*CSXX02*48: Advanced Topics in Machine Learning

L-T-P-Cr: 3-0-0-3

Pre-requisites: Machine Learning, Deep Learning, Reinforcement Learning.

Objective: To expose students to recent and advanced topics in machine learning that go beyond traditional supervised/unsupervised techniques, focusing on interpretability, fairness, causality, meta-learning, graph learning, and scalable ML systems.

Course Outcomes:

CO1: Understand and apply advanced supervised learning techniques such as ensemble models, cost-sensitive learning, and multi-label classification.
CO2: Analyse and evaluate machine learning models for interpretability, fairness, and causal relationships.
CO3: Demonstrate the ability to apply meta-learning, AutoML, and few-shot learning techniques to adapt models for low-data and dynamic environments.
CO4: Design and implement machine learning solutions for graph-structured and sequential data using modern graph-based models.
CO5: Explore scalable and online learning approaches suitable for large-scale, streaming, or decentralized data settings.

CO-PO Mapping

CO \ PO	PO1	PO2	PO3	PO4	PO5	PO6	PO7	PO8	PO9	PO10	PO11	PO12
CO1	3	3	2	2	3	-	-	-	-	1	-	2
CO2	3	3	2	3	3	2	1	2	-	2	-	2
CO3	3	2	2	2	3	-	-	-	-	1	1	3
CO4	3	2	3	2	3	-	-	-	1	2	1	2
CO5	3	2	2	2	3	1	1	-	1	2	2	3

Correlation Levels:

- 3: Strong Correlation
- 2: Moderate Correlation
- 1: Low Correlation
- -: No Correlation

UNIT I: Advanced Supervised Learning Techniques: Lectures: 8

Ensemble Learning: Gradient Boosting (XGBoost, LightGBM), Stacking. Multi-label and Multi-output Learning. Metric Learning. Handling Imbalanced Data: SMOTE, Focal Loss. Cost-sensitive Learning. Calibration of Probabilistic Models (e.g., Platt Scaling, Isotonic Regression).

UNIT II: Interpretability, Fairness, and Causality in ML: Lectures: 10

Explainable AI (XAI): LIME, SHAP, Integrated Gradients. Fairness in Machine Learning: Types of Bias, Fairness Metrics. Causal Inference: Potential Outcomes Framework, DAGs and Causal Graphs,
Do-Calculus (Pearl's framework), Applications in Treatment Effect Estimation.

UNIT III: Meta Learning, AutoML, and Few-shot Learning: Lectures: 8

Meta-Learning Approaches: Model-Agnostic Meta-Learning (MAML), Reptile. Optimization-Based vs Metric-Based Meta Learning. AutoML: Neural Architecture Search (NAS), Hyperparameter Optimization. Few-shot and Zero-shot Learning. Transfer Learning vs Meta Learning.

UNIT IV: Graph-Based and Structured Data Learning: Lectures: 8

Basics of Graph Representation Learning. Graph Neural Networks (GNNs), Graph Convolutional Networks (GCNs). Applications: Social Network Analysis, Recommender Systems, Chemistry.
Structured Prediction (e.g., CRFs, Structured SVMs). Sequence Labelling and Dependency Parsing (overview)

UNIT V: Scalable and Online Machine Learning: Lectures: 8

Online Learning Algorithms: Perceptron, Passive-Aggressive. Federated Learning: Architecture, Challenges, Privacy. Distributed ML Systems: Parameter Servers, Horovod. Streaming Data Learning and Concept Drift. ML at Scale: ML Pipelines, Deployment Tools (brief overview of TensorFlow Extended, MLflow, etc.)

Suggested Texts and References:

1. **Machine Learning: A Probabilistic Perspective – Kevin P. Murphy**
2. **Interpretable Machine Learning – Christoph Molnar**
3. **Causal Inference in Statistics: A Primer – Judea Pearl**

Research papers and tutorials from NeurIPS, ICML, ICLR proceedings.