

Deep Learning With BigDL

Course Summary

Description

This course introduces Deep Learning concepts and BigDL library to students. The abundance of data and affordable cloud scale has led to an explosion of interest in Deep Learning. Intel has released an excellent library called BigDL to open-source, allowing state-of-the-art machine learning done at scale

Objective

- Introduction to Deep Learning
- Deep Learning concepts
- Introducing the BigDL Library
- RDD API versus Pipeline API
- Using BigDL with Tensorflow and Caffe models.
- Visualizing Deep Learning Data with Tensorboard
- Scaling Deep Learning with to massive distributed data

Topics

- Introduction to Deep Learning
- Introducing BigDL
- BigDL Execution Model
- Single Layer Linear Perceptron Classifier With BigDL
- Hidden Layers: Intro to Deep Learning
- High level BigDL: the Pipeline API and Dataframes
- Convolutional Neural Networks in BigDL
- Tensorflow and Caffe Models in BigDL
- Recurrent Neural Networks in BigDL
- Long Short Term Memory (LSTM) in BigDL
- Conclusion

Audience

This course is designed for Developers, Data analysts, and data scientists.

Prerequisite

- 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

Duration

Three Days

Deep Learning With BigDL

Course Outline

- I. **Introduction to Deep Learning**
 - A. Understanding Deep Learning
 - B. Activation Functions, Loss Functions, and Gradient Descent
 - C. Training and Validation
 - D. Regression vs. Classification
- II. **Introducing BigDL**
 - A. BigDL intro
 - B. BigDL Features
 - C. BigDL Versions
 - D. BigDL and Hadoop
 - E. Apache Spark
 - Lab: Setting up and Running BigDL
- III. **BigDL Execution Model**
 - A. Introducing BigDL's Execution Model
 - B. Understanding Basic BigDL Layers
 - C. Understand how BigDL runs on top of Apache Spark workloads.
 - Lab: BigDL Layers
- IV. **Single Layer Linear Perceptron Classifier With BigDL**
 - A. Introducing Perceptrons
 - B. Linear Separability and Xor Problem
 - C. Activation Functions
 - D. Softmax output
 - E. Backpropagation, loss functions, and Gradient Descent
 - Lab: Single-Layer Perceptron in BigDL
- V. **Hidden Layers: Intro to Deep Learning**
 - A. Hidden Layers as a solution to XOR problem
 - B. Distributed Training with BigDL
 - C. Vanishing Gradient Problem and ReLU
 - D. Loss Functions
 - E. Using Tensorboard to Visualize Training
 - Lab: Feedforward Neural Network Classifier in Tensorflow
- VI. **High level BigDL: the Pipeline API and Dataframes**
 - A. Using high level BigDL
 - B. Developing a model with pipeline API
 - Lab: Developing a pipeline API model
- VII. **Convolutional Neural Networks in BigDL**
 - A. Introducing CNNs
 - B. CNNs in BigDL
 - C. CNN's for image classification
 - Lab : Convolutional Neural Networks
- VIII. **Tensorflow and Caffe Models in BigDL**
 - A. How to export and import models in BigDL
 - B. Transfer Learning with BigDL
 - C. ImageNet and other pre-trained models.
 - Lab: Example with a Caffe model
- IX. **Recurrent Neural Networks in BigDL**
 - A. Introducing RNNs
 - B. RNNs in BigDL
 - Lab: RNN
- X. **Long Short Term Memory (LSTM) in BigDL**
 - A. Introducing RNNs
 - B. RNNs in BigDL
 - Lab: RNN
- XI. **Conclusion**
 - A. Summarize features and advantages of BigDL
 - B. Summarize Deep Learning and How BigDL can help
 - C. Next steps