AIPC3005 DEEP LEARNING (3-0-0)

Course Objectives:

1. Understand basic concepts of neural networks and deep learning.
2. Learn to build, train, and optimize deep learning models.
3. Apply CNNs, RNNs, and Transformers to real-world tasks.
4. Develop practical skills using deep learning frameworks.
5. Understand ethical and responsible use of deep learning.

Module - I: Introduction to Deep Learning & Neural Networks (08 Hours)

Overview of machine learning vs. deep learning, Biological vs artificial neurons, Feedforward neural networks, Activation functions (ReLU, Sigmoid, Tanh, Softmax), Loss functions and performance metrics, Backpropagation and gradient computation, Practical considerations: initialization, overfitting, underfitting

Module - II: Optimization & Training Techniques (08 Hours)

Gradient descent and its variants (SGD, Momentum, RMSProp, Adam), Regularization: L1/L2, dropout, early stopping, Batch normalization and layer normalization, learning rate schedules, Hyperparameter tuning, Deep network training challenges (vanishing gradients, overfitting), Debugging deep networks (training curves, visualization)

Module - III: Convolutional Neural Networks (CNNs) (10 Hours)

Convolution operations, padding, stride, pooling, Architectures: LeNet, AlexNet, VGG, GoogLeNet, ResNet, Feature extraction with CNNs, Modern architectures: DenseNet, MobileNet, Inception v3/v4, Image classification and feature learning, Introduction to object detection (Faster R-CNN, YOLO), Introduction to semantic segmentation (FCN, U-Net), Transfer learning and fine-tuning

Module - IV: Sequence Modeling & Transformers (10 Hours)

Recurrent Neural Networks: RNN, LSTM, GRU, Sequence-to-sequence models and encoder–decoder architectures, Attention mechanism, Scaled dot-product attention and multi-head attention, Positional encoding, Transformer architecture: encoder, decoder, applications, Overview of BERT, GPT, Vision Transformers (ViT), Applications to NLP and time-series forecasting

Module - V: Generative Models & Applications (04 Hours)

Autoencoders and representation learning, Denoising autoencoders, Variational Autoencoders (VAE), Generative Adversarial Networks (GANs): Generator, discriminator, Loss functions, Challenges and improvements (DCGAN, CycleGAN), Applications in: Image synthesis, Speech processing, Text generation, Medical imaging, Agriculture & remote sensing, Ethics in deep learning: bias, fairness, responsible AI

Course Outcomes (COs)

After completing this course, students will be able to:

CO1: Understand and explain fundamental concepts of neural networks and deep learning.
CO2: Design and train deep learning models using appropriate architectures and techniques.
CO3: Apply CNNs, RNNs, and Transformers to solve classification, prediction, and sequence modeling problems.
CO4: Implement deep learning solutions using frameworks like TensorFlow/PyTorch.
CO5: Analyze model performance and optimize using regularization, tuning, and evaluation metrics.
CO6: Demonstrate awareness of ethical, responsible, and fair use of AI systems

Text Books:

1. Deep Learning with Python — Third Edition (2025) by François Chollet & Matthew Watson.
2. Neural Networks and Deep Learning: A Textbook — Second Edition (2023) by Charu C. Aggarwal.

References Books:

1. Neural Networks and Deep Learning: Theoretical Insights and Frameworks
2. Dr. Vishwas M. Mishra, Dr. Venumadhava M., B. Lalithadevi, Dr. Vasanthamma H. — published 2024 (India).
3. Deep Learning Amit Kumar Das — published by Pearson India.