

Machine Learning on Google Cloud

Course Code: GCP-MLGC

Duration: 5 days

Instructor-led Training (ILT) | Virtual Instructor-led Training (VILT)

OVERVIEW

What is machine learning, and what kinds of problems can it solve? Why are neural networks so popular right now? How can you improve data quality and perform exploratory data analysis? How can you set up a supervised learning problem and find a good, generalizable solution using gradient descent?

In this course, you'll learn how to write distributed machine learning models that scale in Tensorflow 2.x, perform feature engineering in BQML and Keras, evaluate loss curves and perform hyperparameter tuning, and train models at scale with Cloud AI Platform.

SKILLS COVERED

- Frame a business use case as a machine learning problem.
- Describe how to improve data quality.
- Perform exploratory data analysis.
- Build and train supervised learning models.
- Optimize and evaluate models using loss functions and performance metrics.
- Create repeatable and scalable training, evaluation, and test datasets.
- Implement machine learning models using Keras and TensorFlow 2.x.
- Understand the impact of gradient descent parameters on accuracy, training speed, sparsity, and generalization.
- Represent and transform features.
- Train models at scale with AI Platform.

WHO SHOULD ATTEND

- Aspiring machine learning data scientists and engineers.
- Machine learning scientists, data scientists, and data analysts who want exposure to machine learning in the cloud using TensorFlow 2.x and Keras.
- Data engineers.

PREREQUISITES

- Some familiarity with basic machine learning concepts.
- Basic proficiency with a scripting language - Python preferred.

MODULES

Module 1: How Google Does Machine Learning

- 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 AI Platform.

Module 2: Launching into Machine Learning

- Describe how to improve data quality.
- Perform exploratory data analysis.
- Build and train supervised learning models.
- 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.

Module 3: Introduction to TensorFlow 2.x

- Create TensorFlow 2.x and Keras machine learning models.
- Describe Tensorflow 2.x key components.
- Use the tf.data library to manipulate data and large datasets.
- Use the Keras Sequential and Functional APIs for simple and advanced model creation.
- Train, deploy, and productionalize ML models at scale with Cloud AI Platform.

Module 4: Feature Engineering

- Compare the key required aspects of a good feature.
- Combine and create new feature combinations through feature crosses.
- Perform feature engineering using BQML, Keras, and TensorFlow 2.x.
- Understand how to preprocess and explore features with Cloud Dataflow and Cloud Dataprep.
- Understand and apply how TensorFlow transforms features.

Module 5: The Art and Science of ML

- Optimize model performance with hyperparameter tuning.
- Experiment with neural networks and fine-tune performance.
- Enhance ML model features with embedding layers.