



Deep Learning

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 optimise basic neural networks, convolutional neural networks, and long short term memory networks. Complete learning systems in TensorFlow 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

- Understanding of programming concepts
- Programming knowledge of python is mandatory.

Outline

Session 1 - Deep Learning Concepts

- What is a neural network?
- Why Deep Learning?
- How to choose between deep neural networks?
- An old problem: The Vanishing Gradient
- Deep Belief Networks

Session 2 - Deep Learning Concepts Continued

- Convolutional Networks
 - Recurrent Nets
 - Auto-encoders
 - Recursive Neural Tensor Nets
 - Where is Deep Learning Uses
-

DEEP LEARNING

Session 3 - Platforms for Deep Learning

- What is a Deep Learning Platform?
- H2O.ai
- Dato GraphLab

Session 4 - Deep Learning Software Libraries

- What are several Deep Learning Libraries?
- Tensorflow
- Theano
- Deeplearning4j
- Torch
- Caffe

Session 5: Lab

- Introduction to tensor flow, setting up the Computation Graphs and MLPs
- Tensor flow Language Modeling with LSTMs
- One end to end use case

Session 6: Deep Neural networks

- Theory of deep neural networks.
- Train a simple deep network.
- Effectively regularise a simple deep network.
- Train a competitive deep network via model exploration and hyper-parameter tuning.

Session 7: Convolutional Neural Networks**

- Train a simple convolutional neural net.
- Explore the design space for convolutional nets.

Session 8: Lab

- Setting up ANN
- Setting up DNN
- Setting up CNN
- Setting up RNN

Session 9: Full Case Study: Deep Learning and NLP

- Introduction to NLP and Deep Learning
 - Word Vector Representations: word2vec
 - Word Window Classification and Neural Networks
 - End-to-end models for Speech Processing
-

Session 10: Google Cloud Machine Learning platform

- Setting up Deep Learning algorithms on Google Cloud Platform
- Practical description of the platform
- Further references

COURSE STRUCTURE

Day 1: Understanding of Neural Networks, Deep Learning and Libraries

This day will be about underlying concepts of neural networks, deep learning and other libraries

Topic	Session	From	To
Basics of Neural Networks and Deep Learning	1	9 AM	10:15 AM
Deep Learning Concepts	2	10:30 AM	11:45 AM
Platforms and Languages for Deep Learning	3	12:00 PM	1:15 PM
Open Source Libraries for Deep Learning	4	2:15 PM	3:30 PM
Lab 1: Installation, and Basics of Tensor Flow	5	3:45 PM	5:00 PM

Day 2: Advanced Concepts of Deep Learning and Applications

Application driven advanced concepts understanding of Deep Learning

Topic	Session	From	To
Understanding ANN, CNN, RNN	1	9 AM	10:15 AM
Lab2: Designing ANN using Tensor Flow	2	10:30 AM	11: 45 AM
Lab2 (Continued): Designing CNN, RNN using Tensor Flow	3	12:00 PM	1:15 PM
NLP using Deep Learning	4	2:15 PM	3:30 PM
Deep Deployment in Google Cloud Platform	5	3:45 PM	5:00 PM

ABOUT INSTRUCTORS

Vidyasagar Nallapati

Vidyasagar N is a Data Scientist and Software Architect from Bangalore with close to 10 years of experience. He loves all the things mathematics, statistics, machine learning, developer tools, automation and anything that will make systems better and easier, which brought him to data science, system architecture, building large scale distributed services and designing rock-solid data pipelines. An alumni of IIT BHU and IIM Bangalore, Vidyasagar when not writing code can be found travelling, playing music and indulging in coffee.

Jitendra R

Jitendra, an alumnus of University of Mumbai and IIM Bangalore, is a Data Scientist with over 10 years of experience in Data Analytics and Data Engineering in industries as varied as Telecom, Energy, Pharma, Automotive, and Maritime. He is a joint secretary in the Analytics Society of India (ASI) and additionally guest faculty at IIM Bangalore. He has also presented numerous plenary sessions and papers at the International Conference on Business Analytics and Intelligence. On the academic front his work in analytics has been published and distributed by Harvard as case studies.

Naveen Bhansali

Naveen, 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.
