Name of the Programme: MCA
Course Code: CSA-608
Title of the Course: Deep Learning
Number of Credits: 4(2L+2T+0P)
Effective from AY: 2022-23

Prerequisites for the course	Familiarity with linear algebra, probability theory, machine learning , familiarity with python	
Objectives:	This course is aimed at any one who wishes to explore deep learning from scratch. This course offers a practical hands on exploration of deep learning, avoiding mathematical notation,preferring instead to explain quantitative concepts through programming using python API	
Content:	Introduction :- what is deep learning ?- Artificial Intelligence, machine learning and deep learning -learning representation from data-"the deep " in deep learning -understanding how deep learning works - what deep learning has achieved so far.	2 hours
	Revision of Fundamentals of machine learning- probabilistic modeling – early neural networks- kernel methods-decision tree, random forest and gradient boosting machines -back to neural networks- what makes deep learning different-the modern machine learning landscape .	3 hours
	Four branches of machine learning -supervised -unsupervised-self-supervised – reinforcement learning – evaluating machine learning models – data processing, feature engineering- overfitting and underfitting -universal workflow of machine learning	3 hours
	The mathematical building block of neural networks – a first look at neural networks – data representation for neural networks- the gears of neural networks :Tensor operations- the engine of neural networks : Gradient -based optimization.	3 hours
	Neural networks – anatomy of neural networks- building blocks of deep learning -models of layers -loss functions and optimizers-keys to configuring the learning process.-introduction to keras -keras,tensor flow, theoano and CNTK – developing with keras -setting up a deep learning workstation -case studies – classification movie reviews – classification newswires -predicting house prices.	3 hours
	Deep Learning for computer vision – Introduction to convnets – training convnets from scratch on small data sets – using pre trained convnet – visualizing what convnets learn	3 hours
	Deep learning for text and sequences – working with text -one-hot encoding of words and characters -using word embeddings-understanding recurrent neural networks – A recurrent layer in Keras -understanding LSTM and GRU layers- A concrete LSTM example in Keras.	3 hours
	Advanced use of recurrent neural networks- A temperature-forecasting problem – preparing the data – a common-sense, non machine learning baseline-using recurrent drop out to fight overfitting- stacking recurrent layers-using bidirectional RNNs – sequences processing with convnets	5 hours
	Generative deep learning – text generation with LSTM- deep Dream – neural style transfer-generative images with variational autoencoders- introduction to generative adversarial networks.	5 hours
	Assignments to be carried out during Tutorial Slot - ● Assignment 1 - Logistic Regression with a Neural Network mindset	10 * 3 = 30 hours

	● Assignment 2 - Planar data classification with one hidden layer ● Assignment 3 - Building your Deep Neural Network: Step by Step ● Assignment 4 - Deep Neural Network for Image Classification: Application ● Assignment 5 – Initialization and performance of model, Regularization and whether it helps eliminate overfitting, Gradient Checking with model used, Optimization Methods used for every model ● Assignment 6- TensorFlow Tutorial ● Assignment 7 - Convolution model Step by Step demo ● Assignment 8 - Convolution model Application for image classification ● Assignment 9- Keras Tutorial - Autonomous driving application - Car Detection, Face Recognition ● Assignment 10 - Art Generation with Neural Style transfer	
Pedagogy:	Lectures/ tutorials/lab assignments/self-study	
References/ Readings	Main Reading :- 1. Introduction to Statistical Learning, Gareth James, Daniela Witten, Trevor Hastie, Robert Tibshirani, Springer, 2013. 2. EthemAlpaydin, Introduction to Machine Learning, MIT Press. 3. Richard O. Duda, Peter E. Hart, David G. Stork Pattern Classification,. 4. Peter Flach , Machine Learning , Cambridge 5. Christopher M. Bishop,Pattern recognition and machine Learning, springer. 6. Deep Learning, Ian Good fellow, MIT press 7. Tom Michele, Machine Learning, McGraw-Hill.	
Course Outcomes	By the end of the course , students will be able to: • understand a wide variety of deep learning algorithms. • understand how to apply a variety of learning algorithms to data. • understand how to perform evaluation of learning algorithms and model selection. • equip themselves with a general understanding of deep learning.	

(Back to top)