



E0-334 Aug 3:1

Deep Learning for Natural Language Processing

Instructor

Shirish Shevade

Email: shirish@iisc.ac.in

Teaching Assistant

Email:

Department: Computer Science and Automation

Course Time: Wed., Fri., 14:00-15:30 Hrs

Lecture venue: CSA 117

Detailed Course Page:

Announcements

Brief description of the course

Natural Language Processing (NLP) is an important technology having implications in human-computer interaction. A variety of problems in NLP can be solved using traditional machine learning algorithms. With the availability of a lot of data and advances in high performance computing, Deep Learning models have shown a lot of promise. The aim of this course is to study different types of neural networks and build these networks for solving practical problems in natural language processing.

Prerequisites

A course on Machine Learning or equivalent

Syllabus

Introduction, Multilayer Neural Networks, Back-propagation, Training Deep Networks; Simple word vector representations: word2vec,

GloVe; sentence, paragraph and document representations. Recurrent Neural Networks; Convolutional Networks and Recursive Neural Networks; GRUs and LSTMs; building attention models; memory networks for language understanding. Design and Applications of Deep Nets to Language Modeling, parsing, sentiment analysis, machine translation etc.

Course outcomes

In this course, students will learn to implement, train and invent neural network models and make these models work on practical problems in Natural Language Processing.

Grading policy

10% for assignments

40% for final exam

50% for research papers presentation and course project (continuous evaluation done throughout the semester)

Assignments

Resources

1. Ian Goodfellow and Yoshua Bengio and Aaron Courville. Deep Learning. MIT Press, 2016.

2. Recent Literature