

Advanced Machine Learning with TensorFlow on Google Cloud Platform

Course#: ML-TF
Duration: 5 Days
Price: 0.00

Course Description

This course will give you hands-on experience optimizing, deploying, and scaling a variety of production ML models. You'll learn how to build scalable, accurate, and production-ready models for structured data, image data, time-series, and natural language text, along with recommendation systems.

Objectives

Implement the various flavors of production ML systems: static, dynamic, and continuous training; static and dynamic inference; and batch and online processing.

Solve an ML problem by building an end-to-end pipeline, going from data exploration, preprocessing, feature engineering, model building, hyperparameter tuning, deployment, and serving.

Develop a range of image classification models from simple linear models to high-performing convolutional neural networks (CNNs) with batch normalization, augmentation, and transfer learning.

Forecast time-series values using CNNs, recurrent neural networks (RNNs), and LSTMs.

Apply ML to natural language text using CNNs, RNNs, LSTMs, reusable word embeddings, and encoder-decoder generative models.

Implement content-based, collaborative, hybrid, and neural recommendation models in TensorFlow.

Audience

Data Engineers and programmers interested in learning how to apply machine learning in practice. Anyone interested in learning how to leverage machine learning in their enterprise.

Prerequisites

There are no prerequisites for this course.

Content

The course includes presentations, demonstrations, and hands-on labs.

Module 1: Machine Learning on Google Cloud Platform

Effective ML.

Fully managed ML.

Module 2: Explore the Data

Exploring the dataset.

BigQuery.

BigQuery and AI Platform Notebooks.

Module 3: Creating the dataset

Creating a dataset.

Module 4: Build the Model

Build the model.

Module 5: Operationalize the model

Operationalizing the model.

Cloud AI Platform.

Train and deploy with Cloud AI Platform.

BigQuery ML.

Deploying and Predicting with Cloud AI Platform.

Module 6: Architecting Production ML Systems

The Components of an ML System.

The Components of an ML System: Data Analysis and Validation.

The Components of an ML System: Data Transformation + Trainer.

The Components of an ML System: Tuner + Model Evaluation and Validation.

The Components of an ML System: Serving.

The Components of an ML System: Orchestration + Workflow.

The Components of an ML System: Integrated Frontend + Storage.

Training Design Decisions.

Serving Design Decisions.

Designing from Scratch.

Module 7: Ingesting data for Cloud-based analytics and ML

Data On-Premise.

Large Datasets.

Data on Other Clouds.

Existing Databases.

Module 8: Designing Adaptable ML systems

Adapting to Data.

Changing Distributions.

Right and Wrong Decisions.

System Failure.

Mitigating Training-Serving Skew through Design.

Debugging a Production Model.

Module 9: Designing High-performance ML systems

Training.

Predictions.

Why distributed training?

Distributed training architectures.

Faster input pipelines.

Native TensorFlow Operations.

TensorFlow Records.

Parallel pipelines.

Data parallelism with All Reduce.

Parameter Server Approach.

Inference.

Module 10: Hybrid ML systems

Machine Learning on Hybrid Cloud.

KubeFlow.

Embedded Models.

TensorFlow Lite.

Optimizing for Mobile.

Module 11: Welcome to Image Understanding with TensorFlow on GCP

Images as Visual Data.

Structured vs Unstructured Data.

Module 12: Linear and DNN Models

Linear Models.

DNN Models Review.

Review: What is Dropout?

Module 13: Convolutional Neural Networks (CNNs)

Understanding Convolutions.

CNN Model Parameters.

Working with Pooling Layers.

Implementing CNNs with TensorFlow.

Module 14: Dealing with Data Scarcity

The Data Scarcity Problem.

Data Augmentation.

Transfer Learning.

No Data, No Problem.

Module 15: Going Deeper Faster

Batch Normalization.

Residual Networks.

Accelerators (CPU vs GPU, TPU).

TPU Estimator.

Neural Architecture Search.

Module 16: Pre-built ML Models for Image Classification

Pre-built ML Models.

Cloud Vision API.

AutoML Vision.

AutoML Architecture.

Module 17: Working with Sequences

Sequence data and models.
From sequences to inputs,
Modeling sequences with linear models.
Modeling sequences with DNNs.
Modeling sequences with CNNs.
The variable-length problem.

Module 18: Recurrent Neural Networks

Introducing Recurrent Neural Networks.
How RNNs represent the past.
The limits of what RNNs can represent.
The vanishing gradient problem.

Module 19: Dealing with Longer Sequences

LSTMs and GRUs.
RNNs in TensorFlow.
Deep RNNs.
Improving our Loss Function.
Working with Real Data.

Module 20: Text Classification

Working with Text.
Text Classification.
Selecting a Model.

Python vs Native TensorFlow.

Module 21: Reusable Embeddings

Historical methods of making word embeddings.

Modern methods of making word embeddings.

Introducing TensorFlow Hub.

Using TensorFlow Hub within an estimator.

Module 22: Recurrent Neural Networks Encoder-Decoder Models

Introducing Encoder-Decoder Networks.

Attention Networks.

Training Encoder-Decoder Models with TensorFlow.

Introducing Tensor2Tensor.

AutoML Translation.

Dialogflow.

Module 23: Recommendation Systems Overview

Types of Recommendation Systems.

Content-Based or Collaborative.

Recommendation System Pitfalls.

Module 24: Content-Based Recommendation Systems

Content-Based Recommendation Systems.

Similarity Measures.

Building a User Vector.

Making Recommendations Using a User Vector.

Making Recommendations for Many Users.

Using Neural Networks for Content-Based Recommendation Systems.

Module 25: Collaborative Filtering Recommendation Systems

Types of User Feedback Data.

Embedding Users and Items.

Factorization Approaches.

The ALS Algorithm.

Preparing Input Data for ALS.

Creating Sparse Tensors For Efficient WALS Input.

Instantiating a WALS Estimator: From Input to Estimator.

Instantiating a WALS Estimator: Decoding TFRecords.

Instantiating a WALS Estimator: Recovering Keys.

Instantiating a WALS Estimator: Training and Prediction.

Issues with Collaborative Filtering.

Cold Starts.

Module 26: Neural Networks for Recommendation Systems

Hybrid Recommendation System.

Context-Aware Recommendation Systems.

Context-Aware Algorithms.

Contextual Postfiltering.

Modeling Using Context-Aware Algorithms.

Module 27: Building an End-to-End Recommendation System

Architecture Overview.

Cloud Composer Overview.

Cloud Composer: DAGs.

Cloud Composer: Operators for ML9.

Cloud Composer: Scheduling.

Cloud Composer: Triggering Workflows with Cloud Functions.

Cloud Composer: Monitoring and Logging.