, you will be able to:

- Understand end-to-end machine learning lifecycle on Google Cloud
- Build, train, and optimize ML models using Vertex AI
- Work with structured and unstructured datasets
- Implement feature engineering and model evaluation techniques
- Deploy ML models for real-time and batch predictions
- Implement MLOps practices for automation and scalability
- Monitor and improve model performance
- Apply responsible AI and ethical practices

3. Target Audience

- Machine Learning Engineers
- Data Scientists
- AI/ML Developers
- Cloud Engineers
- Data Engineers transitioning to ML roles

4. Pre-Requisites

Before taking this course, you should have:

- Strong understanding of Python programming
- Basic knowledge of machine learning concepts
- Familiarity with Google Cloud Platform (GCP)
- Understanding of statistics and data analysis

5. Course content

Module 1: Course Introduction

- Course overview and objectives
- Role of a Machine Learning Engineer
- Introduction to ML on Google Cloud

Module 2: Machine Learning Fundamentals

- Supervised, unsupervised, and reinforcement learning
- Model types: regression, classification, clustering
- Bias-variance tradeoff
- Model evaluation metrics

Module 3: Google Cloud ML Ecosystem

- Overview of Vertex AI
- BigQuery ML and AI APIs
- Cloud Storage for datasets
- Integration of GCP services

Module 4: Data Preparation and Feature Engineering

- Data collection and preprocessing
- Data cleaning and transformation
- Feature engineering techniques
- Handling missing and imbalanced data

Module 5: Model Development with Vertex AI

- Vertex AI Workbench
- Training custom models
- Using pre-built containers
- Hyperparameter tuning

Module 6: Using AutoML

- Overview of AutoML
- Training models with minimal coding
- Image, text, and tabular data models
- Evaluating AutoML performance

Module 7: BigQuery ML

- Creating ML models using SQL
- Training and evaluating models
- Predictive analytics use cases
- Integration with datasets

Module 8: Deep Learning Fundamentals

- Neural networks basics
- TensorFlow introduction
- Training deep learning models
- Model optimization techniques

Module 9: Model Deployment

- Deploying models on Vertex AI endpoints
- Real-time vs batch predictions
- API integration
- Versioning models

Module 10: MLOps and Automation

- Introduction to MLOps
- CI/CD pipelines for ML
- Model versioning and tracking
- Automation using pipelines

Module 11: Monitoring and Optimization

- Monitoring model performance
- Detecting data and model drift
- Logging and debugging
- Continuous improvement

Module 12: Responsible AI

- Bias detection and mitigation
- Explainable AI (XAI)
- Ethical considerations
- Governance practices

Module 13: Security and Compliance

- Data security in ML workflows
- IAM roles and permissions
- Secure model deployment
- Compliance standards

Module 14: Advanced ML Topics

- Natural Language Processing (NLP)
- Computer Vision
- Recommendation systems
- Generative AI overview

Module 15: Real-World Use Cases and Capstone Project

- Industry use cases (Retail, Finance, Healthcare)
- End-to-end ML pipeline implementation
- Best practices and architecture design
- Final project and evaluation

