B. Tech with

HONORS

in

COMPUTER SCIENCE AND ENGINEERING – AIML

Academic Regulations, Course Structure and Syllabus

Effective from 2023-24 admitted batches

Offered by

Department of Computer Science and Engineering

KSRM College of Engineering (A) – Kadapa

(Approved by AICTE, Accredited by NAAC with A+ Grade and NBA and Affiliated to JNTUA, Anantapuramu)

ELIGIBILITY / REGISTRATION / AWARD OF HONORS

The objective of introducing B.Tech. (Hons.) is to facilitate the students to choose additionally the specialized courses of their choice and build their competence in a specialized area in the UG level. The program is a best choice for academically excellent students having good academic record and interest towards higher studies and research.

i) Honors is introduced in the curriculum of all B. Tech. programs offering a major degree and is applicable to all B.Tech (Regular and Lateral Entry) students admitted in Engineering & Technology.

ii) Those students with **at least 7.0 CGPA without any course backlogs up to III Semester in the major degree are only eligible to register for Honor degree**.

iii) A student shall earn **additional 18 credits for award of Honors** from same branch / department / discipline registered for major degree. This is in addition to 163 credits by a regular student and 123 Credits by a Lateral Entry student for the award of Major degree.

iv) A student is permitted to register for Honors in IV Semester after the results of III Semester are declared. Students shall register and pass in all the courses prescribed and being offered from V semester under the respective Honor degree.

v) Students have to attend classwork for courses under Honor degree beyond regular academic hours meant for major degree. Students can also undergo the courses under Honor through any proctored online platforms with the prior approval of the BoS Chairman and the HoD of the respective department offering Honor degree.

vi) The attendance for the registered courses under Honors and regular courses offered for Major degree in a Semester will be considered separately.

vii) A student shall have an aggregate of 75% attendance in all courses registered under Honors in that particular semester to become eligible for attending Semester-End examinations.

viii) The registration for the Honor will be cancelled, if the student is detained due to lack of attendance in Major,

ix) The registration for the Honor will be cancelled, if the student fails in any course of either Honor / Major in any semester from V to VIII Semester.

x) A student registered for Honors shall pass in all subjects that constitute the requirement for the Honors degree program. No class/division (i.e., second class, first class and distinction, etc.) will be awarded for Honors degree program.

xi) A separate grade sheet will be issued for the Honor degree courses semester-wise.

xii) If a student drops or is terminated from the Honors program, the additional credits so far earned cannot be converted into open or core electives; they will remain extra.

xiii) The Honors will be mentioned in the degree certificate as Bachelor of Technology (Honors) in XYZ. For example, B.Tech. (Honors) in Mechanical

Engineering.

xiv) There shall be a minimum enrolment of 20% OR 20 enrollments from the list of eligible students to offer Honors program.

xv) There is no fee for registration of courses for Honors program offered.

xvi) A student can register for either Minor / Honor but not both.

xvii) Student shall submit an application for either Minor / Honor at least one week before the commencement of the V Semester.

HONORS PROGRAMS OFFERED

Offering Department	Title	Who can Register
Civil Engineering	Civil Engineering	B.Tech. CE
	Tunnel Engineering	
	Interior Design	
Mechanical Engineering	Mechanical Engineering	B.Tech. ME
Electrical and Electronics Engineering	Electric Vehicles	B.Tech. EEE
Electronics and Communication Engineering	VLSI	B.Tech. ECE
	Embedded Systems and IoT	
Computer Science and Engineering	Computer Science and Engineering	B.Tech. CSE, B.Tech. AIML, B.Tech. CSE(DS), B.Tech. CSE(AIML)
	Artificial Intelligence and Machine Learning	
	Data Science	
	CSE-Artificial Intelligence and Machine Learning	

COURSE STRUCTURE

for

HONORS

in

COMPUTER SCIENCE AND ENGINEERING – AIML

S.No	COURSE CODE	COURSE TITLE	Semester Offered	L	T	P	IM	EM	CR
1	2333571H	Advanced Algorithms for AI and ML	V	3	0	0	30	70	3
2	2333572H	AI Ethics, Fairness and Explainability	V	3	0	0	30	70	3
3	2333671H	Deep Learning and Neural Networks Architectures	VI	3	0	0	30	70	3
4	2333672H	Reinforcement Learning and Decision Making	VI	3	0	0	30	70	3
5	2333771H	AI for Robotics and Automation	VII	3	0	0	30	70	3
6	2333772H	Applied Project Work	VII	0	0	6	60	140	3
Total				**15**	**0**	**6**	**210**	**490**	**18**

		L	T	P	C
2339571H	**HONORS IN COMPUTER SCIENCE AND ENGINEERING – AIML ADVANCED ALGORITHMS FOR AI AND ML (CSE-AIML)**	**3**	**0**	**0**	**3**

Pre-Requisites: Artificial Intelligence, Machine Learning, Graph Neural Networks

Course Outcomes:

After successful completion of the course, students will be able to:

CO1. Analyze and apply classical algorithmic strategies such as divide and conquer, dynamic programming, approximation, and randomized techniques to solve AI/ML-related computational problems.
CO2. Develop and implement graph-based algorithms for shortest path, flow networks, and community detection, and apply them to real-world AI applications like natural language processing and recommender systems.
CO3. Employ convex and non-convex optimization methods, including gradient descent variants and regularization techniques, to train, tune, and improve performance of AI/ML models.
CO4. Explore and apply metaheuristic approaches—evolutionary algorithms, swarm intelligence, and reinforcement learning—for neural architecture search and complex AI optimization problems.
CO5. Design and evaluate scalable, fair, and interpretable algorithmic solutions in AI/ML, drawing insights from advanced systems like AlphaGo, GPT models, and Auto ML frameworks.

SYLLABUS:

UNIT - I: FOUNDATIONS OF ADVANCED ALGORITHMIC TECHNIQUES (10 Periods)

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, Probabilistic Analysis and Tail Bounds, Applications in ML Preprocessing and Feature Selection

UNIT - II: GRAPH ALGORITHMS AND AI APPLICATIONS (09 Periods)

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 (08 Periods)

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 (10 Periods)

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 (08 Periods)

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

Total Periods:45

TEXT BOOKS:

T1. Introduction to Algorithms, Thomas H. Cormen, Charles E. Leiserson, Ronald L. Rivest, Clifford Stein, MIT Press
T2. Algorithms for Machine Learning, Giuseppe Bonaccorso, Packt Publishing
T3. Convex Optimization, Stephen Boyd, Lieven Vandenberghe, Cambridge University Press
T4. Reinforcement Learning: An Introduction, Richard S. Sutton, Andrew G. Barto

REFERENCE BOOKS:

R1. Machine Learning: A Probabilistic Perspective, Kevin P. Murphy
R2. The Elements of Statistical Learning, Trevor Hastie, Robert Tibshirani, Jerome Friedman
R3. Evolutionary Computation, Kenneth A. De Jong
R4. Handbook of Approximation Algorithms and Metaheuristics, Teofilo F. Gonzalez

ONLINE COURSES:

1. Coursera – Advanced Algorithms and Complexity (UC San Diego)
2. edX – Algorithmic Design and Techniques (UC San Diego)
3. MIT OpenCourseWare – Advanced Algorithms
4. Udemy – Optimization Algorithms in Machine Learning
5. Stanford Online – Convex Optimization

		L	T	P	C
2333572H	**HONORS IN COMPUTER SCIENCE AND ENGINEERING – AIML AI ETHICS, FAIRNESS & EXPLAINABILITY (CSE-AIML)**	**3**	**0**	**0**	**3**

Pre-Requisites: Artificial Intelligence, Machine Learning

Course Outcomes:

After successful completion of the course, students will be able to:

CO1. Describe the ethical principles, historical context, and responsibilities associated with AI deployment 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. Demonstrate the need for explainability in AI models and utilize tools such as LIME, SHAP, and Grad-CAM to generate local and global model explanations.
CO4. Design AI systems with accountability by integrating human oversight, ethical documentation (e.g., Model Cards, Datasheets), and adherence to global guidelines.
CO5. Critically assess the broader societal and legal implications of AI in areas such as surveillance, information, and inclusivity, and explore international policy frameworks.

SYLLABUS:

UNIT - I: FOUNDATIONS OF AI ETHICS (09 Periods)

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 (09 Periods)

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) (09 Periods)

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 (09 Periods)

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 CONSIDERATION

(09 Periods)

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

Total Periods: 45

TEXT BOOKS:

T1. Responsible Artificial Intelligence: How to Develop and Use AI in a Responsible Way, Virginia Dignum, Springer
T2. Weapons of Math Destruction, Cathy O'Neil, Crown Publishing
T3. AI Ethics, Mark Coeckelbergh, The MIT Press

REFERENCE BOOKS:

R1. The Ethics of Artificial Intelligence, Nick Bostrom, Eliezer Yudkowsky
R2. Coded Bias (Documentary, 2020), Shalini Kantayya
R3. Ethics of Information, Luciano Floridi, Oxford University Press

ONLINE COURSES & RESOURCES:

1. Coursera – AI For Everyone (Andrew Ng)
2. edX – Ethics of AI and Big Data (Linux Foundation)

		L	T	P	C
2333671H	**HONORS IN COMPUTER SCIENCE AND ENGINEERING – AIML DEEP LEARNING AND NEURAL NETWORK ARCHITECTURES (CSE-AIML)**	**3**	**0**	**0**	**3**

Pre-Requisites: Artificial Intelligence, Machine Learning

Course Outcomes:

After completion of this course, the student will be able to:

CO1. Understand the theoretical foundations of neural networks and deep learning.
CO2. Implement and train multilayer perceptron, 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.

SYLLABUS:

UNIT - I: FOUNDATIONS OF NEURAL NETWORKS (08 Periods)

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) (08 Periods)

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 (10 Periods)

Sequential Data and Time Series, RNN Basics and Backpropagation 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 (10 Periods)

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 (09 Periods)

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

Total Periods: 45

TEXT BOOKS:

T1. Deep Learning, Ian Goodfellow, Yoshua Bengio, and Aaron Courville (MIT Press)
T2. Neural Networks and Deep Learning, Michael Nielsen (Online Book)
T3. Hands-On Machine Learning with Scikit-Learn, Keras & TensorFlow – Aurélien Géron (O'Reilly)

REFERENCE BOOKS:

R1. Pattern Recognition and Machine Learning, Christopher M. Bishop
R2. Deep Learning for Computer Vision, Rajalingappaa Shanmugamani
R3. Natural Language Processing with Transformers, Lewis Tunstall, Leandro von Werra, Thomas Wolf
R4. Reinforcement Learning: An Introduction, Richard S. Sutton and Andrew G. Barto

RECOMMENDED ONLINE COURSES:

1. Deep Learning Specialization – Andrew Ng (Coursera)
2. CS231n: Convolutional Neural Networks for Visual Recognition (Stanford)
3. Fast.ai – Practical Deep Learning for Coders
4. Deep Learning with PyTorch (Udacity)
5. Transformers by Hugging Face (free course)

		L	T	P	C
2333672H	**HONORS IN COMPUTER SCIENCE AND ENGINEERING – AIML REINFORCEMENT LEARNING AND DECISION MAKING (CSE-AIML)**	**3**	**0**	**0**	**3**

Pre-Requisites: Artificial Intelligence, Machine Learning

Course Outcomes:

After successful completion of this course, students will be able to:

CO1. Explore reinforcement learning fundamentals, including agent-environment interaction, RL types, MDPs, and Bellman equations.
CO2. Apply dynamic programming and Monte Carlo methods for policy evaluation, improvement, and control in model-based RL.
CO3. Implement TD learning algorithms such as TD (0), Sarsa, and Q-learning with enhancements like eligibility traces and function approximation.
CO4. Design and analyze policy gradient and actor-critic methods (e.g., REINFORCE, PPO) for continuous and high-dimensional action spaces.
CO5. Utilize deep RL algorithms (DQN, DDPG, A3C, SAC) and exploration strategies to solve complex tasks in robotics, games, and autonomous systems with attention to safety and ethics.

SYLLABUS:

UNIT - I: INTRODUCTION TO REINFORCEMENT LEARNING & MDPS **(09 Periods)**

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 **(10 Periods)**

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 **(09 Periods)**

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 **(08 Periods)**

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
(09 Periods)

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

Total Periods:45

TEXTBOOKS:

T1. Reinforcement Learning: An Introduction, Richard S. Sutton, Andrew G. Barto, 2nd Edition, MIT Press
T2. Deep Learning, Ian Goodfellow, Yoshua Bengio, Aaron Courville, MIT Press

REFERENCE BOOKS:

R1. Reinforcement Learning Course Slides & Lectures, David Silver, DeepMind, University College London
R2. Reinforcement Learning: State of the Art, Marco Wiering, Martijn van Otterlo (Eds.), Springer
R3. Algorithms for Reinforcement Learning, Csaba Szepesvári, Morgan & Claypool
R4. Deep Reinforcement Learning: An Overview, Yuxi Li, arXiv (survey)
R5.

ONLINE COURSES & RESOURCES:

1. DeepMind x UCL Reinforcement Learning Lectures by David Silver
2. Coursera: Reinforcement Learning Specialization – University of Alberta

		L	T	P	C
2333771H	**HONORS IN COMPUTER SCIENCE AND ENGINEERING – AIML AI FOR ROBOTICS AND AUTOMATION (CSE-AIML)**	**3**	**0**	**0**	**3**

Pre-Requisites: Artificial Intelligence, Machine Learning

Course Outcomes:

After completion of the course, student 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.

SYLLABUS:

UNIT - I: FOUNDATIONS OF AI-DRIVEN SOFTWARE ENGINEERING (08 Periods)

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 (10 Periods)

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 (09 Periods)

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 (10 Periods)

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 (08 Periods)

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

Total Periods:45

TEXTBOOKS:

T1. Artificial Intelligence and Software Engineering: Status and Future Directions, Tim Menzies, Diomidis Spinelli's
T2. DevOps: A Software Architect's Perspective, Len Bass, Ingo Weber, Liming Zhu, Addison-Wesley
T3. Machine Learning (AI Systems Engineering Perspective), Andrew Ng

REFERENCE BOOKS:

R1. AI in Software Engineering, Carlos Nunes Silva Engineering
R2. The DevOps Handbook, Gene Kim, Jez Humble, Patrick Debois, John Willis
R3. Machine Learning: (AI Systems Engineering Perspective), Andrew N

ONLINE RESOURCES & COURSES:

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

		L	T	P	C
2333772H	**HONORS IN COMPUTER SCIENCE AND ENGINEERING – AIML APPLIED PROJECT WORK (CSE-AIML)**	**3**	**0**	**0**	**3**

Pre-Requisites: Programming knowledge on C and Python, Data Structures, Artificial Intelligence, Machine Learning, Deep Learning,

Course Outcomes:

On successful completion of the course, student will be able to

CO1. Analyse and formulate complex problems in Artificial Intelligence and Machine Learning domains by applying foundational concepts and techniques.
CO2. Design AI/ML models and algorithms to develop innovative and efficient solutions for real-world applications.
CO3. Implement appropriate tools, programming languages, and frameworks (such as Python, TensorFlow, PyTorch) to execute AI/ML projects effectively.
CO4. Apply data pre-processing, feature engineering, and model optimization techniques to improve the performance of AI/ML systems.
CO5. Demonstrate project management skills, including planning, collaboration, documentation, and presentation of technical work.
CO6. Propose the solutions by considering professional ethics with focus on societal and environmental aspects.