Applied Machine and Deep Learning

Code: CS32113CS	Course Type: Core	Semester: I
Evolution Scheme (TA-MSE-ESE): 20-30-100	Total Marks: 150	L-T-P (Credits): 3-0-0 (3)

Pre-Requisites: Linear Algebra, Probability and Statistics, Python Programming, Data Structures, and introductory Machine Learning.

Course Objective:

1. To provide foundations of machine learning and deep learning for solving real-world data-driven problems.
2. To develop an understanding of supervised, unsupervised, and deep neural learning models.
3. To study training, optimisation, regularisation, and evaluation methods used in modern ML/DL systems.
4. To explore applications of machine learning and deep learning in vision and intelligent systems.

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

CO-1	Explain the principles, workflow, and core algorithms of machine learning and deep learning for data-driven problem solving.
CO-2	Apply supervised, unsupervised, and ensemble learning methods to analyze real-world datasets and build predictive models.
CO-3	Analyze, design, and evaluate deep neural network architectures and optimization strategies for applied learning tasks.
CO-4	Assess modern machine learning and deep learning applications in vision systems, and other intelligent data-driven domains.

Course Articulation Matrix:

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

1 - Slightly; 2 - Moderately; 3 – Substantially

Syllabus:

Unit-1: Foundations of Applied Machine Learning

Introduction to machine learning and deep learning; types of learning; data preprocessing; feature engineering; train, validation, and test sets; regression and classification; k-nearest neighbors; logistic regression; decision trees; support vector machines; model evaluation metrics; resampling methods; cross-validation; hyperparameter tuning.

Unit-2: Core Machine Learning Methods and Model Selection

Ensemble learning: bagging, boosting, random forests, gradient boosting; regularization; model selection; dimensionality reduction using PCA and manifold learning; clustering using k-means, DBSCAN, and hierarchical clustering; and performance measures.

Unit-3: Deep Learning Foundations and Architectures

Artificial neural networks; multilayer perceptrons; activation functions; loss functions; backpropagation; gradient descent and optimization; best practices for training neural networks; regularization and normalization; convolutional neural networks; recurrent neural networks; LSTM

Unit-4: Applied Deep Learning and Emerging Applications

Transfer learning; self-supervised learning; domain adaptation, deep learning applications in computer vision, natural language processing, audio, medical imaging; responsible AI considerations; recent trends in modern ML/DL systems.

Learning Resources:

1. Trevor Hastie, Robert Tibshirani, Jerome Friedman, The Elements of Statistical Learning, Springer.
2. Kevin P. Murphy, Machine Learning: A Probabilistic Perspective, MIT Press.
3. Aurélien Géron, Hands-On Machine Learning with Scikit-Learn, Keras, and TensorFlow, O'Reilly.
4. Charu C. Aggarwal, Neural Networks and Deep Learning, Springer.

Text Books:

1. Ian Goodfellow, Yoshua Bengio, and Aaron Courville, Deep Learning , Adaptive Computation and Machine Learning Series, First Edition, MIT Press, 2016, ISBN: 978-0262035613
2. Christopher M. Bishop, Pattern Recognition and Machine Learning , Information Science and Statistics, First Edition, Springer, 2006, ISBN: 978-0387310732
3. Ethem Alpaydin, Introduction to Machine Learning, Fourth Edition, MIT Press, 2020, ISBN: 978-0262043793
4. Christopher M. Bishop and Hugh Bishop, Deep Learning: Foundations and Concepts, First Edition, Springer, 2024, ISBN: 978-3031454677
1. François Chollet and Matthew Watson, Deep Learning with Python, Third Edition, Manning Publications, 2025, ISBN: 978-1633438262