

PyTorch in Practice: An Applications-First Approach (LFD473)

The course begins with an overview of PyTorch, including model classes, datasets, data loaders and the training loop. Next the role and power of transfer learning is addressed 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: 32 Hours

Prerequisites for this course

- While there are no formal prerequisites, students should have some knowledge of Python (notions of object-oriented programming), PyData Stack (Numpy, Pandas, Matplotlib, Scikit-Learn), and Machine Learning concepts (supervised learning, loss functions, train-validation-test split, evaluation metrics).

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 - Building Your First Hugging Face Dataset

Chapter 7 - Transfer Learning and Pretrained Models

Chapter 8 – Pretrained Models for Computer Vision

Chapter 9 – Natural Language Processing

Chapter 10 - Image Classification with Torchvision

Chapter 11 – Fine-Tuning Pretrained Models for Computer Vision

Chapter 12 - Serving Models with TorchServe

Chapter 13 - Datasets and Transformations for Object Detection and Image Segmentation

Chapter 14 - Models for Object Detection and Image Segmentation

Chapter 15 - Models for Object Detection Evaluation

Chapter 16 - Word Embeddings and Text Classification

Chapter 17 - Contextual Word Embeddings with Transformers