Applied AI and Machine Learning

Code: CS33121CS	Course Type: Elective	Semester: I
Evolution Scheme (TA-MSE-ESE): 20-30-100	Total Marks: 150	L-T-P (Credits): 3-0-0 (3)

Pre-Requisites: Students should have basic knowledge of Programming in Python or any high-level language, Linear algebra and matrices, Probability and statistics, Data structures and basic algorithms, Fundamentals of computer systems and databases

Course Objective:

1. Provide the fundamental concepts of Artificial Intelligence, Machine Learning, and related mathematical foundations.
2. Develop the ability to apply supervised and unsupervised learning techniques to real world problems.
3. Enable students to design, build, and evaluate basic machine learning and neural network models.
4. Create awareness of ethical issues, interpretability, and practical deployment aspects of AI systems.

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

CO-1	Explain the fundamentals of Artificial Intelligence, Machine Learning, and their mathematical foundations.
CO-2	Apply supervised and unsupervised learning techniques to analyze data and solve real-world problems.
CO-3	Develop and evaluate basic machine learning and neural network models using appropriate tools and frameworks.
CO-4	Analyze ethical, societal, and deployment-related issues in AI systems and communicate AI-based solutions effectively.

Course Articulation Matrix:

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

1 - Slightly; 2 - Moderately; 3 – Substantially

CO–PO Mapping:

CO / PO	PO-1	PO-2	PO-3	PO-4	PO-5	PO-6
CO-1	2	1	1	1	1	-
CO-2	3	3	1	2	1	1
CO-3	3	2	2	3	1	1
CO-4	1	1	1	2	3	3

Syllabus:

Unit-1: Foundations of Artificial Intelligence and Machine Learning

Introduction to Artificial Intelligence, Machine Learning, and Deep Learning. Types of learning: supervised, unsupervised, and reinforcement learning. Applications of AI and ML in real-world domains. Mathematical foundations: vectors, matrices, probability, statistics, and optimization basics. Python environment for AI/ML: NumPy, pandas, matplotlib, scikit learn basics.

Unit-2: Supervised Learning

Regression and classification concepts. Linear regression and logistic regression. Decision trees, random forest, k-nearest neighbors, Naive Bayes, and support vector machines. Model training and testing. Overfitting, underfitting, bias-variance tradeoff. Performance evaluation metrics: accuracy, precision, recall, F1-score, confusion matrix, RMSE.

Unit-3: Unsupervised Learning and Neural Networks

Clustering: k-means and hierarchical clustering. Dimensionality reduction using PCA. Anomaly detection basics. Introduction to artificial neural networks. Perceptron, multilayer perceptron, activation functions, backpropagation. Overview of deep learning, CNN, and RNN with simple applications.

Unit-4: Applied AI, Ethics, and Deployment

Data preprocessing and feature engineering. Model selection and interpretation. Applications of AI in healthcare, finance, agriculture, education, and industry. Ethical issues in AI: bias, fairness, privacy, and responsible AI practices. Basics of model deployment and inference pipeline. Mini case study / practical application of AI and ML.

Learning Resources:

1. Aurélien Géron, Hands-On Machine Learning with Scikit-Learn, Keras, and TensorFlow, O'Reilly.
2. Sebastian Raschka and Vahid Mirjalili, Python Machine Learning, Packt.
3. Ian Goodfellow, Yoshua Bengio, Aaron Courville, Deep Learning, MIT Press.
4. Kevin P. Murphy, Machine Learning: A Probabilistic Perspective, MIT Press.

Text Books:

1. Ethem Alpaydin, Introduction to Machine Learning, MIT Press.
2. Tom M. Mitchell, Machine Learning, McGraw-Hill.
3. Christopher M. Bishop, Pattern Recognition and Machine Learning, Springer.
4. Trevor Hastie, Robert Tibshirani, Jerome Friedman, The Elements of Statistical Learning, Springer.