

Deep Learning (Artificial Neural Networks) using Python

Duration: 30 hours live training + 30 hours practice
Hands on, Instructor Led, Use-Case Project Based, Classroom Training
Weekday and Weekend batches

Introduction

What is Deep Learning?

Theano

TensorFlow

Keras

Project: Develop Large Models on GPUs Cheaply In the Cloud

Artificial Neural Networks

Multilayer Perceptrons (MLP)

Develop your first Neural Network with Keras

Evaluate the performance of Deep Learning Models

Use Keras Models with Scikit-Learn for General Machine Learning

Advanced Multilayer Perceptrons and Keras

Reduce Overfitting with Dropout Regularization

Lift Performance with Learning Rate Schedules

Project: Multiclass Classification of Flower Species

Project: Binary Classification of Sonar Returns

Project: Regression of Boston House Prices

Convolutional Neural Networks (CNN)

Basics of ConvNets

Convolutional model development and application

Deep convolutional models

Object detection

Improve Model Performance with Image Augmentation

Project: Handwritten Digit Recognition

Project: Object Recognition in Photographs

Project: Predict Sentiment from Movie Reviews

Recurrent Neural Networks (RNNs and LSTM)

Long Short-Term Memory Networks

Time Series Prediction with Multilayer Perceptrons

Time Series Prediction with LSTM Recurrent Neural Networks

Project: Sequence Classification of Movie Reviews

Understanding Stateful LSTM Recurrent Neural Networks

Project: Text Generation with Alice in Wonderland



Minimum system requirements:

- 4 GB RAM
- i3 and above processors
- 100 GB free hard disk space

Environment: Ubuntu version 14.0 and above

