MCA526: DEEP LEARNING [3-0-1] New Elective

Course Outcomes

On successful completion of this course, the student will be able to:

CO1: describe the feedforward and deep networks
CO2: design single and multi-layer feed-forward deep networks and tune various hyper-parameters.
CO3: analyze the performance of deep networks.

Syllabus:

Introduction: Historical context and motivation for deep learning; basic supervised classification task, optimizing logistic classifier using gradient descent, stochastic gradient descent, momentum, and adaptive sub-gradient method.

Neural Networks: Feedforward neural networks, deep networks, regularizing a deep network, model exploration, and hyperparameter tuning.

Convolution Neural Networks: Introduction to convolution neural networks: stacking, striding and pooling, applications like image, and text classification.

Structuring Machine Learning Projects: Orthogonalization, evaluation metrics, train/dev/test distributions, size of the dev and test sets, cleaning up incorrectly labeled data, bias and variance with mismatched data distributions, transfer learning, multi-task learning

Readings :

1. Jeff Heaton, **Deep Learning and Neural Networks**, Heaton Research Inc, 2015
2. Mindy L Hall , **Deep Learning** , VDM Verlag , 2011
3. Li Deng (Author), Dong Yu , **Deep Learning: Methods and Applications (Foundations and Trends in Signal Processing)** , Now Publishers Inc , 2009
4. Ian Goodfellow, **Deep Learning**, MIT Press, 2016