Semester II
Name of the Programme: M.Sc. in Artificial Intelligence
Course Code: CSI-508
Title of the Course: Deep Learning
Number of Credits: 2(2L-0T-0P)
Effective from AY: 2023-24

Prerequisites for the course	Programme prerequisites	
Objectives	To study the basics of Neural Networks and their various variants such as the Convolutional Neural Networks and Recurrent Neural Networks, to study the different ways in which they can be used to solve problems in various domains such as Computer Vision, Speech and NLP.	
Content	History of Deep Learning, McCulloch Pitts Neuron, Thresholding Logic, Perceptron Learning Algorithm and Convergence	1 hours
	Multilayer Perceptrons (MLPs), Representation Power of MLPs, Sigmoid Neurons, Gradient Descent	1 hours
	Feedforward Neural Networks, Representation Power of Feedforward Neural Networks, Backpropagation	2 hours
	Gradient Descent(GD), Momentum Based GD, Nesterov Accelerated GD, Stochastic GD, Adagrad, AdaDelta,RMSProp, Adam,AdaMax,NAdam, learning rate schedulers	2 hours
	Autoencoders and relation to PCA, Regularization in autoencoders, Denoising autoencoders, Sparse autoencoders, Contractive autoencoders	3 hours
	Bias Variance Tradeoff, L2 regularization, Early stopping, Dataset augmentation, Parameter sharing and tying, Injecting noise at input, Ensemble methods, Dropout	3 hours
	Greedy Layer Wise Pre-training, Better activation functions, Better weight initialization methods, Batch Normalization	3 hours
	Learning Vectorial Representations of Words, Convolutional Neural Networks, LeNet, AlexNet, ZF-Net, VGGNet, GoogLeNet, ResNet	3 hours
	Visualizing Convolutional Neural Networks, Guided Backpropagation, Deep Dream, Deep Art, Fooling Convolutional Neural Networks	3 hours
	Recurrent Neural Networks, Backpropagation Through Time (BPTT), Vanishing and Exploding Gradients, Truncated BPTT	3 hours
	Gated Recurrent Units (GRUs), Long Short Term Memory (LSTM) Cells, Solving the vanishing gradient problem with LSTM	3 hours
	Encoder Decoder Models, Attention Mechanism, Attention over images, Hierarchical Attention, Transformers.	3 hours
Pedagogy	Lectures/ Tutorials/Hands-on assignments/Self-study	
References/ Readings	1. Ian Goodfellow and Yoshua Bengio and Aaron Courville. Deep Learning. An MIT Press book. 2016. 2. Charu C. Aggarwal. Neural Networks and Deep Learning: A Textbook. Springer. 2019. 3. Dive into Deep Learning by Ashton Zang. 4. Introduction to Deep Learning by Sandro Skansi.	
Course Outcomes	1. Understanding of deep learning concepts and principles. 2. Implementation and training of deep learning models. 3. Practical application of deep learning in various domains. 4. Evaluation and interpretation of deep learning model performance.	