DEEP LEARNING – PART 2

COMPUTER SCIENCE & ENGINEERING

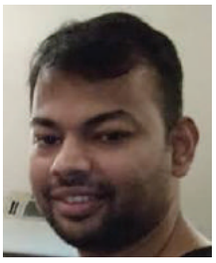

PROF. MITESH M. KHAPRA
Dept. of Computer Science and Engineering
IIT Madras

TYPE OF COURSE : New | Elective | UG/PG/PhD
INTENDED AUDIENCE : BE/BTech/ME/MTech/MS/PhD
PRE-REQUISITES : Deep Learning

COURSE DURATION : 8 weeks (25 Feb'19 - 19 Apr'19)
EXAM DATE : 28 Apr 2019

INDUSTRIES APPLICABLE TO : Google, Microsoft, Amazon, Adobe, IBM

COURSE OUTLINE :

In this course, we will cover topics which lie at the intersection of Deep Learning and Generative Modeling. We will start with basics of joint distributions and build up to Directed and Undirected Graphical Models. We will then make a connection between Graphical Models and Deep Learning by having an in-depth discussion on Restricted Boltzmann Machines, Markov Chains and Gibbs Sampling for training RBMs. Finally, we will cover more recent Deep Generative models such as Variational Autoencoders, Generative Adversarial Networks and Autoregressive Models.

ABOUT INSTRUCTOR :

Mitesh M. Khapra is an Assistant Professor in the Department of Computer Science and Engineering at IIT Madras. While at IIT Madras he plans to pursue his interests in the areas of Deep Learning, Multimodal Multilingual Processing, Dialog systems and Question Answering. Prior to that he worked as a Researcher at IBM Research India. During the four and half years that he spent at IBM he worked on several interesting problems in the areas of Statistical Machine Translation, Cross Language Learning, Multimodal Learning, Argument Mining and Deep Learning. This work led to publications in top conferences in the areas of Computational Linguistics and Machine Learning. Prior to IBM, he completed his PhD and M.Tech from IIT Bombay in Jan 2012 and July 2008 respectively. His PhD thesis dealt with the important problem of reusing resources for multilingual computation. During his PhD he was a recipient of the IBM PhD Fellowship and the Microsoft Rising Star Award. He is also a recipient of the Google Faculty Research Award, 2018.

COURSE PLAN :

Module 01 : A brief introduction to Directed Graphical Models

Module 02 : A brief introduction to Markov Networks, Using joint distributions for classification and sampling, Latent variables

Module 03 : Restricted Boltzmann Machines, Unsupervised Learning, Motivation for Sampling, Markov Chains, Gibbs Sampling for training RBMs, Contrastive Divergence for training RBMs

Module 04 : Variational Autoencoders, Autoregressive models, GANs