CSE4121/CSE6061: Deep Learning

Programme: B. Tech. (CSE) Year: III Semester : 6
Course : Program Elective Credits : 3 Hours : 40

Course Context and Overview (100 words):
This course covers the most successful form of artificial intelligence, deep learning. We will be covering linear regression, logistic regression, deep neural networks, convolutional and recurrent neural networks. The course will also focus on optimization techniques like gradient descent and its variants. We will also cover many regularization techniques for neural networks, including but not limited to L1, L2 and dropout. Programming will be an important component of the course. We will be using Python as our primary language. For implementation of algorithms, we will be using Tensorflow and Keras. The course will be equally inclined towards theory and programming.

Prerequisites Courses: NIL

Course outcomes (COs):

On completion of this course, the students will have the ability to:	
CO1:	Prioritize the collection and usefulness of data for a particular deep learning task
CO2:	Apply theory and implementation learnt in the course to real world problems
CO3:	Judge whether a particular problem can be solved using deep learning or not
CO4:	Critically analyze which architecture to use for a specific problem
CO5:	Design and implement deep learning algorithms using Tensorflow and Keras framework.

Course Topics:

Topics	Lecture Hours
UNIT – I Introduction Linear and logistic regression. Cost function for logistic regression.	4
UNIT – II Deep Neural Networks Generalization of logistic regression to deep neural networks. Cost functions. Optimization algorithms: Gradient descent, Stochastic gradient descent, Momentum, RMSprop, Adam.	9
UNIT – III Regularization Techniques	6

Underfitting and overfitting of neural networks: bias and variance. L1, L2 and dropout regularization techniques, hyperparameter tuning.	
UNIT – IV Deep Learning for Computer Vision Basics of CNN: convolutions and pooling. Detailed understanding of Alexnet, ResNet, VGG-16, VGG-19 and inception architectures. Their implementations. Object recognition and face recognition.	12
UNIT – V Deep Learning for Natural Language Processing Basics of RNN, LSTM, GRU, Bidirectional RNN, deep RNNs. Representations of words as vectors. One hot encoding and word embeddings. Learning word embeddings using word2vec, glove and fasttext.	9

Text Book:

1. I. Goodfellow, Y. Bengio, and A. Courville, "*Deep Learning"*, Cambridge, MA: MIT Press, 2016.

Additional Resources (NPTEL, MIT Video Lectures, Web resources etc.):

1. Stanford CS230: Deep Learning
2. MIT 6.S191 Introduction to Deep Learning (January-March 2018)
3. Stanford CS224d: Deep Learning for Natural Language Processing (January-March, 2018)
4. Stanford CS231n: Convolutional Neural Networks for Visual Recognition (Spring 2018)
5. Coursera specialization on Deep Learning

Evaluation Methods:

Item	Weight Percentage
Midterm	25
Endterm	50
Term paper	15
Quizzes	10

Prepared By: Kshitiz Verma
Last Update: 15 October, 2019