Course Outline

Applied Deep Learning with PyTorch

Course Summary

Description

Starting with the basics of deep learning and their various applications, Applied Deep Learning with PyTorch shows you how to solve trending tasks, such as image classification and natural language processing by understanding the different architectures of the neural networks

Applied Deep Learning with PyTorch takes your understanding of deep learning, its algorithms, and its applications to a higher level. The course begins by helping you browse through the basics of deep learning and PyTorch. Once you are well versed with the PyTorch syntax and capable of building a single-layer neural network, you will gradually learn to tackle more complex data problems by configuring and training a convolutional neural network (CNN) to perform image classification. As you progress through the chapters, you'll discover how you can solve an NLP problem by implementing a recurrent neural network (RNN).

Objectives

By the end of the course, students will be able to:

- Detect a variety of data problems to which you can apply deep learning solutions
- Learn the PyTorch syntax and build a single-layer neural network with it
- Build a deep neural network to solve a classification problem
- Develop a style transfer model
- Implement data augmentation and retrain your model
- Build a system for text processing using a recurrent neural network

Topics

- Introduction to Deep Learning and PyTorch
- Building Blocks of Neural Networks
- A Classification Problem Using DNN
- Convolutional Neural Networks
- Style Transfer
- Analyzing the Sequence of Data with RNNs

Audience

Applied Deep Learning with PyTorch is designed for data scientists, data analysts, and developers who want to work with data using deep learning techniques. Anyone looking to explore and implement advanced algorithms with PyTorch will also find this course useful. Some working knowledge of Python and familiarity with the basics of machine learning are a must. However, knowledge of NumPy and pandas will be beneficial, but not essential.

Duration

Two Days

Course Outline

Applied Deep Learning with PyTorch

Course Outline

I. *Introduction to Deep Learning and PyTorch*
   A. Understanding Deep Learning
   B. PyTorch Introduction

II. *Building Blocks of Neural Networks*
   A. Introduction to Neural Networks
   B. Data Preparation
   C. Building a Neural Network

III. *A Classification Problem Using DNN*
   A. Problem Definition
   B. Dealing with an Underfitted or Overfitted Model
   C. Deploying Your Model

IV. *Convolutional Neural Networks*
   A. Building a CNN
   B. Data Augmentation
   C. Batch Normalization

V. *Style Transfer*
   A. Style transfer
   B. Implementation of Style Transfer Using the VGG-19 Network Architecture

VI. *Analyzing the Sequence of Data with RNNs*
   A. Recurrent Neural Networks
   B. Long Short-Term Memory Networks (LSTMs)
   C. LSTM Networks in PyTorch
   D. Natural Language Processing (NLP)
   E. Sentiment Analysis in PyTorch