Advanced Machine Learning with TensorFlow on Google Cloud Platform

Category: AI, Data & Analytics | **Vendor:** Google Cloud

Duration: 40.00 hours (5 days) **32.5 CPD Hours** **Rating:** ⭐ 4.6 (5,878 reviews)

Course Information

Delivery Format: Instructor Led - Online

Course Overview

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.

About This Course

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.

Who Should Attend

- 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 & Entry Requirements

General Prerequisites:

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

- Knowledge of machine learning and TensorFlow to the level covered in Machine Learning on Google Cloud coursework
- Experience coding in Python
- Knowledge of basic statistics
- Knowledge of SQL and cloud computing (helpful)

Learning Outcomes

Upon successful completion of this course, participants will be able to:

This course teaches participants the following skills:

- 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

Detailed Course Outline

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-Premises
- 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 NetworksEncoder-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 WAL 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

Skills You Will Learn:

This course teaches participants the following skills:

- 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

Frequently Asked Questions

Q: What delivery options are available for Advanced Machine Learning with TensorFlow on Google Cloud Platform?

We offer multiple delivery formats:

- Live Instructor-Led Classroom Online (Virtual/Live Online)
- Traditional Instructor-Led Classroom Training (ILT)
- On-site delivery at your offices anywhere in United Kingdom
- Private dedicated courses customized for your team

Q: How many CPD hours does this course provide?

The 5-day Advanced Machine Learning with TensorFlow on Google Cloud Platform course provides up to 32.5 CPD hours of structured learning. CPD certificates can be provided upon request.

Q: What is the duration of the Advanced Machine Learning with TensorFlow on Google Cloud Platform training?

The training takes place over 5 day(s), with each day lasting approximately 40.00 hours including breaks for lunch and refreshments.

Q: Do you provide corporate training for Advanced Machine Learning with TensorFlow on Google Cloud Platform?

Yes, we provide corporate training, dedicated training, and closed classes for Advanced Machine Learning with TensorFlow on Google Cloud Platform. Training can take place anywhere in United Kingdom including London, Manchester, Birmingham, Edinburgh, or live online allowing teams from across United Kingdom or internationally to attend.

Q: What related terms do people search for?

Popular related searches include: Advanced Machine Learning with TensorFlow on Google Cloud Platform, AI & Data Science, MLTF

Q: Why choose Nexus Human for Advanced Machine Learning with TensorFlow on Google Cloud Platform?

Nexus Human is recognized as one of the leading training providers. Our trainers have won multiple awards including:

- Small Firms Best Trainer Award
- National Training Partner of the Year (Ireland) - Multiple Years
- Global Top 30 Instructor Awards (2012, 2019, 2021)
- Tech Excellence Award Nominations
- Learning Performance Institute (LPI) External Training Provider Sponsor 2024

Q: Are there any discount codes available?

Yes! Use discount code **PENPAL5** when booking your Advanced Machine Learning with TensorFlow on Google Cloud Platform training. Please note that only one discount code can be used per booking and cannot be combined with other special offers.

Nexus Human

Professional Training & Development

Email: info@nexushuman.com

Website: www.nexushuman.com

Phone: +353 1 889 8700 (Ireland) | +44 20 8142 8961 (UK)

© 2026 Nexus Human. All rights reserved. This brochure was generated on 14/11/2026.