HONOURS

S.No	Course Code	Course Title	Scheme of Instructions Hours per Week				Scheme of Examination Maximum Marks		
			L	T	P	C	I	E	Total
1	23HRCAI1	Advanced Machine Learning and AI Systems	3	-	-	3	30	70	100
2	23HRCAI2	Deep Learning & Neural Network Architectures	3	-	-	3	30	70	100
3	23HRCAI3	Reinforcement Learning and Decision Making	3	-	-	3	30	70	100
4	23HRCAI4	AI For Robotics and Automation	3	-	-	3	30	70	100
5	23HRCAI5	AI Ethics, Fairness & Explainability	3	-	-	3	30	70	100
6	23HRCAI6	AI and Machine Learning Lab	-	-	3	1.5	30	70	100
7	23HRCAI7	Robotics & Autonomous Systems Lab	-	-	3	1.5	30	70	100

23HRCAI1	**ADVANCED MACHINE LEARNING & AI SYSTEMS**	**L**	**T**	**P**	**C**
		3	0	0	3

Course Objectives:

- To deepen understanding of advanced machine learning concepts including ensemble learning, probabilistic models, and deep neural architectures.
- To explore scalable machine learning algorithms and their applications in real-world AI systems.
- To equip students with knowledge of interpretability, fairness, and trust in AI.
- To understand deployment, monitoring, and life-cycle management of AI systems.
- To apply machine learning in advanced domains such as natural language processing, vision, and multi-agent systems.

Course Outcomes (COs):

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

1. Apply and evaluate ensemble learning, SVMs, probabilistic models, and clustering techniques to solve complex machine learning problems.
2. Design and implement deep learning architectures including CNNs, RNNs, and transformers for vision and sequential data tasks.
3. Analyze and integrate explainability, fairness, and robustness techniques to build trustworthy and ethical AI systems.
4. Develop scalable ML pipelines and deploy models in production with proper monitoring, tuning, and retraining mechanisms.
5. Implement AI techniques in advanced applications such as natural language processing, computer vision, reinforcement learning, and multi-agent systems.

UNIT I: Advanced Supervised and Unsupervised Learning

Ensemble Learning: Bagging, Boosting, Random Forests, Support Vector Machines: Kernels and Multi-Class SVMs, Probabilistic Graphical Models: Bayesian Networks, HMMs, Expectation-Maximization and Variational Inference, Clustering: Hierarchical, DBSCAN, Gaussian Mixture Models.

UNIT II: Deep Learning Architectures

Deep Neural Networks and Optimization Challenges, Convolutional Neural Networks (CNNs), Recurrent Neural Networks (RNNs), LSTMs, GRUs, Autoencoders, Variational Autoencoders, Attention Mechanisms and Transformer Architectures

UNIT III: Interpretability, Fairness, and Trust in AI

Explainable AI: LIME, SHAP, Saliency Maps, Adversarial Examples and Robustness Techniques, Fairness Metrics and Bias Mitigation. Trustworthy AI Design and Ethical Considerations, Model Compression and Distillation

UNIT IV: Scalable and Production ML Systems

ML Pipelines, Feature Stores, and Model Versioning, Distributed Training with TensorFlow and PyTorch, Hyperparameter Tuning at Scale (Ray Tune, Optuna),Model Deployment with Docker, FastAPI, Flask, Monitoring, Drift Detection, and Model Retraining

UNIT V: Advanced Applications and Multi-Agent AI

Natural Language Understanding with Transformers (BERT, GPT), Vision-based AI Systems (YOLO, Mask R-CNN), Reinforcement Learning and Policy Gradients, AI in Robotics and Autonomous Agents, Multi-Agent Systems and Decentralized Learning

Textbooks:

1. Aurélien Géron – Hands-On Machine Learning with Scikit-Learn, Keras, and TensorFlow, 3rd Edition, O'Reilly
2. Ian Goodfellow, Yoshua Bengio, Aaron Courville – Deep Learning, MIT Press

Reference Books:

1. Kevin P. Murphy – Machine Learning: A Probabilistic Perspective, MIT Press
2. Trevor Hastie, Robert Tibshirani, Jerome Friedman – The Elements of Statistical Learning, Springer
3. Chris Bishop – Pattern Recognition and Machine Learning, Springer
4. Ethem Alpaydin – Introduction to Machine Learning, MIT Press

Online Courses & Resources:

1. CS229 – Machine Learning by Stanford (Andrew Ng)
2. DeepLearning.AI – Advanced Deep Learning Specialization (Coursera)

23HRCAI2	**DEEP LEARNING & NEURAL NETWORK ARCHITECTURES**	**L**	**T**	**P**	**C**
		3	0	0	3

Course Objectives:

- To introduce the fundamental concepts and mathematical foundations of deep learning.
- To explore different neural network architectures including CNNs, RNNs, LSTMs, and Transformers.
- To enable students to implement, train, and optimize deep neural networks.
- To analyze the performance and limitations of various architectures in different AI tasks.
- To develop the ability to apply deep learning models to real-world applications such as image recognition, language modeling, and autonomous systems.

Course Outcomes (COs):

Upon successful 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 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 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

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

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)

23HRCAI3	REINFORCEMENT LEARNING & DECISION MAKING	L	T	P	C
		3	0	0	3

Course Objectives:

- To introduce the fundamentals of reinforcement learning (RL) and its mathematical foundation.
- To understand the Markov Decision Process (MDP) framework for decision making under uncertainty.
- To explore various RL algorithms including value-based, policy-based, and model-based approaches.
- To analyze deep reinforcement learning techniques for real-world applications.
- To study the integration of reinforcement learning with planning, exploration, and control strategies.

Course Outcomes (COs):

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

1. 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.
2. Apply dynamic programming and Monte Carlo methods to perform policy evaluation, policy improvement, and control in model-based RL settings.
3. Implement temporal-difference learning algorithms like TD(0), Sarsa, and Q-learning, and extend them using eligibility traces and function approximation techniques.
4. Develop and analyze policy gradient and actor-critic methods, including REINFORCE and PPO, to optimize policies in continuous and high-dimensional action spaces.
5. 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

Textbooks:

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

Online Courses & Resources:

1. DeepMind x UCL Reinforcement Learning Lectures by David Silver
2. Coursera: Reinforcement Learning Specialization – University of Alberta
3. DeepLearning.AI – Deep Reinforcement Learning with TensorFlow

23HRCAI14	**AI FOR ROBOTICS & AUTOMATION**	**L**	**T**	**P**	**C**
		3	**0**	**0**	**3**

Course Objectives:

- To introduce the fundamental concepts of robotics and its integration with artificial intelligence (AI).
- To understand perception, motion planning, and control strategies using AI for autonomous robots.
- To explore machine learning and deep learning approaches in robotic automation.
- To develop intelligent systems capable of navigation, manipulation, and decision-making.
- To understand real-time robotic applications in industry and research.

Course Outcomes (COs):

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

- CO1: Explain the role of AI in robotics and the architecture of intelligent robotic systems.
- CO2: Apply computer vision and sensor fusion techniques for robotic perception.
- CO3: Design motion planning and control algorithms for robotic navigation.
- CO4: Integrate machine learning models for autonomous behavior and adaptation.
- CO5: Analyze applications of AI-powered robots in industrial automation, agriculture, healthcare, and logistics.
- CO6: Develop and evaluate simple autonomous robotic systems with real-time AI decision-making capabilities.

UNIT I: Introduction to AI and Robotics

Fundamentals of Robotics and Components, AI Techniques for Robotics, Types of Robots: Mobile, Industrial, Collaborative, Humanoids, Architectures for Intelligent Robots, Sensors and Actuators in Robotics

UNIT II: Robotic Perception and Computer Vision

Perception Pipeline in Robots, Image Processing & Object Detection, Depth Estimation, 3D Mapping (SLAM), Sensor Fusion (Camera, LiDAR, IMU), Vision-based Navigation and Obstacle Avoidance

UNIT III: Motion Planning and Control

Path Planning: Dijkstra, A*, RRT, PRM, Control Strategies: PID, Feedback Linearization. Trajectory Generation. Kinematics & Dynamics for Robot Manipulators. Motion Planning in Dynamic Environments

UNIT IV: Machine Learning in Robotics

Reinforcement Learning for Control, Supervised Learning for Object Recognition, Unsupervised Learning for Clustering and Mapping, Behavior Cloning and Imitation Learning, Online and Adaptive Learning in Robots

UNIT V: Applications and Trends in Robotic Automation

AI in Industrial Automation and Smart Factories, AI for Service Robots and Human-Robot Interaction (HRI), Robots in Agriculture, Healthcare, and Delivery, Ethical and Social Implications of AI in Robotics, Case Studies: Boston Dynamics, Tesla Bots, Warehouse Automation

Textbooks:

1. Robin R. Murphy – Introduction to AI Robotics, MIT Press
2. Peter Corke – Robotics, Vision and Control: Fundamental Algorithms in MATLAB, Springer
3. Oussama Khatib and Bruno Siciliano (Eds.) – Springer Handbook of Robotics, Springer

Reference Books:

1. Kevin M. Lynch & Frank C. Park – Modern Robotics: Mechanics, Planning, and Control, Cambridge University Press
2. Siegwart, Nourbakhsh, Scaramuzza – Introduction to Autonomous Mobile Robots, MIT Press
3. Stuart Russell & Peter Norvig – Artificial Intelligence: A Modern Approach, Pearson

Online Courses & Resources:

1. Coursera – Robotics Specialization (University of Pennsylvania)
2. edX – AI for Robotics (Columbia University)
3. Udacity – AI for Robotics by Sebastian Thrun

23HRCAI5	AI ETHICS, FAIRNESS & EXPLAINABILITY	L	T	P	C
		3	0	0	3

Course Objectives:

1. To understand the ethical challenges and responsibilities involved in building and deploying AI systems.
2. To identify and mitigate bias in AI data, algorithms, and models.
3. To explore techniques for interpretability and explainability in machine learning.
4. To critically evaluate AI systems from a legal, social, and philosophical perspective.
5. To analyze real-world case studies involving ethical dilemmas in AI deployment.

Course Outcomes:

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

CO1: Demonstrate awareness of ethical and societal concerns associated with AI technologies.

CO2: Detect and reduce bias in datasets and machine learning models.

CO3: Apply explainable AI (XAI) techniques for model transparency.

CO4: Evaluate AI systems based on fairness, accountability, and transparency.

CO5: Reflect on policy, legal, and human-centric implications of AI deployment.

Unit I: Foundations of AI Ethics

Introduction to AI Ethics and Responsible AI, Ethical Theories: Utilitarianism, Deontology, Virtue Ethics, Key Ethical Principles: Fairness, Accountability, Transparency, Privacy (FATP), Human-in-the-loop and Ethical Decision Making, AI and the SDGs (Sustainable Development Goals)

Unit II: Bias in Data and Algorithms

Types of Bias: Historical, Representation, Measurement, Aggregation, Sources of Bias in AI: Data Collection, Annotation, Model Training, Metrics for Fairness: Demographic Parity, Equal Opportunity, Predictive Parity, Bias Mitigation Techniques: Pre-processing, In-processing, Post-processing, Case Studies: COMPAS, Hiring Algorithms, Face Recognition Bias

Unit III: Explainable AI (XAI) Techniques

Need for Explainability and Transparency, Global vs Local Explanations, Methods: LIME, SHAP, Anchors, Integrated Gradients, Model-specific vs Model-agnostic Explanations, Visual Explanations and Human-Centric Interpretability

Unit IV: Legal, Regulatory & Societal Aspects

Data Protection Laws: GDPR, CCPA, Indian Digital Personal Data Protection Act, Ethical Guidelines: IEEE, UNESCO, OECD AI Principles, Algorithmic Accountability and Auditing, Intellectual Property and Liability in AI, Ethical Considerations in Surveillance, Military, and Social Scoring

Unit V: Building Responsible AI Systems

Designing Ethical AI Systems: Frameworks and Toolkits, Human-Centered AI and Value Alignment, Responsible AI Lifecycle and Documentation (Model Cards, Data Sheets), AI for Good and Ethical Innovation, Industry Case Studies: Google, Microsoft, IBM‘s AI Governance

Textbooks:

1. **Virginia Dignum**, Responsible Artificial Intelligence: How to Develop and Use AI in a Responsible Way, Springer, 2019.
2. **Cathryn Carson and John Zerilli**, Ethics and Data Science, O'Reilly Media, 2021.
3. **Patrick Lin, Keith Abney, Ryan Jenkins**, Robot Ethics 2.0: From Autonomous Cars to Artificial Intelligence, Oxford University Press, 2017.

Reference Books:

1. **Shalini Sharma, B. Ravindran**, Responsible AI: An Indian Perspective, Springer, 2023.
2. **Christopher Kuner et al.**, The GDPR: General Data Protection Regulation (EU) Regulation 2016/679, Oxford University Press.

23HRCAI6	AI & MACHINE LEARNING LAB	L	T	P	C
		0	0	3	1.5

Course Objectives

1. To provide hands-on experience in implementing AI and machine learning algorithms.
2. To develop and evaluate models using real-world datasets.
3. To introduce optimization and hyperparameter tuning techniques.
4. To build intelligent systems for classification, prediction, and clustering.

Course Outcomes (CO)

After completing this lab, 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.

Tools Required

- Python (NumPy, Pandas, Scikit-learn, TensorFlow/Keras, OpenCV)
- Jupyter Notebook / Google Colab
- Datasets from UCI, Kaggle, Scikit-learn
- Anaconda / VS Code

List of 12 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.

23HRCAI7	ROBOTICS & AUTONOMOUS SYSTEMS LAB	L	T	P	C
		0	0	3	1.5

Course Objectives:

- To introduce students to the fundamental concepts of robotics, control, and autonomous navigation.
- To provide hands-on experience with robotic simulation tools and real-time robot programming.
- To explore sensor integration, motion planning, and autonomous decision-making.
- To familiarize students with ROS (Robot Operating System) and robotic hardware platforms.
- To apply AI and machine learning concepts in robotics for perception and autonomy.

Course Outcomes:

By the end of this course, students will be able to:

- Understand and implement kinematics and control algorithms for robotic systems.
- Program robots using ROS and simulate them in environments like Gazebo or Webots.
- Integrate sensors such as LIDAR, cameras, and IMUs for perception.
- Develop algorithms for autonomous navigation, obstacle avoidance, and mapping.
- Apply AI and computer vision techniques in robotic decision-making.

List of 12 Lab Experiments:

1. Experiment 1: Introduction to Robot Operating System (ROS) and workspace setup.
2. Experiment 2: Build a basic ROS publisher and subscriber for robot control.
3. Experiment 3: Simulate a differential drive robot in Gazebo or Webots.
4. Experiment 4: Implement forward and inverse kinematics for a 2-link robotic arm.
5. Experiment 5: Control robot movement using PID control in simulation.
6. Experiment 6: Interface and process data from ultrasonic/IR sensors.
7. Experiment 7: Integrate and visualize LIDAR data for environment sensing.
8. Experiment 8: Implement SLAM (Simultaneous Localization and Mapping) using Gmapping or Cartographer.
9. Experiment 9: Develop a path planning algorithm using A* or Dijkstra.
10. Experiment 10: Obstacle avoidance using sensor data and reactive behavior.
11. Experiment 11: Vision-based object detection and tracking using OpenCV.
12. Experiment 12: Mini project – Build a complete pipeline for autonomous navigation in a mapped environment.

Textbooks:

1. Roland Siegwart, Illah Nourbakhsh, and Davide Scaramuzza, Introduction to Autonomous Mobile Robots, MIT Press.
2. John J. Craig, Introduction to Robotics: Mechanics and Control, Pearson.