

Intro to Deep Learning With TensorFlow



Course Description

The abundance of data and affordable cloud scale has led to an explosion of interest in Deep Learning. Google has released an excellent library called TensorFlow to open-source, allowing state-of-the-art machine learning done at scale, complete with GPU-based acceleration. This course introduces Deep Learning concepts and TensorFlow library to students. Tired of shallow courses that just show code without explaining more? We strive to provide a fundamental understanding of Deep Learning concepts and how TensorFlow implements them.

Learning Objectives

You will learn:

- + Introduction to Machine Learning
- + Deep Learning concepts
- + TensorFlow library
- + Writing TensorFlow applications (CNN, RNN)
- + Using TF tools
- + High level libraries : Keras

Who Should Attend

- + Developers
- + Data Analysts
- + Data Scientists

Prerequisites

- + Basic knowledge of Python language and Jupyter notebooks is assumed
- + Basic knowledge of Linux environment would be beneficial
- + Some Machine Learning familiarity would be nice, but not necessary

Lab Environment

- + Cloud servers will be provided students for installation, administration, and lab work. Students would need a SSH client and a browser to access the cluster. Most labs will be in jupyter notebook format.

Intro to Deep Learning With TensorFlow

(3 Days)

Course Content

1. Introduction to Machine Learning

- + Understanding Machine Learning
- + Supervised vs. Unsupervised Learning
- + Regression
- + Classification
- + Clustering

2. Introducing TensorFlow

- + TensorFlow intro
- + TensorFlow Features
- + TensorFlow Versions
- + GPU and TPU scalability
- + Lab: Setting up and Running TensorFlow

3. The Tensor: The Basic Unit of TensorFlow

- + Introducing Tensors
- + TensorFlow Execution Model
- + Lab: Learning about Tensors

4. Single Layer Linear Perceptron Classifier With TensorFlow

- + Introducing Perceptrons
- + Linear Separability and XOR Problem
- + Activation Functions
- + Softmax output
- + Backpropagation, loss functions, and Gradient Descent
- + Lab: Single-Layer Perceptron in TensorFlow

5. Hidden Layers: Intro to Deep Learning

- + Hidden Layers as a solution to XOR problem
- + Distributed Training with TensorFlow
- + Vanishing Gradient Problem and ReLU
- + Loss Functions
- + Lab: Feedforward Neural Network Classifier in TensorFlow

6. High level TensorFlow: tf.learn

- + Using high level TensorFlow
- + Developing a model with tf.learn
- + Lab: Developing a tf.learn model

7. Convolutional Neural Networks in TensorFlow

- + Introducing CNNs
- + CNNs in TensorFlow
- + Lab : CNN apps

8. Introducing Keras

- + What is Keras?
- + Using Keras with a TensorFlow Backend
- + Lab: Example with a Keras

9. Recurrent Neural Networks in TensorFlow

- + Introducing RNNs
- + RNNs in TensorFlow
- + Lab: RNN

10. Long Short Term Memory (LSTM) in TensorFlow

- + Introducing RNNs
- + RNNs in TensorFlow
- + Lab: RNN

11. Conclusion

- + Summarize features and advantages of TensorFlow
- + Summarize Deep Learning and How TensorFlow can help
- + Next steps

To register for an ePlus cloud training course, contact us today.

Call: 888.482.1122

Email: CloudServices@eplus.com

Web: www.eplus.com/cloud


**Where Technology
Means More®**