Deep Learning Techniques

Code: CS33212CS	Course Type: Core	Semester: II
Evolution Scheme (TA-MSE-ESE): 20-30-100	Total Marks: 150	L-T-P (Credits): 2-0-2 (3)

Pre-Requisites: Linear Algebra, Probability & Statistics, Machine Learning, Python Programming

Course Objective: To provide a comprehensive understanding of deep learning architectures and techniques, enabling students to design, train, and optimize deep neural networks for complex data-driven applications such as computer vision, natural language processing, and sequential modeling, with emphasis on performance, scalability, and generalization.

Course Outcomes: At the end of the course, the student will be able to

CO-1	Analyze and implement deep learning architectures for structured and unstructured data
CO-2	Design and optimize neural network models using appropriate training strategies and regularization techniques
CO-3	Develop advanced models for image, sequence, and language processing tasks
CO-4	Evaluate deep learning systems for performance, generalization, and deployment feasibility

Course Articulation Matrix:

PO	PEO-1	PEO-2	PEO-3	PEO-4	PEO-5
CO-1	3	2	1	-	-
CO-2	3	2	1	-	1
CO-3	3	2	-	-	3
CO-4	3	2	1	1	-

1 - Slightly; 2 - Moderately; 3 – Substantially

Syllabus:

Unit-1: Foundations of Deep Learning: Review of neural networks, perceptron and multi layer networks, activation functions, loss functions, gradient descent and backpropagation, optimization algorithms (SGD, Adam, RMSProp), weight initialization, vanishing and exploding gradient problems, regularization techniques (dropout, batch normalization), model evaluation strategies.

Unit-2: Convolutional Neural Networks: Convolution operations, pooling layers, CNN architectures (LeNet, AlexNet, VGG, ResNet), transfer learning, fine-tuning, feature extraction, object detection basics, image classification and segmentation tasks, applications in computer vision.

Unit-3: Sequence Models and Language Processing: Recurrent neural networks, LSTM and GRU architectures, sequence-to-sequence models, attention mechanisms, transformer architecture, embeddings (word2vec, GloVe), applications in NLP such as machine translation, text generation, and sentiment analysis.

Unit-4: Advanced Deep Learning Techniques: Generative models (Autoencoders, Variational Autoencoders, GANs), model compression and pruning, distributed and parallel deep learning, edge deployment considerations, ethical concerns in deep learning, robustness and adversarial attacks.

Learning Resources:

Text Books:

1. Ian Goodfellow et al., Deep Learning, MIT Press
2. François Chollet, Deep Learning with Python, Manning Dive into Deep Learning, Aston Zhang et al.