

Sela.

GCPML

Machine Learning

college@sela.co.il

03-6176666





Machine Learning

GCPML - Version: 1

 5 days Course

Description:

This course teaches you how to build Vertex AI AutoML models without writing a single line of code; build BigQuery ML models knowing basic SQL; create Vertex AI custom training jobs you deploy using containers (with little knowledge of Docker); use Feature Store for data management and governance; use feature engineering for model improvement; determine the appropriate data preprocessing options for your use case; write distributed ML models that scale in TensorFlow; and leverage best practices to implement machine learning on Google Cloud. Learn all this and more!

Intended audience:

This class is primarily intended for the following participants:

- Aspiring machine learning data analysts, data scientists and data engineers
- Learners who want exposure to ML using Vertex AI AutoML, BQML, Feature Store, Workbench, Dataflow, Vizier for hyperparameter tuning, and TensorFlow/Keras

Prerequisites:

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

- Some familiarity with basic machine learning concepts:
- Basic proficiency with a scripting language, preferably Python

Objectives:

Topics:



How Google does Machine Learning

- What are best practices for implementing machine learning on Google Cloud? What is Vertex AI and how can you use the platform to quickly build, train, and deploy AutoML machine learning models without writing a single line of code? What is machine learning, and what kinds of problems can it solve?
- Google thinks about machine learning slightly differently: it's about providing a unified platform for managed datasets, a feature store, a way to build, train, and deploy machine learning models without writing a single line of code, providing the ability to label data, create Workbench notebooks using frameworks such as TensorFlow, SciKit Learn, Pytorch, R, and others. Our Vertex AI Platform also includes the ability to train custom models, build component pipelines, and perform both online and batch predictions. We also discuss the five phases of converting a candidate use case to be driven by machine learning, and consider why it is important to not skip the phases. We end with a recognition of the biases that machine learning can amplify and how to recognize them.

Launching into Machine Learning

- The course begins with a discussion about data: how to improve data quality and perform exploratory data analysis. We describe Vertex AI AutoML and how to build, train, and deploy an ML model without writing a single line of code. You will understand the benefits of Big Query ML. We then discuss how to optimize a machine learning (ML) model and how generalization and sampling can help assess the quality of ML models for custom training.

TensorFlow on Google Cloud

- This course covers designing and building a TensorFlow input data pipeline, building ML models with TensorFlow and Keras, improving the accuracy of ML models, writing ML models for scaled use, and writing specialized ML models.



Feature Engineering

- Want to know about Vertex AI Feature Store? Want to know how you can improve the accuracy of your ML models? What about how to find which data columns make the most useful features? Welcome to Feature Engineering, where we discuss good versus bad features and how you can preprocess and transform them for optimal use in your models. This course includes content and labs on feature engineering using BigQuery ML, Keras, and TensorFlow.

Machine Learning in the Enterprise

- This course encompasses a real-world practical approach to the ML Workflow: a case study approach that presents an ML team faced with several ML business requirements and use cases. This team must understand the tools required for data management and governance and consider the best approach for data preprocessing: from providing an overview of Dataflow and Dataprep to using BigQuery for preprocessing tasks.
- The team is presented with three options to build machine learning models for two specific use cases. This course explains why the team would use AutoML, BigQuery ML, or custom training to achieve their objectives. A deeper dive into custom training is presented in this course. We describe custom training requirements from training code structure, storage, and loading large datasets to exporting a trained model.
- You will build a custom training machine learning model, which allows you to build a container image with little knowledge of Docker.
- The case study team examines hyperparameter tuning using Vertex Vizier and how it can be used to improve model performance. To understand more about model improvement, we dive into a bit of theory: we discuss regularization, dealing with sparsity, and many other essential concepts and principles. We end with an overview of prediction and model monitoring and how Vertex AI can be used to manage ML models.