III B.Tech II Semester

23A39601	**ADVANCED MACHINE LEARNING (Professional Core)**	**L**	**T**	**P**	**C**
		3	**0**	**0**	**3**

Course Objectives

- Provide an in-depth understanding of advanced topics in machine learning beyond traditional models.
- Explore ensemble learning, kernel methods, structured prediction, and deep generative models.
- Equip students with the ability to evaluate and implement state-of-the-art machine learning algorithms for real-world problems.
- Introduce recent advancements in model generalization, regularization techniques, and probabilistic reasoning.
- Foster analytical thinking to solve complex AI problems through advanced modeling and optimization techniques.

Course Outcomes

- Demonstrate a clear understanding of advanced machine learning algorithms and their theoretical underpinnings.
- Apply ensemble methods, kernel tricks, and probabilistic models to diverse data science problems.
- Analyze and implement optimization techniques for large-scale machine learning problems.
- Design and evaluate structured models and generative approaches for complex datasets.
- Develop and test machine learning systems using modern frameworks and evaluate performance using robust metrics.

UNIT I – Ensemble Learning and Model Generalization

Bias-Variance Trade-off Revisited, Bagging and Bootstrap Aggregating, Random Forests – Theory and Implementation, Boosting Techniques – AdaBoost, Gradient Boosting, XGBoost and LightGBM, Stacking and Blending of Models, Regularization Methods – L1, L2, Dropout, Early Stopping and Cross-Validation Strategies.

UNIT II – Kernel Methods and Support Vector Machines

Linear and Non-linear Classification, Kernel Trick – Polynomial, RBF, and Custom Kernels, Soft Margin SVMs, Dual Form and Optimization of SVM, Support Vector Regression, Kernel PCA for Non-linear Dimensionality Reduction, Practical Issues with Kernel Methods, Applications in Text and Image Classification.

UNIT III – Probabilistic Graphical Models and Bayesian Learning

Introduction to Probabilistic Graphical Models, Bayesian Networks – Construction and Inference, Markov Random Fields (MRFs), Conditional Random Fields (CRFs), Expectation-Maximization (EM) Algorithm, Variational Inference, Bayesian Linear Regression, Gaussian Processes for Regression and Classification.

UNIT IV – Structured Prediction and Unsupervised Learning
Hidden Markov Models (HMMs) and Sequence Modeling, Structured SVMs and CRFs, Clustering Revisited – Hierarchical and Spectral Methods, Dimensionality Reduction – ICA, t-SNE, UMAP, Matrix Factorization and Collaborative Filtering, Autoencoders and Variational Autoencoders (VAEs), Clustering Evaluation Metrics, Advanced Use Cases in NLP and Computer Vision.

UNIT V – Optimization and Deep Generative Models
Convex vs. Non-convex Optimization, Gradient Descent Variants – SGD, Adam, RMSProp, Deep Generative Models – VAEs and GANs, Conditional GANs and StyleGANs, Reinforcement Learning Introduction, Policy Gradient Methods, Generative Pre-trained Transformers (GPT) Overview, Advanced Topics – Meta Learning, Few-shot Learning.

Textbooks

1. "Pattern Recognition and Machine Learning" by Christopher M. Bishop
2. "Machine Learning: A Probabilistic Perspective" by Kevin P. Murphy
3. "Deep Learning" by Ian Goodfellow, Yoshua Bengio, and Aaron Courville

Reference Books

1. "Understanding Machine Learning: From Theory to Algorithms" by Shai Shalev-Shwartz and Shai Ben-David
2. "The Elements of Statistical Learning" by Trevor Hastie, Robert Tibshirani, and Jerome Friedman
3. "Bayesian Reasoning and Machine Learning" by David Barber
4. Recent IEEE Transactions and ACM journals on ML

Online Courses

1. Advanced Machine Learning Specialization – Coursera (HSE University)
2. Probabilistic Graphical Models – Stanford (Coursera)
3. Advanced Deep Learning – DeepLearning.AI (Coursera)
4. Bayesian Methods for Machine Learning – Coursera
5. Advanced Machine Learning – NPTEL