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.

Explainable AI

Code: CS33122CS	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: Machine Learning, Deep Learning, Probability & Statistics

Course Objective: To develop an understanding of interpretability and transparency in AI systems, enabling students to design explainable models, evaluate fairness and bias, and ensure accountability in AI-driven decision-making systems.

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

CO-1	Analyze the need for interpretability and transparency in AI models across application domains
CO-2	Apply model-agnostic and model-specific explainability techniques to interpret machine learning and deep learning models
CO-3	Evaluate fairness, bias, and ethical implications of AI systems
CO-4	Design trustworthy AI systems incorporating explainability, accountability, and compliance principles

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: Introduction to Explainable AI: Motivation for interpretability, black-box vs white box models, levels of explainability, human-centered AI, regulatory requirements, trust and accountability in AI systems, evaluation of interpretability.

Unit-2: Model-Agnostic Explainability Techniques: Feature importance methods, partial dependence plots, surrogate models, LIME, SHAP, sensitivity analysis, local vs global explanations, visualization techniques for interpretability.

Unit-3: Model-Specific Interpretability: Interpretable models (decision trees, rule-based systems), explainability in deep learning (saliency maps, Grad-CAM, attention visualization), interpretability in NLP and vision models, counterfactual explanations.

Unit-4: Fairness, Ethics and Trustworthy AI: Bias detection and mitigation techniques, fairness metrics, accountability and governance frameworks, privacy-preserving explainability, explainability in high-stakes domains (healthcare, finance), case studies and real-world deployments.

Learning Resources:

Text Books:

1. Christoph Molnar, Interpretable Machine Learning
2. Samek, Montavon et al., Explainable AI: Interpreting, Explaining and Visualizing Deep Learning
3. Ribeiro et al. (LIME), Lundberg & Lee (SHAP) – research papers

Optimization Techniques

Code: CS33123CS	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: Programming Knowledge, Design and Analysis of Algorithm, Discrete Mathematics, Engineering Mathematics / Numerical Methods, Artificial Intelligence & Data Science

Course Objective: To equip students with fundamental and advanced **optimization techniques** for modeling and solving engineering problems, and to develop the ability to apply mathematical and computational methods for efficient decision-making.

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

CO-1	To develop a strong foundation in convex optimization and duality theory.
CO-2	To understand and analyse modern optimization algorithms.
CO-3	To apply optimization methods to machine learning, data science, and engineering problems.
CO-4	To explore both convex and non-convex optimization techniques.

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 Optimization

Background Concepts: Convex sets and convex functions, Convex analysis fundamentals, Linear algebra and matrix theory (norms, eigenvalues, PSD matrices), Basics of probability theory for optimization

Preliminaries: Applications in engineering, AI, and data science, Formulation of optimization problems, Optimality conditions (first-order and second-order), Duality theory: Lagrangian, weak and strong duality, KKT conditions

Unit-2: First-Order Optimization Methods:

Gradient-Based Methods: Gradient descent and its variants, Convergence analysis Sub-gradient and

Proximal Methods: Sub-gradient methods for non-differentiable optimization, Proximal gradient methods, Regularization and composite optimization Advanced First-Order Techniques: Mirror descent, Multiplicative weights update method, Accelerated first-order methods (e.g., Nesterov acceleration)

Unit-3: Second-Order and Stochastic Optimization

Second-Order Methods: Newton's method, Quasi-Newton methods (BFGS, L-BFGS), Convergence and computational trade-offs

Stochastic Optimization: Stochastic optimization problem formulation, Notion of regret and performance measures, Online-to-batch conversion

Online Learning Methods: Online Gradient Descent (OGD), Exponentiated Gradient (EG), Online Mirror Descent (OMD), Methods with vanishing regret.

Unit-4: Non-Convex Optimization and Advanced Methods:

Non-Convex Optimization: Nature and challenges of non-convex problems, Multiple local minima and saddle points, Lack of global optimality guarantees. Applications: Sparse recovery, Affine rank minimization, Low-rank matrix completion. Solution Approaches: Convex relaxation methods. Non-convex methods (Projected gradient descent, alternating minimization). Interior Point Methods (for large-scale convex problems), Coordinate Descent Methods, Optimization for Deep Learning (SGD and mini-batch methods), Parallel and Distributed Optimization (basic concepts), Variance-Reduced Methods (SVRG, SAGA)

Learning Resources:

Text Books:

1. Dimitri P. Bertsekas, Nonlinear Programming, 3rd Edition, Athena Scientific, 2016
2. S. Boyd and L. Vandenberghe, Convex Optimization, The Cambridge University Press, 2003.
3. S. Bubeck, Convex Optimization: Algorithms and Complexity, Foundations and Trends in Machine Learning, 8(3-4): 231-357, 2015.
4. S. Sra, S. Nowozin, and S. Wright (eds). Optimization for Machine Learning, The MIT Press, 2011.

References:

1. T. Hastie, R. Tibshirani and M. J. Wainwright, Statistical Learning with Sparsity: the Lasso and Generalizations, Chapman and Hall/CRC Press, 2015. http://web.stanford.edu/~hastie/StatLearnSparsity/
2. E. Hazan. Introduction to Online Convex Optimization. Draft, 2015. http://ocobook.cs.princeton.edu/ 3. Y. Nesterov, Introductory lectures on convex optimization, Kluwer-Academic, 2003.