

Deep Learning in Python

Length: 5 Days

Summary: Deep learning is the machine learning technique behind the most exciting capabilities in diverse areas like robotics, natural language processing, image recognition and artificial intelligence. In this course, you'll gain hands-on, practical knowledge of how to use deep learning with Keras. In addition, you will also learn Python libraries suitable for Machine Learning, the differences between Supervised vs. Unsupervised Machine Learning, how Statistical Modeling relates to Machine Learning, and how to do a comparison of each along with different ways machine learning affects society

Experience Needed: Class participants should have experience in object-oriented programming or at the minimum understand the principles of object-oriented programming. Attendees should have taken a college-level algebra course, which covered the following:

- Continuous functions of one or more variables
- Linear and non-linear functions
- Exponential functions
- Basic differential calculus
- Basic statistics
- Basic linear algebra

COURSE CONTENT

DAY 1

1: Comprehensive Intro to Python and Basics of deep learning & neural networks

In this module, you'll become familiar with the fundamental concepts and terminology used in deep learning and understand why deep learning techniques are so powerful today. We will also review tools we will use throughout the course include (but not limited to) Numpy, Matplotlib, and Jupyter. You'll build simple neural networks yourself and generate predictions with them.

DAY 2

2: Building deep learning models with TensorFlow, TensorBoard, and Keras

In this module, you'll use some of the common deep learning APIs including Keras, TensorFlow, and to visualize learning using TensorBoard libraries to build deep learning models for both regression as well as classification! You'll learn about the Specify-Compile-Fit workflow that you can use to make predictions and by the end of this module, you'll have all the tools necessary to build deep neural networks.

3: Optimizing a neural network with backward propagation

Here, you'll learn how to optimize the predictions generated by your neural networks. You'll do this using a method called backward propagation, which is one of the most important techniques in deep learning. Understanding how it works will give you a strong foundation to build from in the second half of the course.

DAY 3 & (part of) DAY 4

4: Optimizing datasets for deep learning using Pandas and Fine-tuning Keras models

Here, you'll learn how to optimize your deep learning models in Keras. You'll learn how to validate your models, understand the concept of model capacity, and experiment with wider and deeper networks.

DAY 4

5: Regression

Here, you will learn about Linear Regression, Non-linear Regression, and Model evaluation methods.

6: Classification

Here, you will learn about K-Nearest Neighbor, Decision Trees, Logistic Regression, Support Vector Machines, and Model Evaluation.

DAY 5

7: Unsupervised Learning

Here, we will discuss and demonstrate K-Means Clustering, Hierarchical Clustering, and Density-Based Clustering

8: Review

Here we will review everything we have learned during the week and demonstrate ways we can use what we have covered and put it into real-world scenarios

