

Introduction to Deep Learning with TensorFlow and Keras – 3 Days

Course Summary

Description

This course introduces Deep Learning concepts and Tensorflow and Keras libraries to students.

The abundance of data and affordable cloud scale has led to an explosion of interest in Deep Learning. Google has open sourced a library called TensorFlow which has become the de-facto standard, allowing state-of-the-art machine learning done at scale, complete with GPU-based acceleration.

Objectives

After taking this course, students will learn:

- Deep Learning concepts
- TensorFlow and Keras
- Create neural networks with Tensorflow and Keras
- Learn to use tools like Tensorboard to help with training neural networks
- We will build neural networks to solve the following problems
- Regression
- Classification
- Computer vision / Image analytics
- Text analytics
- Time series

Topics

- Introduction to Deep Learning
- Introducing TensorFlow
- Introducing Keras
- Deep Learning Concepts
- Feedforward Network
- Computer Vision
- Recurrent Neural Networks
- Transfer Learning
- Workshop (Time permitting)

Audience

This course is designed for Developers, Data analysts, and Data Scientists

Prerequisites

- Basic knowledge of Python language and Jupyter notebooks is assumed. Even if you haven't done any Python programming, Python is such an easy language to learn quickly. We will provide Python resources.

Duration

Three Days

Introduction to Deep Learning with TensorFlow and Keras – 3 Days

Course Outline

I. Introduction to Deep Learning

- A. Understanding Deep Learning use cases
- B. Understanding AI / Machine Learning / Deep Learning
- C. Data and AI
- D. AI vocabulary
- E. Hardware and software ecosystem
- F. Understanding types of Machine Learning (Supervised / Unsupervised / Reinforcement)

II. Introducing TensorFlow

- A. TensorFlow intro
- B. TensorFlow features
- C. Execution graph
- D. TensorFlow on GPU and TPU
- E. TensorFlow API
 - Lab: Setting up and Running TensorFlow

III. Introducing Keras

- A. Keras Intro
- B. Keras concepts (models, layers)
- C. Using Keras API
 - Lab

IV. Deep Learning Concepts

- A. Introducing Perceptrons
- B. Linear Perceptrons
- C. Activation Functions (Sigmoid, Tanh, Relu, Softmax)
- D. Backpropagation
- E. Optimizers (Gradient Descent, Adam, RMSProp)
- F. Loss functions for regressions and classifications
- G. Vanishing/exploding gradient problem
 - Lab: Tensorflow playground

V. Feedforward Network

- A. FFNN architecture
- B. Input layer, output layer
- C. Hidden layers and Deep neural networks
- D. Sizing neural networks
 - Lab: Feedforward Neural Networks

VI. Computer Vision

- A. Introducing Convolutional Neural Networks (CNN)
- B. CNN architecture
- C. CNN concepts
 - Lab: Image recognition using CNNs

VII. Recurrent Neural Networks

- A. Introducing RNNs
- B. RNN architecture
- C. RNN concepts
- D. LSTM (Long Short-Term Memory) networks
- E. LSTM architecture
 - Lab: RNNs for text and sequence prediction

VIII. Transfer Learning

- A. Understanding transfer learning
- B. Customizing available models
 - Lab: transfer learning lab
 - Lab: Benchmarking performance on CPU and GPU

IX. Workshop (Time permitting)

- A. Students will work in teams to solve a real-world use case