

# Machine Learning with TensorFlow on Google Cloud Platform

ID GO-MLTF Price on request Duration 5 days

## Course Overview

What is machine learning, and what kinds of problems can it solve? What are the five phases of converting a candidate use case to be driven by machine learning, and why is it important that the phases not be skipped? Why are neural networks so popular now? How can you set up a supervised learning problem and find a good, generalizable solution using gradient descent and a thoughtful way of creating datasets?

## Who should attend

- Data Engineers and programmers interested in learning how to apply machine learning in practice.
- Anyone interested in learning how to build and operationalize TensorFlow models.

## This course is part of the following Certifications

Google Cloud Certified Professional Machine Learning Engineer (PMLE)

## Prerequisites

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

- Experience coding in Python
- Knowledge of basic statistics
- Knowledge of SQL and cloud computing (helpful)

## Course Objectives

Learn how to write distributed machine learning models that scale in Tensorflow, scale out the training of those models, and offer high-performance predictions. Convert raw data to features in a way that allows ML to learn important characteristics from the data and bring human insight to bear on the problem. Finally, learn how to incorporate the right mix of parameters that yields accurate, generalized models and knowledge of the theory to solve specific types of ML problems. You will experiment with end-to-end ML, starting from building an ML-focused strategy and progressing into model training, optimization, and productionalization with hands-on labs using

## Google Cloud Platform

This course teaches participants the following skills:

- Frame a business use case as a machine learning problem
- Create machine learning datasets that are capable of achieving generalization
- Implement machine learning models using TensorFlow
- Understand the impact of gradient descent parameters on accuracy, training speed, sparsity, and generalization
- Build and operationalize distributed TensorFlow models
- Represent and transform features

## Course Content

- How Google Does Machine Learning
- Launching into Machine Learning
- Intro to TensorFlow
- Feature Engineering
- The Art and Science of ML

## Detailed Course Outline

### 1: How Google Does Machine Learning

What is machine learning, and what kinds of problems can it solve? Google thinks about machine learning slightly differently: it's about logic, rather than just data. We talk about why such a framing is useful when thinking about building a pipeline of machine learning models. Then we discuss the five phases of converting a candidate use case to be driven by machine learning, and consider why it is important not to skip the phases. We end with a recognition of the biases that machine learning can amplify and how to recognize them.

- Develop a data strategy around machine learning
- Examine use cases that are then reimaged through an ML lens
- Recognize biases that ML can amplify
- 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
- Invoke pre-trained ML models from Cloud Datalab

## 2: Launching into Machine Learning

Starting from a history of machine learning, we discuss why neural networks today perform so well in a variety of problems. We then discuss how to set up a supervised learning problem and find a good solution using gradient descent. This involves creating datasets that permit generalization; we talk about methods of doing so in a repeatable way that supports experimentation.

- Identify why deep learning is currently popular
- 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

## 3: Intro to TensorFlow

We introduce low-level TensorFlow and work our way through the necessary concepts and APIs so as to be able to write distributed machine learning models. Given a TensorFlow model, we explain how to scale out the training of that model and offer high-performance predictions using Cloud Machine Learning Engine.

- Create machine learning models in TensorFlow
- Use the TensorFlow libraries to solve numerical problems
- Troubleshoot and debug common TensorFlow code pitfalls
- Use `tf_estimator` to create, train, and evaluate an ML model
- Train, deploy, and productionalize ML models at scale with Cloud ML Engine

## 4: Feature Engineering

A key component of building effective machine learning models is to convert raw data to features in a way that allows ML to learn important characteristics from the data. We discuss how to represent features and code this up in TensorFlow. Human insight can be brought to bear in machine learning problems through the use of custom feature transformations. In this module, we talk about common types of transformations and how to implement them at scale.

- Turn raw data into feature vectors
- Preprocess and create new feature pipelines with Cloud Dataflow
- Create and implement feature crosses and assess their impact
- Write TensorFlow Transform code for feature engineering

## 5: The Art and Science of ML

Machine Learning is both an art that involves knowledge of the right mix of parameters that yields accurate, generalized models and a science that involves knowledge of the theory to solve specific types of ML problems. We discuss regularization, dealing with sparsity, multi-class neural networks, reusable embeddings, and many other essential concepts and principles.

- Optimize model performance with hyperparameter tuning
- Experiment with neural networks and fine-tune performance
- Enhance ML model features with embedding layers
- Create reusable custom model code with the Custom Estimator

**Training Centres worldwide**

