

PyTorch Essentials: An Applications-First Approach (LFD273)

The course begins with an overview of PyTorch, including model classes, datasets, data loaders and the training loop. Next, it covers the role and power of transfer learning, along with how to use it with pretrained models. Practical lab exercises cover multiple topics including: image classification, object detection, sentiment analysis, text classification, and text generation/completion. Learners also will use their data to fine-tune existing models and leverage third-party APIs.

Duration: 5 Days

Prerequisites for this course

In order to complete this course, learners should be able to:

- Python (notions of Object-Oriented Programming (OOP))
- PyData Stack (Numpy – arrays, slicing, vectorized operations – , Pandas – series, slicing, indexing, transformations – , Matplotlib – basic plotting only – , Scikit-Learn – linear regression, pipelines, one-hot encoding, normalization/scaling, grid search, hyper-parameter optimization)
- Machine Learning Concepts (supervised learning: regression and classification; loss functions: RMSE, cross-entropy; train-validation-test split; evaluation metrics (R-squared, precision, recall, accuracy, confusion matrix)).

Outline for this course

Chapter 1 - Course Introduction

Chapter 2 - PyTorch, Datasets, and Models

Chapter 3 - Building Your First Dataset

Chapter 4 - Training Your First Model

Chapter 5 - Building Your First Hugging Face Dataset

Chapter 6 - Transfer Learning and Pretrained Models

Chapter 7 - Pretrained Models for Computer Vision

Chapter 8 - Pretrained Models for Natural Language Processing

Chapter 9 - Image Classification with Torchvision

Chapter 10 - Fine-Tuning Pretrained Models for Computer Vision

Chapter 11 - Serving Models with TorchServe

Chapter 12 - Datasets and Transformations for Object Detection and Image Segmentation

Chapter 13 - Models for Object Detection and Image Segmentation

Chapter 14 - Object Detection Evaluation

Chapter 15 - Word Embeddings and Text Classification

Chapter 16 - Contextual Word Embeddings with Transformers

Chapter 17 - Hugging Face Pipelines for NLP Tasks

Chapter 18 - Q&A, Summarization, and LLMs