

Updated on 05/05/2026

Sign up

Deep Learning Training with TensorFlow

3 days (21 hours)

Overview

TensorFlow is one of the most powerful and mature frameworks for Deep Learning. Developed by Google, it offers a comprehensive ecosystem for moving from research to large-scale production, thanks to its flexibility and multi-platform deployment capabilities.

Our Deep Learning with TensorFlow training will enable you to master the Keras API to design high-performance models, while exploring low-level features for full control over your architectures.

You will learn to structure your data with the `tf.data` API, optimize your training, and use TensorBoard for monitoring and debugging. The course also covers crucial aspects of industrialization, from server deployment with TensorFlow Serving to mobile optimization with TFLite.

By the end of the course, you will be able to build complex neural architectures, train them efficiently on GPUs, and ensure their monitoring in a production environment.

Like all our courses, this one will introduce you to **the latest stable version** of the technology and its new features.

Objectives

- Create neural networks
- Train and evaluate models
- Use advanced TensorFlow APIs
- Deploy and monitor DL models

Target Audience

- Data scientists
- AI engineers

Prerequisites

- Knowledge of Python, ML, and mathematics

Software requirements

- At least 8 GB of RAM, 16 GB recommended
- Python environment (Anaconda or venv) with TensorFlow installed
- Access to a GPU (via Google Colab or a local CUDA-compatible machine)
- A code editor such as VS Code or Jupyter Notebook

Deep Learning with TensorFlow Training

[Day 1 - Morning]

TensorFlow Architecture and Fundamentals

- Understanding the TensorFlow 2.x ecosystem and Keras integration
- Tensor manipulation and distributed mathematical operations
- Memory Management and Hardware Acceleration (GPU/TPU)
- The tf.data API: Creating High-Performance Data Pipelines
- Reading complex formats (TFRecord) and on-the-fly transformation
- Hands-on workshop: Building an optimized data pipeline for a massive dataset.

[Day 1 - Afternoon]

Dense neural networks and the Keras API

- Sequential API vs. Functional API: when and why to choose?
- Dense layers, activation functions, and weight initialization
- Model training: Loss functions, optimizers (Adam, SGD), and metrics
- Callback management: EarlyStopping, ModelCheckpoint, and TensorBoard
- Hands-on workshop: Creating and training a multi-class classifier on tabular data.

Advanced training and customization

- Writing custom training loops with tf.GradientTape
- Creating custom layers and models via subclassing

- Gradient management and learning stabilization
- Hands-on workshop: Implementing a custom cost function for a specific problem.

[Day 2 - Morning]

Computer Vision with CNNs

- Convolutional Neural Network (CNN) Architecture
- Convolutions, Pooling, Padding, and Stride
- Real-time data augmentation techniques
- Using Pre-trained Models and Transfer Learning
- Strategic fine-tuning of the final layers
- Hands-on Workshop: Complex Object Recognition with a Pre-trained ResNet Model.

[Day 2 - Afternoon]

Sequences and Time Series

- Recurrent Neural Networks (RNNs) and LSTM/GRU cells
- Processing variable-length sequences
- Attention mechanisms and introduction to Transformers
- Time Series Forecasting
- Hands-on workshop: Trend prediction on time series data.

Optimization and Regularization

- Combating Overfitting: Dropout, Batch Normalization, L1/L2
- Hyperparameter tuning with Keras Tuner
- Advanced visualization with TensorBoard diagnostic tools
- Hands-on workshop: Fine-tuning a model to maximize generalization.

[Day 3 - Morning]

Advanced Deep Learning and NLP

- Embeddings and text vectorization
- Using TensorFlow Hub to Reuse Components
- Introduction to autoencoder architectures
- Handling multiple inputs and multiple outputs
- Hands-on Workshop: Sentiment Analysis on Text Reviews.

[Day 3 - Afternoon] Deployment

and Industrialization

- Exporting models to SavedModel format
- Introduction to TensorFlow Serving for Inference via API
- Optimization for mobile and embedded devices with TensorFlow Lite
- Model quantization to reduce size and increase speed
- Hands-on Workshop: Exporting and Deploying a Model to an Inference Container

Monitoring and Maintenance

- Monitoring Data Drift in Production
- Setting up regression tests for models
- Automating pipelines with TensorFlow Extended (TFX) - Overview
- Hands-on workshop: Final project—Design, training, and export of a complete pipeline.

Target Audience

This training is intended for both individuals and companies, large or small, wishing to train their teams in new advanced IT technologies or to acquire specific business knowledge or modern methods.

Entry-level assessment

The pre-training assessment complies with Qualiopi quality standards. Upon final registration, the learner receives a self-assessment questionnaire that allows us to evaluate their estimated proficiency in various types of technologies, as well as their expectations and personal goals regarding the upcoming training, within the limits imposed by the selected format. This questionnaire also allows us to anticipate certain connection or internal security issues within the company (intra-company or virtual classroom) that could pose challenges for monitoring and ensuring the smooth running of the training session.

Teaching Methods

Practical Course: 60% Practical, 40% Theory. Training materials distributed in digital format to all participants.

Organization

The course alternates between theoretical input from the trainer, supported by examples and reflection sessions, and group work.

Assessment

At the end of the session, a multiple-choice questionnaire is used to verify the correct acquisition

of the skills.

Certification

A certificate will be issued to each trainee who has completed the entire training program.