




Deep Learning Workshop



Deep Learning is a new area of Machine Learning research, which has been introduced with the objective of moving Machine Learning closer to one of its original goals: Artificial Intelligence.

This course is intended to give a holistic understanding on Deep Learning and its applications. It would cover:

- Basics of Neural Networks and Deep Learning
- Mathematics and Programming for Deep Learning
- Detailed application oriented approach towards learning Deep Learning
- Tensor Flow and Keras based practicals

SUMMARY

Objective

Machine learning is one of the fastest-growing and most exciting fields out there, and deep learning represents its true bleeding edge. In this course, you'll develop a clear understanding of the motivation for deep learning, and design intelligent systems that learn from complex and/or large-scale datasets.

We'll show you how to train and optimize basic neural networks, convolutional neural networks, and long short term memory networks. Complete learning systems in TensorFlow and Keras will be introduced with working examples.

You will learn to solve new classes of problems that were once thought prohibitively challenging, and come to better appreciate the complex nature of human intelligence as you solve these same problems effortlessly using deep learning methods.

Hardware and Software

- A laptop running 64 bit OS (Linux/OSX/Windows)
- Minimum 8 GB RAM
- There are several installers necessary, we would be sharing the installations in advance.

Pre-requisites

- A basic understanding of programming concepts
- Programming knowledge of python is a plus

Outline

Introduction to Deep Learning

- Deep learning in the context of Machine Learning and Artificial Intelligence
- What is a neural network?
- Why Deep Learning?

Dive deep into Deep Learning

- Train a deep network using TensorFlow/Keras
- How to choose between deep neural networks?
- Effectively regularize a simple deep network.
- Train a competitive deep network via model exploration and hyper-parameter tuning.

Convolutional Neural Networks**

- Train a simple convolutional neural net.
- Explore the design space for convolutional nets.
- Transfer Learning
- Overview of different architectures – VGG; AlexNet etc.

Recurrent Neural Networks

- BPTT; Vanishing Gradient Problem
- LSTM; GRU
- Character Prediction using Keras

Auto-encoders

- Single layer auto-encoders; Stacked Auto-encoders
- Dimension Reduction
- Anomaly Detection

Deep Natural Language Processing

- Embedding Concept in the context of Recommendation Engines
- Word Embedding
- Word Vector Representations: Glove
- Word Vector Representations: word2vec
- Text Classification using Word Embedding, CNN and RNN

Model Debugging and deployment

- TensorFlow Debugger
- Tensorboard
- Deploying models on AWS Platform

About Instructor



Naveen Kumar Bhansali an alumnus IIM Bangalore, is a Big Data Analyst with over 12 years of experience in Data Analytics and Data Engineering in industries as varied as Telecom, Energy, Automotive, and Finance.

He is a guest faculty at IIM Bangalore and has also presented numerous plenary sessions and papers at the International Conference on Business Analytics and Intelligence.

Linkedin: <https://www.linkedin.com/in/nkbhansali/>