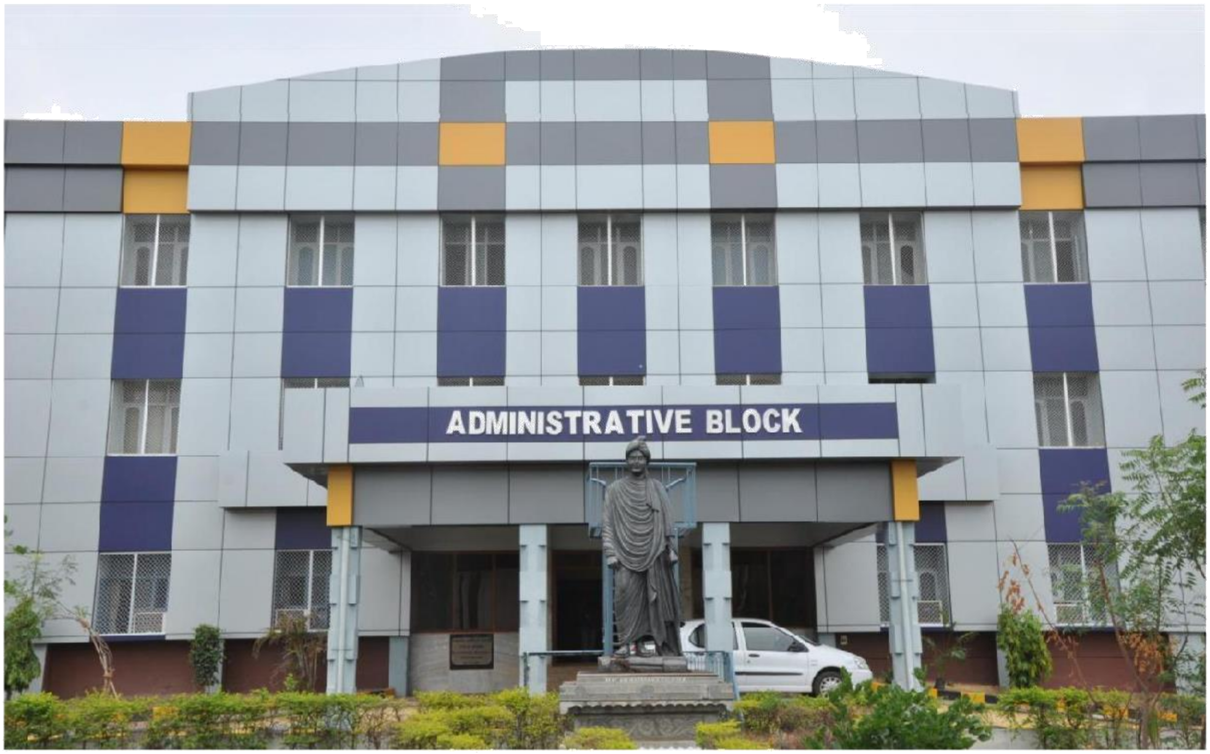

Scheme–2023

Department of Emerging Technologies in Computer Science

G.Pulla Reddy Engineering College (Autonomous): Kurnool

Accredited by NBA of AICTE and NAAC of UGC

Affiliated to JNTUA, Anantapuramu

Scheme and Syllabus for

Honors in

Computer Science and Engineering (AI&ML)

(With Effect from the Batch Admitted in 2023-24)

G. PULLA REDDY ENGINEERING COLLEGE (Autonomous) : KURNOOL
SCHEME -23
B. TECH – CSE (AI&ML) Honors
Applicable from the Academic Year 2023-24 onwards

B.Tech- CSE(AI & ML) Honors

S.No	Title	L	T	P	Credits	CIA	End Exam	Total Marks
1.	Advanced Algorithms for AI & ML/MOOCs	3	0	0	3	30	70	100
2.	Deep Learning & Neural Network Architectures / MOOCs	3	0	0	3	30	70	100
3.	Reinforcement Learning & Decision Making	3	0	0	3	30	70	100
4.	AI for Robotics & Automation	3	0	0	3	30	70	100
5.	AI Ethics, Fairness & Explainability	3	0	0	3	30	70	100
6.	Advanced Algorithms for AI & ML Lab	0	0	3	1.5	30	70	100
7.	Robotics & Autonomous Systems Lab	0	0	3	1.5	30	70	100
				Total	**18**	**-**	**-**	**-**

ADVANCED ALGORITHMS FOR AI & ML(AAML)								
Honors in AI&ML					Scheme: 2023			
Course Code	Category	Hours/Week			Credits	Maximum Marks		
HCM01	H	L	T	P	C	Continuous Internal Assessment	End Exam	TOTAL
		3	0	0	3	30	70	100
Sessional Exam Duration: 2 Hours					End Exam Duration: 3 Hours			

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

CO1:	Analyze and apply classical algorithmic techniques including divide and conquer, dynamic programming, approximation, and randomized algorithms in the context of AI/ML.
CO2:	Implement advanced graph algorithms for shortest paths, flows, and community detection, and apply them to AI problems like NLP and recommender systems.
CO3:	Apply convex and non-convex optimization strategies, gradient-based learning, and regularization techniques to train and tune AI/ML models effectively.
CO4:	Use evolutionary, swarm intelligence, and reinforcement learning-based metaheuristic methods for neural architecture search and complex optimization tasks in AI.
CO5:	Evaluate and design scalable algorithmic solutions with fairness and interpretability for AI/ML applications, referencing case studies like AlphaGo, GPT, and Auto ML systems.

UNIT– I

Foundations of Advanced Algorithmic Techniques: Review of Time and Space Complexity, Divide and Conquer, Dynamic Programming, and Greedy Algorithms, Recurrence Relations and Master Theorem, Approximation Algorithms: Vertex Cover, TSP, Set Cover, Randomized Algorithms: Monte Carlo and Las Vegas Types, Applications in ML Preprocessing and Feature Selection.

UNIT– II

Graph Algorithms and AI Applications: Graph Representations and Traversal Algorithms, Shortest Path: Dijkstra's, Bellman-Ford, Floyd Warshall, Minimum, Spanning Trees: Kruskal and Prim, Network Flows and Max Flow-Min Cut Theorem, Graph-Based Semi-Supervised Learning, PageRank, Centrality, and Community Detection, Applications in NLP,Vision and Recommender systems

UNIT– III

Optimization in AI/ML: Convex and Non-Convex Optimization, Gradient Descent Variants: SGD, Momentum, Adam, Convergence Analysis and Learning Rates, Duality and Lagrange Multipliers, Regularization: L1, L2, ElasticNet, Hyperparameter Optimization: Grid, Random, Bayesian, Constrained Optimization in SVMs and Deep Learning.

UNIT– IV

Evolutionary & Metaheuristic Algorithms: Genetic Algorithms and Evolutionary Strategies, Swarm Intelligence: PSO, Ant Colony Optimization, Simulated Annealing and Tabu Search, Multi-objective Optimization, Reinforcement Learning and Policy Gradient Methods, Neuroevolution: Evolving Neural Networks, Use Cases in Feature Engineering and Neural Architecture Search (NAS).

UNIT– V

Advanced Topics and Case Studies: Online Learning and Regret Minimization, Bandit Algorithms: Multi-Armed Bandits, Thompson Sampling, Large-Scale Algorithms: MapReduce, Apache Spark MLlib, Case Studies: AlphaGo, GPT, BERT, Recommendation Engines, Research Trends in Algorithmic ML and AutoML, Capstone Problem Solving using Hybrid Algorithms.

Text Books:

1. Introduction to Algorithms – Thomas H. Cormen, Charles E. Leiserson, Ronald L. Rivest, Clifford Stein (MIT Press).

2. Algorithms for Machine Learning – Giuseppe Bonaccorso, Packt Publishing.
3. Convex Optimization – Stephen Boyd and Lieven Vandenberghe, Cambridge University Press.
4. Reinforcement Learning: An Introduction – Richard S. Sutton and Andrew G. Barto.

Reference Books:

1. Machine Learning: A Probabilistic Perspective – Kevin P. Murphy.
2. The Elements of Statistical Learning – Trevor Hastie, Robert Tibshirani, Jerome Friedman.
3. Evolutionary Computation – Kenneth A. De Jong
4. Handbook of Approximation Algorithms and Metaheuristics – Teofilo F. Gonzalez

Web References:

1. Advanced Algorithms and Complexity | Coursera

Question Paper Pattern:

Sessional Examination:

The question paper for Sessional Examination shall be for 40 marks. The question paper shall consist of four questions and all questions are compulsory. Question No.1 contains five short answer questions (2 marks each) for a total of ten marks. Remaining three questions shall be EITHER/OR type descriptive questions for ten marks each. Each of these descriptive questions may contain sub-questions.

End Examination:

The question paper for End Examination shall be for 70 marks. The question paper shall contain six questions and all questions are compulsory. Question No.1 shall contain ten short answer questions (2 marks each) for a total of twenty marks, with two short answer questions from each unit. Remaining five questions (Each question covering one unit of syllabus) carrying 10 marks each shall be EITHER/OR type descriptive questions and may contain sub-questions.

DEEP LEARNING & NEURAL NETWORK ARCHITECTURES(DL&NA)								
Honors in AI&ML					Scheme: 2023			
Course Code	Category	Hours/Week			Credits	Maximum Marks		
HCM02	H	L	T	P	C	Continuous Internal Assessment	End Exam	TOTAL
		3	0	0	3	30	70	100
Sessional Exam Duration: 2 Hours					End Exam Duration: 3 Hours			

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

CO1:	Understand the theoretical foundations of neural networks and deep learning.
CO2:	Implement and train multilayer perceptrons, CNNs, RNNs, and other architectures.
CO3:	Analyze and optimize deep learning models using advanced regularization and tuning techniques.
CO4:	Evaluate the applicability of different neural network architectures for various AI problems.
CO5:	Apply state-of-the-art models such as Transformers and GANs in real-world domains.

UNIT– I

Foundations of Neural Networks: Introduction to Artificial Neural Networks, Biological Neuron vs. Artificial Neuron, Perceptron, Multilayer Perceptron (MLP), Activation Functions: ReLU, Sigmoid, Tanh, Softmax, Backpropagation and Gradient Descent, Loss Functions: MSE, Cross Entropy, Overfitting, Regularization (L1/L2), Dropout

UNIT– II

Convolutional Neural Networks (CNNs): Convolution Operation and Feature Maps, Pooling Layers: Max and Average Pooling, CNN Architectures: LeNet, AlexNet, VGG, ResNet, Transfer Learning and Fine-tuning, Image Classification, Object Detection Basics, Implementation with TensorFlow/PyTorch

UNIT– III

Recurrent Neural Networks (RNNs) and Variants: Sequential Data and Time Series, RNN Basics and Back propagation Through Time (BPTT), Vanishing and Exploding Gradients, LSTM and GRU Architectures, Applications in Text, Speech, and Music, Sequence-to-Sequence Models

UNIT– IV

Advanced Architectures & Optimization: Autoencoders and Variational Autoencoders (VAEs), Generative Adversarial Networks (GANs), Deep Reinforcement Learning Overview, Batch Normalization, Early Stopping, Hyperparameter Tuning and Optimization, Performance Metrics and Evaluation

UNIT– V

Transformer Models & Applications: Attention Mechanism and Self-Attention, Transformers and BERT Architecture, Positional Encoding, Multi-head Attention, Pre-trained Language Models and Fine-Tuning, Applications in NLP: Text Classification, Translation, Large Language Models and Transfer Learning

Text Books:

1. Deep Learning – Ian Goodfellow, Yoshua Bengio, and Aaron Courville (MIT Press)
2. Neural Networks and Deep Learning – Michael Nielsen (Online Book)
3. Hands-On Machine Learning with Scikit-Learn, Keras & TensorFlow – Aurélien Géron (O'Reilly)

Reference Books:

1. Pattern Recognition and Machine Learning – Christopher M. Bishop
2. Deep Learning for Computer Vision – Rajalingappaa Shanmugamani
3. Natural Language Processing with Transformers – Lewis Tunstall, Leandro von Werra, Thomas Wolf
4. Reinforcement Learning: An Introduction – Richard S. Sutton and Andrew G. Barto

Web References:

1. Deep Learning Specialization on Coursera
2. CS231n – Stanford Course Website
3. Practical Deep Learning for Coders – Fast.ai
4. Intro to Deep Learning with PyTorch – Udacity
5. Transformers Course – Hugging Face Learn

Question Paper Pattern:

Sessional Examination:

The question paper for Sessional Examination shall be for 40 marks. The question paper shall consist of four questions and all questions are compulsory. Question No.1 contains five short answer questions (2 marks each) for a total of ten marks. Remaining three questions shall be EITHER/OR type descriptive questions for ten marks each. Each of these descriptive questions may contain sub-questions.

End Examination:

The question paper for End Examination shall be for 70 marks. The question paper shall contain six questions and all questions are compulsory. Question No.1 shall contain ten short answer questions (2 marks each) for a total of twenty marks, with two short answer questions from each unit. Remaining five questions (Each question covering one unit of syllabus) carrying 10 marks each shall be EITHER/OR type descriptive questions and may contain sub-questions.

REINFORCEMENT LEARNING & DECISION MAKING(RLDM)

Honors in AI&ML								Scheme: 2023
Course Code	Category	Hours/Week			Credits	Maximum Marks		
HCM03	H	L	T	P	C	Continuous Internal Assessment	End Exam	TOTAL
		3	0	0	3	30	70	100
Sessional Exam Duration: 2 Hours						End Exam Duration: 3 Hours		

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

CO1:	Understand the fundamentals of reinforcement learning, including agent environment interaction, types of RL, and solving decision-making problems using Markov Decision Processes and Bellman equations.
CO2:	Apply dynamic programming and Monte Carlo methods to perform policy evaluation, policy improvement, and control in model-based RL settings.
CO3:	Implement temporal-difference learning algorithms like TD(0), Sarsa, and Q-learning and extend them using eligibility traces and function approximation techniques.
CO4:	Develop and analyze policy gradient and actor-critic methods, including REINFORCE and PPO, to optimize policies in continuous and high-dimensional action spaces.
CO5:	Employ deep reinforcement learning techniques (DQN, DDPG, A3C, SAC) and exploration strategies to solve complex tasks in robotics, games, and autonomous systems, considering safety and ethical decision-making.

UNIT– I

Introduction to Reinforcement Learning & MDPs Foundations of RL: Agent-Environment Interaction, Types of RL: Model-based vs. Model-free, Reward Signals, Return, and Discounting, Markov Decision Processes (MDPs), Bellman Equations and Optimality.

UNIT– II

Dynamic Programming & Monte Carlo Methods: Policy Evaluation and Policy Improvement, Value Iteration and Policy Iteration, Monte Carlo Prediction and Control, First-visit and Every-visit Methods, Limitations of DP and MC Approaches.

UNIT– III

Temporal-Difference Learning & Function Approximation: TD(0), Sarsa, and Q-Learning Algorithms, Eligibility Traces: TD(λ), Sarsa(λ), Off-policy vs. On-policy Learning, Linear Function Approximation, Generalization in RL.

UNIT– IV

Policy Gradient Methods and Actor-Critic Algorithms: Policy Gradient Theorem, REINFORCE Algorithm, Baselines and Variance Reduction, Actor-Critic Architectures, Trust Region and Proximal Policy Optimization (PPO).

UNIT– V

Deep Reinforcement Learning and Applications: Deep Q-Networks (DQN) and Experience Replay, DDPG, A3C, and SAC Algorithms, Exploration Techniques: ε-greedy, UCB, Intrinsic Rewards, RL in Robotics, Game AI, and Autonomous Systems, Safety, Ethics, and Fairness in Decision Making.

Text Books:

1. Richard S. Sutton and Andrew G. Barto – Reinforcement Learning: An Introduction, 2nd Edition, MIT Press
2. Ian Goodfellow, Yoshua Bengio, Aaron Courville – Deep Learning, MIT Press

Reference Books:

1. David Silver's RL Course Slides & Lectures – DeepMind, University College London
2. Marco Wiering & Martijn van Otterlo (Eds.) – Reinforcement Learning: State of the Art, Springer
3. Csaba Szepesvári – Algorithms for Reinforcement Learning, Morgan & Claypool
4. Yuxi Li – Deep Reinforcement Learning: An Overview, arXiv survey

Web References:

1. DeepMind x UCL Reinforcement Learning Lectures by David Silver
2. Coursera: Reinforcement Learning Specialization – University of Alberta : https://www.coursera.org/specializations/reinforcement-learning

Question Paper Pattern:

Sessional Examination:

The question paper for Sessional Examination shall be for 40 marks. The question paper shall consist of four questions and all questions are compulsory. Question No.1 contains five short answer questions (2 marks each) for a total of ten marks. Remaining three questions shall be EITHER/OR type descriptive questions for ten marks each. Each of these descriptive questions may contain sub-questions.

End Examination:

The question paper for End Examination shall be for 70 marks. The question paper shall contain six questions and all questions are compulsory. Question No.1 shall contain ten short answer questions (2 marks each) for a total of twenty marks, with two short answer questions from each unit. Remaining five questions (Each question covering one unit of syllabus) carrying 10 marks each shall be EITHER/OR type descriptive questions and may contain sub-questions.

AI FOR ROBOTICS & AUTOMATION(AIR&A)

Honors in AI&ML						Scheme: 2023		
Course Code	Category	Hours/Week			Credits	Maximum Marks		
HCM04	H	L	T	P	C	Continuous Internal Assessment	End Exam	TOTAL
		3	0	0	3	30	70	100
Sessional Exam Duration: 2 Hours						End Exam Duration: 3 Hours		

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

CO1:	Describe the role of AI in modern software engineering processes and lifecycle stages.
CO2:	Apply AI/ML models for requirements gathering, code generation, defect prediction, and testing.
CO3:	Implement intelligent DevOps practices including CI/CD, release automation, and anomaly detection.
CO4:	Analyze data from software pipelines to drive informed decisions and improve quality.
CO5:	Develop an end-to-end AI-enabled software delivery pipeline with automated learning based optimizations.

UNIT– I

Foundations of AI-Driven Software Engineering: Introduction to Software Engineering Lifecycle, Traditional vs. AI-Driven Software Development, AI/ML in Software Engineering: Overview and Scope, Natural Language Processing (NLP) for Requirements Engineering, AI for Software Design Recommendation, Intelligent Code Completion (e.g., GitHub Copilot).

UNIT– II

AI in Testing and Defect Prediction: Static and Dynamic Testing with AI, Automated Test Case Generation, Defect Detection and Prediction using ML Models, Sentiment and Bug Report Analysis, AI in Refactoring and Code Review, Tools: SonarQube, DeepCode.

UNIT– III

DevOps Principles and Practices: DevOps Overview: CI/CD Pipelines, Infrastructure as Code (IaC), Configuration Management Tools: Ansible, Puppet,Monitoring and Logging Tools: Prometheus, Grafana, Containerization and Orchestration: Docker, Kubernetes, Agile and Lean Practices in DevOps.

UNIT– IV

AI for DevOps Automation and Intelligence: Predictive Analytics for Deployment Success, AI for Log Analytics and Root Cause Analysis, Self-Healing Systems and Auto-Scaling, Feedback Loops in DevOps using Reinforcement Learning, Data-Driven Decision Making in Release Management, ChatOps and AIOps Platforms.

UNIT– V

Case Studies and Emerging Trends : Case Study: AI-Augmented DevOps in Enterprises, ML-Ops vs. DevOps vs. DataOps, Security in DevOps (DevSecOps), Explainability and Ethics in AI-Driven Software Engineering, Generative AI in Software Development, Future Trends and Industry Standards.

Text Books:

1. Tim Menzies, Diomidis Spinellis – Artificial Intelligence and Software Engineering: Status and Future Directions
2. Len Bass, Ingo Weber, Liming Zhu – DevOps: A Software Architect's Perspective, Addison-Wesley
3. Thomas Erl, Ricardo Puttini, Zaigham Mahmood – AI & Analytics for DevOps, Pearson

Reference Books:

1. Carlos Nunes Silva – AI in Software Engineering
2. Gene Kim, Jez Humble, Patrick Debois, John Willis – The DevOps Handbook
3. Andrew Ng – Machine Learning B.Techning (AI Systems Engineering Perspective)

Web References:

1. Coursera – AI for Software Engineering (IBM)
2. DevOps with Microsoft Azure – edX
3. Udacity – AI for DevOps Engineers Nanodegree

Question Paper Pattern:

Sessional Examination:

The question paper for Sessional Examination shall be for 40 marks. The question paper shall consist of four questions and all questions are compulsory. Question No.1 contains five short answer questions (2 marks each) for a total of ten marks. Remaining three questions shall be EITHER/OR type descriptive questions for ten marks each. Each of these descriptive questions may contain sub-questions.

End Examination:

The question paper for End Examination shall be for 70 marks. The question paper shall contain six questions and all questions are compulsory. Question No.1 shall contain ten short answer questions (2 marks each) for a total of twenty marks, with two short answer questions from each unit. Remaining five questions (Each question covering one unit of syllabus) carrying 10 marks each shall be EITHER/OR type descriptive questions and may contain sub-questions.

AI ETHICS, FAIRNESS & EXPLAINABILITY(AIEFE)								
Honors in AI&ML					Scheme: 2023			
Course Code	Category	Hours/Week			Credits	Maximum Marks		
HCM05	H	L	T	P	C	Continuous Internal Assessment	End Exam	TOTAL
		3	0	0	3	30	70	100
Sessional Exam Duration: 2 Hours					End Exam Duration: 3 Hours			

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

CO1:	Describe ethical principles, history, and responsibilities in AI use across domains like healthcare and law enforcement.
CO2:	Identify different forms of bias in datasets and algorithms, and apply fairness metrics and mitigation strategies to ensure equitable AI systems.
CO3:	Explain need for AI model interpretability and use LIME, SHAP, Grad-CAM for local & global explanations.
CO4:	Apply and Design AI systems with accountability by integrating human oversight.
CO5:	Critically assess the broader societal and legal implications of AI.

UNIT– I

Foundations of AI Ethics : Historical background of AI ethics, Core principles: beneficence, non-maleficence, autonomy, justice, Moral and legal responsibilities in AI systems, Risk assessment and governance in AI, Ethical AI case studies from healthcare, policing, hiring.

UNIT– II

Fairness and Bias in AI: Types of bias: dataset bias, label bias, historical bias, Fairness definitions: demographic parity, equal opportunity, individual fairness, Disparate impact and fairness metrics, Algorithmic audits and bias detection, Fairness-aware learning and mitigation strategies.

UNIT– III

Explainable Artificial Intelligence (XAI) : Need for interpretability in AI models, Taxonomy of XAI methods: model-agnostic, model-specific, LIME, SHAP, Grad-CAM, Partial Dependence Plots, Local vs Global explanations, Trade-offs: accuracy vs interpretability.

UNIT– IV

Accountability and Responsible AI Design : Transparent AI systems, Human-in-the-loop and AI-assisted decision-making, Accountability frameworks (e.g., IEEE, NIST, EU Guidelines), Documentation tools: Datasheets for datasets, Model Cards, Responsible AI lifecycle management.

UNIT– V

Societal Impacts and Policy Considerations :AI in surveillance, misinformation, and social manipulation, Ethical implications in autonomous systems (vehicles, weapons), AI and inclusion: accessibility, gender, race, socioeconomic impacts, Public policy, legal frameworks, and global initiatives, Future challenges and global governance of AI.

Text Books:

1. Fundamentals Joe Reis, Matt Housley, Fundamentals of Data Engineering, O'Reilly Media, Inc.,June 2022,ISBN: 9781098108304

Reference Books:

1. Paul Crickard , Data Engineering with Python, Packt Publishing, October 2020.
2. Ralph Kimball, Margy Ross, The Data Warehouse Toolkit: The Definitive Guide to Dimensional Modeling, Wiley, 3rd Edition, 2013
3. James Densmore, Data Pipelines Pocket Reference: Moving and Processing Data for Analytics, O'Reilly Media, 1st Edition, 2021

Web References:

1. https://technicalsummaries.com/docs/books/fundamentals
2. https://medium.com/@dom.n/the
3. https://www.qlik.com/us/data

Question Paper Pattern:

Sessional Examination:

The question paper for Sessional Examination shall be for 40 marks. The question paper shall consist of four questions and all questions are compulsory. Question No.1 contains five short answer questions (2 marks each) for a total of ten marks. Remaining three questions shall be EITHER/OR type descriptive questions for ten marks each. Each of these descriptive questions may contain sub-questions.

End Examination:

The question paper for End Examination shall be for 70 marks. The question paper shall contain six questions and all questions are compulsory. Question No.1 shall contain ten short answer questions (2 marks each) for a total of twenty marks, with two short answer questions from each unit. Remaining five questions (Each question covering one unit of syllabus) carrying 10 marks each shall be EITHER/OR type descriptive questions and may contain sub-questions.

Advanced Algorithms for AI & ML LAB(AAML(P))								
Honors in AI&ML					**Scheme:** 2023			
Course Code	**Category**	**Hours/Week**			**Credits**	**Maximum Marks**		
HCM06	**H**	**L**	**T**	**P**	**C**	**Continuous Internal Assessment**	**End Exam**	**TOTAL**
		0	**0**	**3**	**1.5**	**30**	**70**	**100**
					End Exam Duration: 3 Hours			

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

CO1:	Implement key machine learning algorithms from scratch and using libraries.
CO2:	Preprocess data and select suitable features for modeling.
CO3:	Train, test, and evaluate models for accuracy and performance.
CO4:	Apply AI techniques to solve classification, regression, and decision-making problems.
CO5:	Develop simple AI agents and use neural networks for predictive tasks.

List of Experiments

1. Data Preprocessing – Cleaning, normalization, encoding, and splitting data.
2. Linear Regression – Implement simple and multiple linear regression.
3. Logistic Regression – Binary classification on datasets like breast cancer or Titanic.
4. K-Nearest Neighbors (KNN) – Classification task with evaluation metrics.
5. Decision Trees and Random Forests – Tree-based classification and visualization.
6. Support Vector Machines (SVM) – Margin classification with kernel trick.
7. Naive Bayes – Text classification with spam dataset.
8. K-Means Clustering – Unsupervised clustering with customer segmentation.
9. Principal Component Analysis (PCA) – Dimensionality reduction and visualization.
10. Artificial Neural Networks (ANNs) – Implement basic neural network using Keras.
11. Model Evaluation & Tuning – Use cross-validation, GridSearchCV, and confusion matrices.
12. AI Agent Search Algorithms – Implement A*, DFS, BFS for pathfinding problems.

Web References

1. https://www.kaggle.com/code/pythonafroz/data-preprocessing-methods-used-in-machinelearning
2. https://www.geeksforgeeks.org/machine-learning/house-price-prediction-using-machine-learning-in-python/
3. https://medium.com/@anibahs/data-preprocessing-for-classification-bb42398b0363
4. https://www.accel.ai/anthology/2024/10/30/automating-model-selection-and-hyperparameter-tuning-with-gridsearchcv
5. https://amueller.github.io/aml/01-ml-workflow/10-model-validation-and-tuning.html

ROBOTICS & AUTONOMOUS SYSTEMS LAB(RAS(P))

Honors in AI&ML | **Scheme:** 2023

Course Code	Category	Hours/Week			Credits	Maximum Marks		
HCM07	H	L	T	P	C	Continuous Internal Assessment	End Exam	TOTAL
		0	0	3	1.5	30	70	100

End Exam Duration: 3 Hours

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

CO1:	Package and deploy AI models using tools such as Flask, FastAPI, and Docker.
CO2:	Automate machine learning workflows using CI/CD pipelines and MLOps tools.
CO3:	Monitor and manage deployed models in real-time environments.
CO4:	Apply version control and model registry techniques effectively.
CO5:	Deploy models on cloud platforms like AWS, Azure, or GCP and use MLflow, Kubeflow.

List of Experiments

1. Build a simple ML model and serve it via Flask or FastAPI.
2. Containerize the model application using Docker.
3. Deploy a Dockerized model on a local or cloud-based Kubernetes cluster
4. Implement CI/CD pipeline using GitHub Actions or GitLab CI.
5. Track experiments and manage model versions using MLflow.
6. Use DVC (Data Version Control) for tracking data and pipeline stages.
7. Automate model retraining and deployment using Jenkins.
8. Model monitoring using Prometheus and Grafana.
9. Introduce model drift detection and retraining triggers.
10. Deploy a model on a cloud platform (e.g., AWS SageMaker, GCP AI Platform).

List of Additional Experiments

1. Use Kubeflow pipelines for end-to-end ML workflow management.
2. Capstone: Full-cycle ML project from training to monitoring using MLOps best practices

References:

1. Mark Treveil and Alok Shukla, AI and Analytics in Production: How to Implement Successful AI and Analytics Applications, O'Reilly Media.
2. Emmanuel Ameisen, Building Machine Learning Powered Applications, O'Reilly Media.
3. Chris Fregly and Antje Barth, Data Science on AWS: Building End-to-End Applications, O'Reilly
4. Alfredo Deza and Noah Gift, Practical MLOps, O'Reilly Media.
5. Soham Kamani, Learning MLOps, Packt Publishing