CURRICULUM & SYLLABI

M.Tech. (Artificial Intelligence)

COMPUTER SCIENCE & ENGINEERING

Effective from AY: 2026-2027

NATIONAL INSTITUTE OF TECHNOLOGY WARANGAL

WARANGAL, TELANGANA

Table of Contents

Vision and Mission of the Institute **3**

Vision and Mission of the Department **3**

Program Educational Objectives **4**

Program Outcomes **5**

CURRICULUM **6**

SYLLABI **10**

Vision and Mission of the Institute

National Institute of Technology Warangal

VISION

Towards a Global Knowledge Hub, striving continuously in pursuit of excellence in Education, Research, Entrepreneurship and Technological services to the society

MISSION

- Imparting total quality education to develop innovative, entrepreneurial and ethical future professionals fit for globally competitive environment.
- Allowing stake holders to share our reservoir of experience in education and knowledge for mutual enrichment in the field of technical education.
- Fostering product-oriented research for establishing a self-sustaining and wealth creating centre to serve the societal needs.

Vision and Mission of the Department

Computer Science & Engineering

VISION

Attaining global recognition in Computer Science & Engineering education, research and training to meet the growing needs of the industry and society.

MISSION

- MS1: Imparting quality education through well-designed curriculum in tune with the challenging software needs of the industry.
- MS2: Providing state-of-art research facilities to generate knowledge and develop technologies in the thrust areas of computer science and engineering.
- MS3: Developing linkages with world class organizations to strengthen industry-academia relationships for mutual benefit.

Program: M.Tech. (Artificial Intelligence)
Program Educational Objectives

PEO-1	Design, develop, and test Artificial Intelligence solutions for societal and industrial problems
PEO-2	Design and evaluate efficient artificial intelligence algorithms to address complex problems
PEO-3	Work in teams using common Artificial Intelligence and Data science tools to achieve project objectives
PEO-4	Communicate effectively and demonstrate professional ethics with societal responsibilities.
PEO-5	Engage in lifelong learning to keep pace with the changing landscape of AI tools and technologies for professional advancement.

Program Articulation Matrix

PEO / Mission Statements	PEO-1	PEO-2	PEO-3	PEO-4	PEO-5
Imparting quality education through well-designed curriculum in tune with the challenging software needs of the industry	3	2	1	1	2
Providing state-of-art research facilities to generate knowledge and develop technologies in the thrust areas of computer science and engineering.	2	2	2	-	2
Developing linkages with world class organizations to strengthen industry-academia relationships for mutual benefit	2	2	2	2	1

1 - Slightly; 2 - Moderately; 3 - Substantially

Program: M. Tech. (Artificial Intelligence)
Program Outcomes

PO-1	Engage in critical thinking and pursue investigations / research and development to solve practical problems.
PO-2	Communicate effectively, write and present technical reports on complex engineering activities by interacting with the engineering fraternity and with society at large.
PO-3	Demonstrate a higher level of professional skills to tackle multidisciplinary and complex problems related to Artificial Intelligence.
PO-4	Design efficient Artificial Intelligence and Data Science algorithms to solve the industrial and societal problems.
PO-5	Apply AI-driven high-performance computing platforms to efficiently solve complex problems with optimized resource utilization.
PO-6	Develop customized, robust, resilient, secure, and scalable AI solutions that adhere to ethical and responsible practices for real-world problems.

MAPPING OF PROGRAM OUTCOMES WITH PROGRAME EDUCATIONAL OBJECTIVES

PO	PEO-1	PEO-2	PEO-3	PEO-4	PEO-5
1	2	3	2	-	-
2	1	-	-	3	1
3	2	2	3	-	1
4	3	3	2	3	2
5	3	3	2	3	1
6	2	3	3	2	2

CURRICULUM
Program: M.Tech. (Artificial Intelligence)

1st Semester

S.No.	Code	Course Title	L-T-P	Credits
1	CS36001	Mathematical Foundations for Data Science	3-0-0	3
2	CS36003	Foundations of AI Models	3-0-0	3
3	CS36005	Computational Methods of Optimization	3-0-0	3
4	CS360XX	Professional Elective – I	3-0-0	3
5	CS360XX	Professional Elective – II	3-0-0	3
6	CS36007	AI and MLOps Laboratory	0-0-3	2
7	CS36009	Data Science Laboratory	0-1-2	2
		Total Credits		**19**

2nd Semester

S.No.	Code	Course Title	L-T-P	Credits
1	CS36002	Reinforcement Learning	3-0-0	3
2	CS36004	Deep Learning Models	3-0-0	3
3	CS360XX	Professional Elective - III	3-0-0	3
4	CS360XX	Professional Elective - IV	3-0-0	3
5	CS360XX	Professional Elective - V	3-0-0	3
6	CS36006	Advanced Data Structures Practice	1-0-4	3
7	CS36008	Deep Learning Laboratory	0-0-2	1
8	CS36095	Minor Project	0-0-0	2
		Total Credits		**21**

3rd Semester

S.No.	Code	Course Title	L-T-P	Credits
1	CS36089	Seminar and Technical Writing	0-0-0	2
2	CS36091	Summer Internship/Research Experience	0-0-0	2
3	CS36093	Comprehensive Viva-Voce	0-0-0	2
4	CS36097	Dissertation: Part A	0-0-0	8
			Total Credits	**14**

4th Semester

S.No.	Code	Course Title	L-T-P	Credits
1	CS36098	Dissertation: Part B	0-0-0	12
			Total Credits	**12**

Professional Elective Courses:

Professional Elective-I/II				
S.No.	**Code**	**Course Title**	**L-T-P**	**Credits**
1	CS36021	Recommender Systems	3-0-0	3
2	CS36023	Advanced Machine Learning	3-0-0	3
3	CS36025	Computer Vision	3-0-0	3
4	CS36027	Big Data Analytics	3-0-0	3
5	CS36029	IoT Data Analytics	3-0-0	3
6	CS36031	Advanced Natural Language Processing	3-0-0	3
7	CS36033	Smart Edge Computing	3-0-0	3
8	CS36035	AI in FinTech	3-0-0	3
9	CS36037	AI for Software Engineering	3-0-0	3
10	CS36039	Intelligent Data Networks	3-0-0	3
11	CS36041	Data Lakes and Data Fabric	3-0-0	3
12	CS36043	HPC Systems for AI	3-0-0	3
13	CS36045	Agentic AI	3-0-0	3
14	CS36047	Responsible AI	3-0-0	3

Professional Elective-III/ IV/ V				
S.No.	**Code**	**Course Title**	**L-T-P**	**Credits**
1	CS36022	Quantum Machine Learning	3-0-0	3
2	CS36024	Cloud AI	3-0-0	3
3	CS36026	Autonomous Vehicles and Drone Technologies	3-0-0	3
4	CS36028	AI for Cyber Security	3-0-0	3
5	CS36030	AI for Digital Forensics	3-0-0	3

6	CS36032	Generative AI	3-0-0	3
7	CS36034	AI for Graph Optimization	3-0-0	3
8	CS36036	AI for Healthcare	3-0-0	3
9	CS36038	Social Network Analysis	3-0-0	3
10	CS36040	AI for Cyber Physical Systems	3-0-0	3
11	CS36042	Explainable AI	3-0-0	3
12	CS36044	Hardware Acceleration and Optimization for AI	3-0-0	3
13	CS36046	AI for Data Privacy	3-0-0	3
14	CS36048	Federated Learning	3-0-0	3
15	CS36050	Algorithmic Game Theory	3-0-0	3
16	CS36052	Testing AI Models	3-0-0	3
17	CS36054	Robotics	3-0-0	3
18	CS36056	Multi-modal Data Analytics	3-0-0	3
19	CS36058	eXtended Reality (XR)	3-0-0	3
20	CS36060	AI for Malware Analysis	3-0-0	3

CS36001	MATHEMATICAL FOUNDATIONS FOR DATA SCIENCE	Credits 3-0-0: 3

Pre-Requisites: **None**

Course Outcomes:

CO-1	Apply statistical methods to data for inferences.
CO-2	Analyze data using classification, graphical, and computational methods
CO-3	Understand Data Wrangling approaches.
CO-4	Perform descriptive analytics over massive data.

Course Articulation Matrix:

	PO1	PO2	PO3	PO4	PO5	PO6
CO1	3	2	3	1	-	-
CO2	2	-	1	2	-	-
CO3	1	3	2	3	-	-
CO4	3	-	3	3	-	-

1 - Slightly; 2 - Moderately; 3 - Substantially

Syllabus:

Overview of Random variables and distributions, Statistical learning, Assessing model accuracy, Descriptive Statistics, Dependent and Independent events.

Linear Algebra: Vectors and tensors; Matrices; Basis; Norms; Orthonormality; Linear System, Rank, and Solution; Linear Transformation; Matrix Decomposition: QR Factorization; Cholesky Decomposition; PCA and Singular Value Decomposition; Eigen Decomposition; Programming for basic computational methods such as Eigen values and Eigen vectors, sparse matrices, and Interpolation by divided differences.

Analytic Geometry- Inner Products, Lengths and Distances, Angles and Orthogonality. Matrix Decompositions-Eigenvalues and Eigenvectors, Cholesky Decomposition, Eigen decomposition and Diagonalization, Singular Value Decomposition, Matrix Approximation, matrix Phylogeny.

Vector Calculus: Differentiation of Univariate Functions, Partial Differentiation and Gradients, Gradients of Vector-Valued Functions, Gradients of Matrices.

Probability and Distributions: Construction of a Probability Space, Discrete and Continuous Probabilities, Sum Rule, Product Rule, and Bayes Theorem, Summary Statistics and Independence, Gaussian Distribution.

Dimensionality Reduction with Principal Component Analysis- Problem Setting, Maximum Variance Perspective, Projection Perspective, Eigen vector Computation and Low-Rank, Approximations, PCA in High Dimensions.

Linear Regression: Simple and multiple linear regressions, Comparison of Linear regression with K-nearest neighbors. Simple Hypothesis Testing, Student's t-test, paired t and U test, correlation and covariance, tests for association. Classification: Linear and Logistic Regression, LDA, and comparison of classification methods.

Graphical Analysis: Histograms and frequency polygons, Box-plots, Quartiles, Scatter Plots, Heat Maps, t-Distributed Stochastic Neighbor Embedding (t-SNE).

Learning Resources:

Text Books:

1. Marc Peter Deisenroth, A. Aldo Faisal, and Cheng Soon Ong. *"Mathematics for Machine Learning".* Cambridge University Press, 2020
2. Gibert Strang. "*Linear Algebra and Learning from Data*. Wellesley-Cambridge Press, 2019
3. Bertsekas, Dimitri P. and John N. Tsitsiklis. "*Introduction to Probability"*. Vol. 1. Belmont, MA: Athena Scientific, 2002.
4. Robert Nowak, "Mathematical Foundations of Machine Learning", 2022

CS36003	FOUNDATIONS OF AI MODELS	Credits 3-0-0: 3

Pre-Requisites: **None**

Course Outcomes:

At the end of the course, the student will be able to

CO1	Create, learn various agent models and design knowledge-based agents.
CO2	Learn and implement various searching techniques on live examples.
CO3	Represent knowledge using Propositional Logic and First Order Logic and perform logical inference.
CO4	Build optimal classifiers using Naïve Bayesian, Decision Trees, SVM and Regression
CO5	Design Unsupervised Learning Methods and Dimensionality reduction methods

Course Articulation Matrix:

	PO1	PO2	PO3	PO4	PO5	PO6
CO1	1	2	3	3	-	-
CO2	-	2	3	3	-	-
CO3	1	-	3	3	-	-
CO4	-	2	3	3	-	-
CO5	2	-	3	3	-	-

1 - Slightly; 2 - Moderately; 3 – Substantially

Syllabus:

INTRODUCTION – AI overview, Intelligent Agents – Agents and Environments, Good Behavior: The Concept of Rationality, The Structure of Agents; Solving Problems By Search – Problem-Solving Agents, Formulating problems, Searching for Solutions, Uninformed Search Strategies, Breadth-first search, Depth-first search, Depth-limited search, Bidirectional search, Informed (Heuristic) Search Strategies, Greedy best-first search, A* Search: Minimizing the total estimated solution cost, Heuristic Functions, Local Search Algorithms and Optimization Problems, –Adversarial Search – Games, The minimax algorithm, Optimal decisions in multiplayer games, Alpha-Beta Pruning, Evaluation functions; Logical Agents – Knowledge-Based agents, The Wumpus World, Logic, Propositional Logic: A Very Simple Logic, Reasoning Patterns in Propositional Logic; First Order Logic – Syntax and Semantics of First-Order Logic, Using First-Order Logic, Knowledge Engineering in First-Order Logic; Inference In First Order Logic – Propositional vs. First-Order Inference, Unification, Introduction to Machine Learning, Supervised and Unsupervised Learning Methods, Bayes Theorem, Naïve Bayesian Classification, Decision Trees, SVM; Clustering – Partition-based, Hierarchal and density based; Regression Models, Evaluation of Classification and Clustering approaches.

Text Books:

1. Stuart Russell and Peter Norvig, *Artificial Intelligence, A Modern Approach*, Pearson Publishers, 2020, 4th Edition.
2. Nils J Nilsson, "*Artificial Intelligence: A New Synthesis*", Morgan Kaufmann Publications, 2000.
3. Tom Mitchell, *Machine Learning*", McGrawHill publications, 2017
4. Chritopher Bishop, *Pattern recognition and machine learning, 2009*

CS36005	COMPUTATIONAL METHODS OF OPTIMIZATION	Credits 3-0-0: 3

Pre-Requisites: **None**

Course Outcomes:

At the end of the course, the student will be able to

CO1	Formulate and solve linear programming model.
CO2	Formulate and solve Non- linear programming model.
CO3	Solve real life problems using Meta-heuristic techniques.
CO4	Generate random numbers and random variates.

Course Articulation Matrix:

	PO1	PO2	PO3	PO4	PO5	PO6
CO1	2	2	3	2	-	3
CO2	2	2	3	2	-	3
CO3	2	2	3	2	-	3
CO4	2	2	3	2	-	3

1 - Slightly; 2 - Moderately; 3 – Substantially

Syllabus:

Need for unconstrained methods in solving constrained problems. Necessary conditions of unconstrained optimization, Methods of line search, Goldstein and Wolfe conditions for partial line search. Global convergence theorem, Steepest descent method. Quasi-Newton methods: DFP, BFGS, Broyden family. Conjugate-direction methods: Fletcher-Reeves, Polak-Ribière. KKT conditions, Duality, Linear Programming, Simplex method, Gradient Projection, Penalty methods.

Introduction to Quadratic Programming: Constrained Optimization Problem Solving, Convex Optimization Methods.

Heuristic and Meta Heuristic Programming: Simulated Annealing, Genetic Algorithm, Particle Swarm Optimization algorithm, Ant Colony Optimization, and Teaching Learning-Based Optimization algorithm - Non-Linear Programming algorithms.

Simulation Modelling: Random number generation, Random variate generation – Verification and Validation of simulation models, Simulation of Computer Systems and Computer Networks.

Text Books

1. Hamdy A Taha – "*Operations Research-An Introduction*", 10th Edition, Pearson, 2022 (Chapters 1-8, 12, 14, 17)
2. Jerry Banks, Hon S Carson, Barry L Nelson, David M Nicol, "*Discrete Event Simulation*", 5th Edition, Pearson, 2013 (Chapters 8 – 12, 14, 15).
3. Fletcher R., *Practical Methods of Optimization*, John Wiley, 2000.

Online Resources:

1. David G. Luenberger and Yinyu Ye, "*Linear and Nonlinear Programming"*, Fourth edition, Springer, 2008
2. Yurii Nesterov,*Introductory, "Lectures on Convex Optimization"*, Second edition, 2018, Springer.

CS36007	AI and MLOps LABORATORY	Credits 0-0-3: 2

Pre-Requisites: **None**

Course Outcomes:

At the end of the course, the student will be able to

CO1	Design various knowledge-based agent models and search techniques
CO2	Implement various end-to-end ML production workflows
CO3	Represent knowledge using Propositional Logic and First Order Logic and perform logical inference.
CO4	Design Supervised and Unsupervised Learning Methods and Dimensionality Reduction Methods
CO5	Implement CI/CD pipeline to testing, training, and deployment

Course Articulation Matrix:

	PO1	PO2	PO3	PO4	PO5	PO6
CO1	1	-	3	3	3	3
CO2	-	-	3	3	3	3
CO3	-	-	3	3	3	3
CO4	1	1	3	3	3	3
CO5	1	2	3	3	3	3

1 - Slightly; 2 - Moderately; 3 – Substantially

Syllabus:

1. Design a 3×3 digital board game where a human play against the computer. The system should evaluate all possible future outcomes of a move before selecting the best one. Ensure the computer never loses if a draw is possible.
2. Create a simplified 8×8 strategy board where a computer chooses moves by predicting opponent responses up to 2–3 future turns. The system should score board positions and select the move leading to the most favourable outcome.
3. Build a puzzle solver that rearranges a scrambled 3×3 tile board into its goal configuration. The system should measure how close each move is to the goal and systematically explore states while avoiding repetitions.
4. Uniform Cost Search (UCS) - Design a navigation system that finds the lowest travel cost between two cities in a weighted road network. The system must always expand the partial route with the smallest accumulated travel cost and update paths if a cheaper option appears.
5. Depth Limited Search (DLS) - Create a network crawler that searches for a target server but stops exploring beyond a fixed depth limit. Display visited nodes and indicate when the search stops due to reaching the maximum depth.
6. Iterative Deepening Depth First Search (IDDFS) - Modify the crawler so it starts searching with depth 0 and gradually increases the depth limit until the target is found. Compare its memory usage with a full level-by-level exploration strategy.
7. Bidirectional Search - Design a routing system that simultaneously searches from the accident location and the hospital location in a road network. Stop when both searches meet and reconstruct the shortest connecting path.
8. Greedy Best First Search (GBFS) - Develop a delivery planner that chooses the next city based only on estimated closeness to the destination. Demonstrate a scenario where this strategy fails to give the shortest overall route.
9. A* Search Algorithm - Enhance the planner so it considers both distances already travelled and estimated remaining distance. Show through experiments that the system finds the optimal path when the estimate never overstates the true distance.

10. Optimal Decision Making in Multiplayer Games - A company is building an AI agent to compete in a strategic multiplayer bidding game where multiple participants try to maximize their profit. Design an intelligent decision-making module that computes the optimal move considering rational opponents.
11. An autonomous exploration robot is deployed in a hazardous mining field where dangerous zones emit warning signals in nearby cells. Design a knowledge-based agent that uses local percepts to infer safe paths and reach a target location without prior knowledge of hazard positions. The hazards and signals must be randomly initialized each run. Demonstrate logical reasoning and safe decision-making under uncertainty.
12. A startup wants an AI assistant that can automatically organize a student's daily academic tasks based on priority, deadlines, and available time. Design a simple agent that takes task inputs, reasons about constraints, and generates an optimized action plan for the day. The agent should revise its plan if new tasks are added. Demonstrate autonomous decision-making and rule-based reasoning.
13. Build a decision tree model for a classification problem.
14. Implement the Naïve Bayesian and SVM algorithm to obtain the classifier for the given dataset.
15. Design and implement PCA and regression model for any given dataset.
16. Implement the Clustering methods for any given dataset.

MLOps-Model development, model deployment, monitoring, and optimization using MLOps tools and frameworks. CI/CD pipeline construction to automate training and testing. To track model drift and performance using tools. Implement MLOps on cloud platforms

Text Books

1. Mark Treveil and the Dataiku Team, "Introducing MLOps", O'Reilly books, 2020.
2. Noah Gift, Alfredo Deza, "Practical MLOps", O'Reilly books, 2022.
3. Stuart Russell and Peter Norvig, *Artificial Intelligence, A Modern Approach*, Pearson Publishers, 2020, 4th Edition.
4. Nils J Nilsson, "*Artificial Intelligence: A New Synthesis*", Morgan Kaufmann Publications, 2000.
5. Tom Mitchell, "*Machine Learning",* , McGrawHill publications, 2025
6. Chritopher Bishop, "*Pattern recognition and machine learning*", First edition, 2006

CS36009	DATA SCIENCE LABORATORY	Credits 0-1-2: 2

Pre-Requisites: **None**

Course Outcomes:

At the end of the course, the student will be able to

CO1	Apply statistical methods to data for inferences.
CO2	Data Curation, data wrangling, data acquisition and feature engineering
CO3	Handling Large Scale Data and Hypothesis Testing
CO4	Data Visualization

Course Articulation Matrix:

	PO1	PO2	PO3	PO4	PO5	PO6
CO1	1	-	3	2	-	3
CO2	1	-	3	2	-	3
CO3	-	-	3	2	-	3
CO4	-	2	3	2	-	3

1 - Slightly; 2 - Moderately; 3 – Substantially

Syllabus:

1. Descriptive statistics and confidence interval estimation (including bootstrapping).
2. Correlation/covariance analysis and regression inference using p-values and confidence intervals.
3. Logistic regression inference (odds ratio, significance) with ROC-AUC and calibration.
4. Data acquisition from API/CSV/database and integration into a curated dataset.
5. Data cleaning pipeline: missing values, outliers, duplicates, inconsistent categories.
6. Feature engineering and feature selection (encoding, scaling, mutual information/chi-square).
7. Large-scale data processing using Spark/Dask: joins, group-bys, window functions with benchmarking.
8. Hypothesis testing on large datasets (t-test/ANOVA/chi-square) with multiple-testing correction.
9. Exploratory data visualization report/dashboard (distributions, boxplots, heatmaps, insights).
10. High-dimensional visualization using PCA + t-SNE/UMAP and interpretation of separability
11. Exploration of latest tool like WEKA, Tableau, Graph Databases, Semantic Search Databases.

Text Books

1. Wes McKinney**,** *Python for Data Analysis (3rd Edition)***,** O'Reilly, 2022.
2. Mark Gardener, Beginning R, *The Statistical Programming Language,* Wiley, 2015.
3. Trevor Hastie, Robert Tibshirani, Jerome Friedman, *The Elements of Statistical Learning (2nd Edition)***,** *Springer, 2009.*
4. Gareth James, Daniela Witten, Trevor Hastie, Robert Tibshirani, *An Introduction to Statistical Learning (2nd Edition)***,** Springer, 2021.
5. Jiawei Han, Micheline Kamber, Jian Pei, *Data Mining: Concepts and Techniques* (3rd Edition)**,** Morgan Kaufmann, 2012.
6. Jules J. Berman**,** *Principles of Big Data: Preparing, Sharing, and Analyzing Complex Information (2nd Edition)***,** Morgan Kaufmann, 2016.

Reference Books

1. Peter Bruce, Andrew Bruce, Peter Gedeck**,** *Practical Statistics for Data Scientists* (2nd Edition)**,** O'Reilly, 2020.
2. Alberto Cairo, *The Truthful Art: Data, Charts, and Maps for Communication***,** New Riders, 2016**.**
3. Claus O. Wilke**,** *Fundamentals of Data Visualization***,** O'Reilly, 2019.
4. Bill Chambers, Matei Zaharia**,** *Spark: The Definitive Guide***,** O'Reilly, 2018.
5. Kevin P. Murphy**,** *Probabilistic Machine Learning: An Introduction***,** MIT Press, 2022.

Online Resources

1. scikit-learn User Guide (models, metrics, CV, feature selection, PCA): https://scikit-learn.org/stable/user_guide.html
2. pandas Documentation (wrangling, joins, missing values): https://pandas.pydata.org/docs/
3. SciPy Stats (hypothesis tests): https://docs.scipy.org/doc/scipy/reference/stats.html
4. statsmodels (regression inference, p-values, CIs): https://www.statsmodels.org/stable/
5. Matplotlib Tutorials (core plotting): https://matplotlib.org/stable/tutorials/
6. Seaborn Tutorial (statistical visualization): https://seaborn.pydata.org/tutorial.html
7. Plotly Python (interactive dashboards): https://plotly.com/python/
8. Apache Spark Documentation (large-scale processing): https://spark.apache.org/docs/latest/
9. Dask Documentation (parallel / large-scale pandas-like computing): https://docs.dask.org/en/stable/
10. Kaggle Learn (hands-on DS mini-courses + datasets): https://www.kaggle.com/learn

2nd Semester

CS36002	REINFORCEMENT LEARNING	Credits 3-0-0: 3

Pre-Requisites: **NIL**

Course Outcomes:

CO-1	Understand elements of RL and solutions to Multi-Arm Bandit Problem.
CO-2	Formulate and solve problems modelled with Markov Reward Process.
CO-3	Comprehend solutions for problems with Markov Decision Process.
CO-4	Apply Dynamic Programming for Markov Decision Process
CO-5	Analyze Temporal Difference Methods

Course Articulation Matrix:

	PO-1	PO-2	PO-3	PO-4	PO-5	PO-6
CO-1	1	-	3	2	3	2
CO-2	3	-	3	-	3	-
CO-3	3	-	3	-	3	2
CO-4	2	-	3	3	3	2
CO-5	2	-	3	2	3	1

1 - Slightly; **2 - Moderately;** **3 - Substantially**

Syllabus:

Introduction: Introduction to Reinforcement Learning (RL) – Difference between RL and Supervised Learning, RL and Unsupervised Learning. Elements of RL, Markov property, Markov chains, Markov reward process (MRP). Evaluative Feedback - Multi-Arm Bandit Problem: An n-Armed Bandit Problem, Exploration vs Exploitation principles, Action value methods, Incremental Implementation, tracking a non-stationary problem, optimistic initial values, upper-confidence-bound action selection, Gradient Bandits. Introduction to and proof of Bellman equations for MRPs. Introduction to Markov decision process (MDP), state and action value functions, Bellman expectation equations, optimality of value functions and policies, Bellman optimality equations. Dynamic Programming (DP): Overview of dynamic programming for MDP, principle of optimality, Policy Evaluation, Policy Improvement, policy iteration, value iteration, asynchronous DP, Generalized Policy Iteration. Monte Carlo Methods for Prediction and Control: Overview of Monte Carlo methods for model free RL, Monte Carlo Prediction, Monte Carlo estimation of action values, Monto Carlo Control, On policy and off policy learning, Importance sampling. Temporal Difference Methods: TD Prediction, Optimality of TD (0), TD Control methods - SARSA, QLearning and their variants.

Learning Resources:

Text Books:

1. Richard S. Sutton and Andrew G. Barto, "*Reinforcement Learning: An Introduction"*, MIT Press, Second Edition, 2018.

Reference Books:

1. Csaba Szepesvari, "*Algorithms for Reinforcement Learning"*, Morgan & Claypool, First edition, 2010.
2. Warren B. Powell, "*Reinforcement Learning and Stochastic Optimization*: A Unified Framework for Sequential Decisions", First Edition, Wiley, 2022.

CS 36004	DEEP LEARNING MODELS	Credits 3-0-0: 3

Pre-Requisites: **None**

Course Outcomes:

At the end of the course, the student will be able to

CO1	Identify Convolutional Neural Networks models to solve Supervised Learning Problems
CO2	Design Autoencoders to solve Unsupervised Learning problems
CO3	Apply Long Shot Term Memory (LSTM), Bi-LSTM and Transformer Networks for time series analysis classification problems
CO4	Apply Classical Supervised Tasks for Image Denoising, Segmentation and Object detection problems

Course Articulation Matrix:

	PO1	PO2	PO3	PO4	PO5	PO6
CO1	-	-	2	-	3	3
CO2	-	1	-	2	3	3
CO3	-	1	2	3	3	3
CO4	-	-	3	3	3	3

1 - Slightly; 2 - Moderately; 3 – Substantially

Syllabus:

Introduction to Deep learning, optimization algorithms: Gradient descent methods, RMSProp, Adam, Transfer learning, Regularization Methods, Gradients Vanishing and Exploding Problem. Convolutional Neural Networks and CNN architectures: AlexNet, VGGNet, ResNet, GoogLeNet, InceptionNet, DenseNet, MobileNet, EfficientNet; Early stopping, dropout, batch normalization, instance normalization, group normalization; Advanced Sequence Models: RNN, LSTM, Bi-LSTM, and Transformer Modes: multi-headed self-attention, cross-attention and Transformer types, Language Models, Generative models: Autoencoders and Types, Generative Adversarial Networks, Diffusion Models and applications. Graph Neural Networks (GNNs). Object detection: RCNN, FRCNN and YOLO types and Segmentation methods. Recent trends: zero-shot, one-shot, few-shot learning.

Learning Resources:

Text Books

1. C.M.Bishop and H.Bishop, *Deep Learning: Foundations and Concepts*, Springer, 2024
2. Charu C. Aggarwal, *Neural Networks and Deep Learning*, Springer, 2nd Ed., 2023
3. A. Zhang, Z.C.Lipton, M.Li and A.J.Smola, *Dive into Deep Learning*, Cambridge University Press, 2024

CS36006	ADVANCED DATA STRUCTURES PRACTICE	Credits 1-0-4: 3

Pre-Requisites: Data Structures, Design and Analysis of Algorithms, Programming in C/C++/Java/Python

Course Outcomes:

At the end of the course, the student will be able to

CO1	Implement and analyze advanced tree, string, and probabilistic data structures.
CO2	Compare deterministic and randomized data structures through empirical evaluation.
CO3	Apply advanced structures to real-world problems like indexing, streaming, and range queries.
CO4	Design optimized data structures for large-scale and high-performance systems.

Course Articulation Matrix:

	PO1	PO2	PO3	PO4	PO5	PO6
CO1	3	1	2	2	2	1
CO2	2	2	2	2	1	1
CO3	2	2	3	3	2	1
CO4	3	1	3	3	3	2

1 - Slightly; 2 - Moderately; 3 – Substantially

Syllabus:

Advanced Balanced Trees: AVL Trees, Red-Black Trees, Order Statistic Trees, Splay Trees, Treaps.
Probabilistic and Randomized Data Structures: Skip Lists, Randomized Binary Search Trees, Cuckoo Hashing, Bloom Filters, Count-Min Sketch.
Trie and String Data Structures: Standard Trie, Compressed Trie (Radix Tree), Patricia Trie, Suffix Trie, Suffix Tree, Suffix Array with LCP, Applications in text indexing and pattern matching.
Advanced Range Query Data Structures: Segment Trees, Sparse Tables, Persistent Segment Trees, Interval Trees.
Heaps and Advanced Priority Structures: Binomial Heaps, Fibonacci Heaps, Pairing Heaps, Applications in graph algorithms.
Graph and Dynamic Data Structures: Disjoint Set Union (Union-Find with path compression and union by rank), Dynamic Connectivity Structures, Heavy-Light Decomposition, Link-Cut Trees.
Geometric and Spatial Data Structures: k-d Trees, Quad Trees, Range Trees, R-Trees and their applications in spatial indexing and multidimensional search.

Learning Resources:

Text Books

1. Thomas H. Cormen, Charles E. Leiserson, Ronald L. Rivest, Clifford Stein, *Introduction to Algorithms*, 3rd Edition, MIT Press, 2009.
2. Mark Allen Weiss, *Data Structures and Algorithm Analysis in C++*, 4th Edition, Pearson Education, 2014.
3. Robert Sedgewick, Kevin Wayne, *Algorithms*, 4th Edition, Addison-Wesley Professional, 2011.

Reference Books

1. Michael T. Goodrich, Roberto Tamassia, Michael H. Goldwasser, *Data Structures and Algorithms in Java*, 6th Edition, Wiley, 2014.
2. Erik Demaine, MIT 6.851 Advanced Data Structures Lecture Notes, Massachusetts Institute of Technology, OpenCourseWare.

Online Resources:

1. MIT OpenCourseWare – Advanced Data Structures
2. Stanford CS166 Lecture Notes
3. CMU 15-451/651 Course Materials
4. GeeksforGeeks Advanced DS Section (for practice)

CS36008	DEEP LEARNING LABORATORY	Credits 0-0-2: 1

Pre-Requisites: **Foundations of AI Models**

Course Outcomes:

At the end of the course, the student will be able to

CO1	Feed Forward Neural Network Learning
CO2	Models for Sequence learning tasks.
CO3	Optimization and Regularization Techniques for Deep Neural Networks
CO4	Fine Tuning Pre-trained Models

Course Articulation Matrix:

	PO1	PO2	PO3	PO4	PO5	PO6
CO1	-	2	3	3	3	3
CO2	-	2	3	3	1	-
CO3	-	-	1	3	2	2
CO4	2	1	1	3	1	-

1 - Slightly; 2 - Moderately; 3 – Substantially

Syllabus:

1. Implement Feed Forward Neural Network with backpropagation from scratch with gradient checking on MNIST. Show the effect of learning rate, optimization function, size of the training dataset.
2. Design a Deep learning model and test it for overfitting and under fitting.Analyze the number of parameters in the model.
3. Compare RNN vs LSTM vs GRU vs BiLSTM for movie review sentiment classification.
4. Train a CNN on CIFAR-10 and generate Grad-CAM heatmaps for model explainability.
5. Handle imbalanced classification using class-weights and focal loss, comparing F1/PR-AUC.
6. Add an attention layer on BiLSTM and visualize attention weights over words.
7. Design a deep learning model for multi-task learning.
8. Perform transfer learning with ResNet/MobileNet and compare frozen vs fine-tuned layers.
9. Fine-tune BERT for sentiment classification and perform detailed error analysis.
10. Fine-tune BERT for sentiment classification for any Indian low resource language (Telugu, Bengali, Marathi. etc.) and perform detailed error analysis.

Text Books

1. Ian Goodfellow, Yoshua Bengio, Aaron Courville, *Deep Learning*, MIT Press, 2016.
2. François Chollet, *Deep Learning with Python*, Manning, 2nd Edition.
3. Aurélien Géron, *Hands-On Machine Learning with Scikit-Learn, Keras, and TensorFlow*, O'Reilly, 3rd Edition

Reference Books

1. Christopher M. Bishop, *Pattern Recognition and Machine Learning*, Springer, 2006 (selected chapters).
2. Dive into Deep Learning (D2L) – *Interactive book* (deep learning theory + labs).
3. Sebastian Raschka et al., *Machine Learning with PyTorch and Scikit-Learn*, Packt.

Online Resources

1. PyTorch Tutorials: https://pytorch.org/tutorials/
2. TensorFlow Tutorials: https://www.tensorflow.org/tutorials
3. Hugging Face Course (Transformers): https://huggingface.co/learn
4. Papers With Code (benchmarks + implementations): https://paperswithcode.com/
5. D2L Book: https://d2l.ai/

CS36095	MINOR PROJECT	Credits 0-0-0: 2

Course Outcomes:

CO-1	Identifying the problem of a research project through literature survey.
CO-2	Analyze the given solution and reproduce it for better experience.

Course Articulation Matrix:

	PO-1	PO-2	PO-3	PO-4	PO-5	PO-6
CO-1	--	1	3	3	3	3
CO-2	-	2	3	3	3	3

1 - Slightly; 2 - Moderately; 3 - Substantially

3rd Semester

CS36089	SEMINAR AND TECHNICAL WRITING	Credits 0-0-0: 2

Pre-Requisites: **None**

Course Outcomes:

CO-1	Analyze the selected topic, organize the content and communication to audience in an effective manner.
CO-2	Practice the learning by self study.

Course Articulation Matrix:

	PO-1	PO-2	PO-3	PO-4	PO-5	PO-6
CO-1	3	2	1	-	-	-
CO-2	3	2	1	-	-	-

1 - Slightly; **2 - Moderately;** **3 - Substantially**

Description:

In Seminar and Technical Writing, every student is expected to prepare a well-organized report based on one / all of the following:

- by attending at least 5 expert lectures/ invited talks/ Seminar/ Popular lectures etc. organized by the institute/any of the departments, ideally in a specific domain or with the same theme.
- prepare a business or marketing plan based on patent search

The student is expected to consolidate the ideas from these lectures/patent searches and may even include material from other sources to strengthen the content of the report. The student should prepare a well-organized report based on the above and present it to the panel constituted by the department, for evaluation.

Evaluation Criteria:

The student will be evaluated by the panel based on the below criteria.

Criteria	Description	Weightages
I	Clarity on the topic	
II	List of lectures attended	
III	Report	
IV	Presentation	
V	Response to questions	

Evaluation Criteria-CO Mapping

CO / Criteria	CO1	CO2
I	X	
II	X	
III		X
IV		
V		

CS36091	SUMMER INTERNSHIP/RESEARCH EXPERIENCE	Credits 0-0-0: 2

Pre-Requisites: None

Course Outcomes:

CO-1	Analyze the practical application of all courses in post graduate curriculum of Computer Science and Engineering.

Course Articulation Matrix:

	PO-1	PO-2	PO-3	PO-4	PO-5	PO-6
CO-1	1	1	3	3	3	3

1 - Slightly; 2 - Moderately; 3 - Substantially

Description:

Every student has to undergo either a Summer Internship / Research project. The summer internship may be undergone in an Industry/Research organization. The research project shall be registered under the guidance of any faculty member in the institute. The student is required to submit a report and present the work before an evaluation committee, nominated by the Head of the Department.

Evaluation Criteria:

The student will be evaluated by the panel based on the below criteria. Weightage for each criterion will be determined by the panel and will be informed to the students.

Criteria	Description	Weightages
I	Relevance of the area of work	15 Marks
II	Performance of the Task	15 Marks
III	Crucial learnings from the work	20 Marks
IV	Report Preparation	25 Marks
V	Presentation and Response to questions	25 Marks

Evaluation Criteria-CO Mapping

CO / Criteria	CO1
I	X
II	X
III	X
IV	X
V	X

CS36093	COMPREHENSIVE VIVA-VOCE	Credits 0-0-0: 2

Pre-Requisites: **None**

Course Outcomes:

CO-1	Comprehend and correlate the understanding of all courses in post graduate curriculum of Computer Science and Engineering.

Course Articulation Matrix:

	PO-1	PO-2	PO-3	PO-4	PO-5	PO-6
CO-1	3	3	3	3	3	3

1 - Slightly; 2 - Moderately; 3 - Substantially

Description:

In Comprehensive viva-voce each student will be evaluated for their overall comprehension of the course work and laboratory training that they have undergone. The students will be expected to answer questions orally, write down simple equations, draw plots, schematics, write simple code etc. as questioned by the panel. Assessment will be done by the panel based on the student response.

CS36097	Dissertation: Part A	Credits 0-0-0: 8

Pre-Requisites: **None**

Course Outcomes:

CO-1	Identify the problem of a research project through literature survey.
CO-2	Analyze the technical feasibility of the project.
CO-3	Propose a solution for the research problem.
CO-4	Analyze, design and implement the proposed solution using software engineering practices.

Course Articulation Matrix:

	PO1	PO2	PO3	PO4	PO5	PO6
CO1	2	3	2	2	2	1
CO2	2	3	2	2	2	1
CO3	2	3	2	2	2	1
CO4	2	3	2	2	2	1

1 - Slightly; 2 - Moderately; 3 - Substantially

Description:

Students are expected to choose real-world contemporary problem and apply the engineering principles learned, to solve the problem through building prototypes or simulations or writing codes or establishing processes/synthesis/correlations etc. The department constituted panel will decide the suitability and worthiness of the project

Dissertation Evaluation:

i. The dissertation shall be submitted as per the schedule given in the academic calendar.
ii. The dissertation supervisor will periodically review the progress of the student and finally give his/her assessment of the work done by the student.
iii. The Dissertation Part – A will be evaluated for 100 marks, with the following weightages:

Sub-component	Weightage
a) Periodic evaluation by Guide	40 marks
b) Mid-term review	20 marks
c) End Semester viva-voce examination	40 marks

The midterm review and the end semester viva-voce examination will be conducted by a committee constituted by the Head of the Department.

Evaluation Criteria

The student will be evaluated by the panel based on the below criteria. Weightage for each criterion will be determined by the panel and will be informed to the students.

Criteria	Description	Weightages
I	Selection of Topic	
II	Literature Survey	
III	Defining the Objectives and Solution Methodology	
IV	Performance of the Task	

Evaluation Criteria - CO Mapping:

CO / Criteria	CO1	CO2	CO3	CO4
I	X			
II		X		
III			X	
IV				X

Refer to MTech. – Regulations for any further information regarding Mid-term review, End Sem evaluation, Template for report preparation and plagiarism.

4th Semester

CS36098	Dissertation: Part B	Credits 0-0-0: 8

Pre-Requisites: **None**

Course Outcomes:

CO-1	Synthesize and apply prior knowledge to designing and implementing solutions to open-ended computational problems while considering multiple realistic constraints.
CO-2	Design and Develop the software with software engineering practices and standards.
CO-3	Analyze Database, Network and Application Design methods.
CO-4	Evaluate the solution through various validation and verification methods.
CO-5	Practice CASE tools for solving software engineering CASE Studies.
CO-6	Analyze professional issues, including ethical, legal and security issues, related to computing projects.

Course Articulation Matrix:

	PO1	PO2	PO3	PO4	PO5	PO6
CO1	3	-	2	2	2	1
CO2	2	-	2	2	2	1
CO3	2	-	2	2	2	1
CO4	-	-	2	2	2	1
CO-5	-	-	2	2	2	1
CO-6	-	-	2	2	2	1

1 - Slightly; 2 - Moderately; 3 - Substantially

Description:

Students are expected to choose real-world contemporary problem and apply the engineering principles learned, to solve the problem through building prototypes or simulations or writing codes or establishing processes/synthesis/correlations etc. The department constituted panel will decide the suitability and worthiness of the project.

Dissertation Evaluation:

i. The dissertation shall be submitted as per the schedule given in the academic calendar.
ii. The dissertation supervisor will periodically review the progress of the student and finally give his/her assessment of the work done by the student.
iii. The Dissertation Part – B will be evaluated for 100 marks, with the following weightages:

Sub-component	**Weightage**
a) Periodic evaluation by Guide	40 marks
b) Mid-term review	20 marks
c) End Semester viva-voce examination	40 marks

The midterm review will be conducted by a committee constituted by the Head of the Department. The end semester examination will be conducted by an External Examiner along with the evaluation committee constituted by the Head of the Department.

Evaluation Criteria:

The student will be evaluated by the panel based on the below criteria as a continuation of Dissertation Part A. Weightage for each criterion will be determined by the panel and will be informed to the students.

Task	Description	Weightages (%)
IV	Performance of the Task	20
V	Dissertation Preparation	25
VI	Review (Presentation & Understanding)	30
VII	Viva-Voce	15
VIII	Publications /Possibility of publication	10

Evaluation Criteria-CO Mapping:

CO / Criteria	CO1	CO2	CO3	CO4	CO5	CO6
IV	X	X				
V				X		
VI			X			
VII				X		
VIII					X	X

Refer to M.Tech. – Regulations for any further information regarding Mid-term review, End Sem evaluation, Template for report preparation and plagiarism.

Professional Electives

CS36021	RECOMMENDER SYSTEMS	Credits 3-0-0: 3

Pre-Requisites: **Machine Learning**

Course Outcomes:

At the end of the course, the student will be able to

CO1	Introduce the fundamental concepts, architectures, and applications of recommender systems in real-world personalization scenarios.
CO2	Understand and analyze collaborative filtering, content-based, and hybrid recommendation techniques.
CO3	To address practical challenges such as data sparsity, scalability, cold start, and robustness in large-scale recommendation systems
CO4	Explore advanced and research-oriented topics, including context-aware, sequential, and online recommender systems.

Course Articulation Matrix:

	PO1	PO2	PO3	PO4	PO5	PO6
CO1	2	1	2	2	–	–
CO2	2	1	3	3	–	–
CO3	3	1	3	3	2	–
CO4	2	2	2	2	–	–

1 - Slightly; 2 - Moderately; 3 – Substantially

Syllabus:

Introduction to Recommender Systems: Overview and motivation of recommender systems, historical evolution and applications, personalization and information overload, types of recommender systems, user–item interaction data, explicit and implicit feedback, challenges such as sparsity, scalability, and cold start, and real-world recommendation scenarios.

Collaborative Filtering: User-based and item-based collaborative filtering approaches, similarity measures, rating prediction and ranking, handling sparsity and scalability, memory-based vs model-based methods, advantages and limitations of neighborhood models, and practical considerations in large-scale systems. Latent factor models, matrix factorization techniques, optimization objectives, regularization, probabilistic and Bayesian interpretations, handling implicit feedback, temporal dynamics, scalability issues, and connections to statistical learning and dimensionality reduction.

Content-Based and Knowledge-Based Recommendation: Content representation and feature extraction, similarity computation, learning user profiles, classification and regression approaches for content-based recommendation, knowledge-based and constraint-based recommenders, comparison with collaborative filtering, and hybrid recommendation strategies.

Statistical and Learning-Based Recommendation Models: Statistical foundations of recommender systems, high-dimensional prediction under sparsity, regression and classification models for recommendation, probabilistic modeling, uncertainty estimation, context-aware recommendation, and sequential decision-making perspectives.

Evaluation of Recommender Systems: Offline evaluation metrics for rating prediction and ranking, top-N recommendation evaluation, online evaluation and A/B testing, user engagement metrics, biases in evaluation, reproducibility, and limitations of standard evaluation protocols.

Advanced Topics in Recommender Systems: Context-aware and session-based recommendation, social and trust-aware recommender systems, time-aware and sequential models, robustness to shilling and adversarial attacks, diversity, novelty, and serendipity in recommendation, ethical and fairness considerations, Reinforcement Learning for Recommender Systems, Explainable Recommender Systems.

Learning Resources:

Textbooks:

1. Charu C. Aggarwal, "*Recommender Systems"*, Springer, 2016,
2. Deepak K. Agarwal & Bee-Chung Chen, "*Statistical Methods for Recommender Systems",* Cambridge University Press, *2016.*

CS36023	Advanced Machine Learning	Credits 3-0-0: 3

Pre-Requisites: **None**

Course Outcomes:

At the end of the course, the student will be able to

CO1	Understand the mathematical perspective of machine learning algorithms
CO2	Design and implement ensemble architectures such as XGBoost, LightGBM, and stacking for complex datasets
CO3	Understand Fuzzy logic and Rough set, and analyze large datasets
CO4	Design and deploy HMM-based solutions for real-world tasks

Course Articulation Matrix:

	PO1	PO2	PO3	PO4	PO5	PO6
CO1	1	-	3	-	2	3
CO2	3	2	-	1	3	1
CO3	2	-	-	3	-	1
CO4	-	1	-	2	1	2

1 - Slightly; 2 - Moderately; 3 – Substantially

Syllabus:

Advanced Supervised Learning: Introduction to Ensemble Learning, Bagging Techniques, Boosting Methods, Stacking and Blending, Advanced Ensemble Techniques, Gradient Boosting Machines- GBM, XGBoost, LightGBM, Kernel methods & Support Vector Machines, Probabilistic models -Naïve Bayes, Bayesian Networks, Calibration & uncertainty estimation, Ensemble Methods for Different Tasks.

Bayesian statistics-Bayesian model selection, Priors, Hierarchical Bayes, Empirical Bayes, Bayesian decision theory.

Fuzzy Logic and Rough Sets, Fuzzy Sets and Membership Functions, Fuzzy Inference Systems, Introduction to Rough Set Theory, Hybrid Approaches.

Markov and hidden Markov models-Markov models, Hidden Markov models, Inference in HMMs, Learning for HMMs, Generalizations of HMMs.

Learning Resources:

Text Books

1. Marc Peter Deisenroth, A. Aldo Faisal, and Cheng Soon Ong. *“Mathematics for Machine Learning”.* Cambridge University Press, 2020
2. Kevin P. Murphy, “Machine Learning: A Probabilistic Perspective”. MIT Press, 2012.
3. Gautam Kunapuli, “Ensemble Methods for Machine Learning”, 1st edition, Pearson, 2023
4. George J. Klir, Bo Yauan, “Fuzzy sets and Fuzzy Logic”, Prentice Hall, 1995

CS36025	COMPUTER VISION	Credits 3-0-0: 3

Pre-Requisites: **None**

Course Outcomes:

CO1	Classify Image representations..
CO2	Apply Image transformation methods.
CO3	Design face detection and recognition algorithms
CO4	Recover the information, knowledge about the objects in the scene..

Course Articulation Matrix:

	PO1	PO2	PO3	PO4	PO5	PO6
CO1	3	1	2	1	-	-
CO2	-	2	3	-	-	-
CO3	2	-		3	3	3
CO4	-	-	3	3	3	3

1 - Slightly; 2 - Moderately; 3 – Substantially

Syllabus:

The image model and acquisition, image shape, sampling, intensity images, color images, range images, image capture, scanners. Statistical and spatial operations, Gray level transformations, histogram equalization, multi image operations. Spatially dependent transformations, templates and convolution, window operations, directional smoothing, other smoothing techniques. Segmentation and Edge detection, region operations, Basic edge detection, second order detection, crack edge detection, edge following, gradient operators, compass & Laplace operators. Morphological and other area operations, basic morphological operations, opening and closing operations, area operations, morphological transformations. Image compression: Types and requirements, statistical compression, spatial compression, contour coding, quantizing compression. Representation and Description, Object Recognition, 3-D vision and Geometry, Digital Watermarking. Texture Analysis.

Learning Resources:

Text Books

1. D. A. Forsyth, J. Ponce, "*Computer Vision: A Modern Approach*", Second edition, Pearson, 2012.
2. Milan Soanka, Vaclav Hlavac and Roger Boyle, "*Digital Image Processing and Computer Vision", Fifth edition,* Cengage Learning, 2015
3. R.C. Gonzalez and R.E. Woods, "*Digital Image Processing*", 4th edition, Pearson, 2018.

CS36027	BIG DATA ANALYTICS	Credits 3-0-0: 3

Pre-Requisites: None

Course Outcomes:

CO-1	Understand big data challenges in different domains including social media, transportation, finance and medicine.
CO-2	Analyze the capability of No-SQL systems.
CO-3	Apply machine learning algorithms for data analytics.
CO-4	Analyze MAP-REDUCE programming model for better optimization.

Course Articulation Matrix:

	PO1	PO2	PO3	PO4	PO5	PO6
CO1	1	-	-	3	3	3
CO2	3	-	2	1	3	1
CO3	1	3	-	2	2	2
CO4	2	2	3	1	3	2

1 - Slightly; 2 - Moderately; 3 - Substantially

Syllabus:

Overview of Big Data Analytics: Big Data Characteristics, Big Data Challenges, Big Data Analytics system Architecture, Big Data Storage Technologies, Big Data Analytics Life Cycle; No SQL databases: Introduction to No SQL databases, Types of No SQL databases, CAP Theorem, Replication & Shading, MongoDB basics; Overview of Hadoop, Hadoop Ecosystem, Reading and Writing Files on Hadoop Distributed File System. MapReduce basics, MapReduce algorithm design. Graph Algorithms, Data Mining with Big Data, Stream Data Processing, Stream Data Processing System and Tools, Overview of Spark, Spark SQL-Mllib-GraphX, Apache Kafka, Big Data Analytics in Industry Verticals, Operationalizing Basic Data Analytic Methods Using R, Analytics for Unstructured Data.

Learning Resources:

Text Books:

1. Jure Leskovec, Anand Rajaraman, J D Ullman, "*Mining Massive Datasets"*, Cambridge University Press, 3rd Edition, 2020
2. Jimmy Lin and Chris Dyer, *Data Intensive Text Processing with Map Reduce*, Morgan & Claypool Publishers,1st Edition, 2010.

Reference Books:

1. Bill Franks, *Taming The Big Data Tidal Wave*, 1st Edition, Wiley, 2012.
2. Johannes Ledolter, *Data Mining and Business Analytics* with R, Wiley, 2013.

CS36029	IoT DATA ANALYTICS	Credits 3-0-0: 3

Pre-Requisites: **Computer Networks Protocols and ML**

Course Outcomes:

At the end of the course, the student will be able to

CO1	Understand how IoT-generated data can be processed and analyzed to support informed decision-making and create measurable business value
CO2	Analyse the IoT data life cycle
CO3	Assess IoT data in terms of data volume, velocity, variety, real-time processing requirements, and architectural considerations.
CO4	Implement analytics solutions using modern tools.

Course Articulation Matrix:

	PO1	PO2	PO3	PO4	PO5	PO6
CO1	1	-	-	-	2	-
CO2	-	2	1	3	3	3
CO3	2	-	2	3	3	3
CO4	-	3	3	3	3	3

1 - Slightly; 2 - Moderately; 3 – Substantially

Syllabus:

Introduction to Data Analytics & Tools: Overview of IoT architecture (Perception, Network, Application layers) Characteristics of IoT data (Volume, Velocity, Variety, Veracity) Introduction to Data Analytics,
Types of Analytics: Descriptive, Diagnostic, Predictive, Prescriptive, IoT Data Lifecycle, IoT analytics use cases (Smart cities, healthcare, Industry 4.0, smart grids).
IoT Data Acquisition and Management: Sensor data generation and formats (JSON, CSV, XML), IoT communication protocols (MQTT, CoAP, HTTP), Data ingestion techniques, Edge vs Cloud data collection, Data storage models.
Data Preprocessing and Cleaning, IoT Data Analytics Techniques-Statistical analysis for IoT data, Time-series analysis, Machine Learning for IoT, Anomaly detection, Predictive maintenance models, Stream analytics vs batch analytics.
Visualization and Business Insights: Data visualization principles, Dashboards and reporting,
KPI identification, Translating analytics into business decisions, Case studies (Industry 4.0, Smart agriculture, Healthcare IoT).

Learning Resources:

Text Books:

1. Harry G. Perros, "*An Introduction to IoT Analytic",* Chapman & Hall / CRC Press, First edition, 2021.
2. Andrew Minteer, "*Analytics for the Internet of Things (IoT): Intelligent analytics for your intelligent devices"*, Packt Publishing, 1st Edition, 2017

CS36031	ADVANCED NATURAL LANGUAGE PROCESSING	Credits 3-0-0: 3

Pre-Requisites: **Nil**

Course Outcomes:

At the end of the course, the student will be able to

CO1	Understand approaches to syntax and semantics in NLP.
CO2	Understand approaches to discourse, generation, dialogue and summarization within NLP
CO3	Understand current methods for statistical approaches to machine translation.
CO4	Understand machine learning techniques used in NLP, including hidden Markov models and probabilistic context-free grammars, clustering and unsupervised methods, log-linear and discriminative models, and the EM algorithm as applied within NLP.

Course Articulation Matrix:

	PO1	PO2	PO3	PO4	PO5	PO6
CO1	1	2	2	2	1	2
CO2	1	2	3	2	2	2
CO3	3	3	3	3	3	3
CO4	2	3	3	2	2	3

1 - Slightly; 2 - Moderately; 3 – Substaintially;

Syllabus:

Foundations & Language Modeling - Introduction to NLP, Statistical NLP paradigm, Generative vs Discriminative models, Structured prediction, Evaluation metrics, Estimation Techniques - Maximum Likelihood Estimation, Parameter estimation under sparsity, N-gram models, RNN, Evaluation & Error Analysis- Intrinsic vs Extrinsic evaluation

Structured Prediction & Parsing: Probabilistic Grammars **-** CFG review, PCFGs, Tree probabilities, CKY algorithm, Lexicalized & Dependency Parsing, Sequence Tagging Models- Hidden Markov Models (HMM), Forward–Backward algorithm, Viterbi decoding, conditional random field (CRF).

Similarity, Representation & Clustering - Probabilistic Similarity Measures, KL Divergence, Jensen-Shannon, Divergence, Cosine similarity, Distributional Semantics- Word embeddings (Word2Vec, GloVe), Contextual similarity concepts, Clustering Techniques in NLP - K-means, Hierarchical clustering, Topic Modeling- Latent Dirichlet Allocation, Autoencoders.

Machine Translation & Dialogue Systems- Statistical Machine Translation - Noisy Channel Model, Neural Machine Translation - Encoder–Decoder models, Attention mechanism, Evaluation of MT Systems - BLEU score, Human evaluation, Error categories, Dialogue Systems - Intent detection, Slot filling, dialogue stage tracking.

Discourse, Generation & Advanced Topics - Discourse Processing- Coreference & Anaphora Resolution, Mention-pair models, Text Summarization & Advanced Topics- Extractive vs Abstractive summarization, ROUGE evaluation, introduction to large language models.

Learning Resources:

Text Books

1. Jurafsky, David, and James H. Martin, "*Speech and Language Processing: An Introduction to Natural Language Processing, Computational Linguistics and Speech Recognition" ,* 3rd Edition, 2026
2. Manning, Christopher, and Hinrich Schutze, " *Foundations of statistical natural language processing*", MIT press, 2006.
3. Deng, Li, and Yang Liu, eds, "*Deep learning in natural language processing*", Springer, First edition, 2018.

Reference Books

1. Deng, Li, and Yang Liu, eds. "*Deep learning in natural language processing".* Springer, 2018.
2. Goldberg, Yoav, and Graeme Hirst. "Neural network methods for natural language processing", 2018.

CS36033	SMART EDGE COMPUTING	Credits 3-0-0: 3

Pre-Requisites: **None**

Course Outcomes:

At the end of the course, the student will be able to

CO1	Understand cloud limitations and edge evolution
CO2	Apply distributed systems in edge computing
CO3	Design edge pipelines using containers
CO4	Develop edge analytics and ML models

Course Articulation Matrix:

	PO1	PO2	PO3	PO4	PO5	PO6
CO1	3	-	1	2	3	3
CO2	1	2	-	-	-	1
CO3	-	3	2	2	1	3
CO4	-	-	3	3	3	3

1 - Slightly; 2 - Moderately; 3 – Substantially

Syllabus:

Introduction to Cloud and its limitations to support low latency and RTT. From Cloud to Edge computing: Waves of innovation, Introduction to Edge Computing Architectures, Edge Computing to support User Applications (5G-Slicing, self-driving cars and more), Concepts of distributed systems in edge computing such as time ordering and clock synchronization, distributed snapshot, etc., Introduction to Edge Data Center, Lightweight Edge Clouds and its services provided by different service providers, Introduction to docker container and Kubernetes in edge computing. Design of edge storage systems like key-value stores, Introduction to MQTT and Kafka for end-to-end edge pipeline. Edge analytics topologies for M2M and WSN network (MQTT), Use cases of machine learning for edge sensor data in predictive maintenance, image classifier and self-driving cars. Deep Learning On-Device inference at the edge to support latency-based application.

Learning Resources:

Text Books

1. Buyya R., Srirama S.N., "*Fog and Edge Computing*", Wiley, 2019.
2. Taheri J. & Deng S. (eds.): "*Edge Computing: Models, technologies and applications*", IET, 2020

Reference Books

1. Sabella D., Reznik A., Frazao R., "*Multi-access Edge Computing in Action*", 1st edition, Kindle, 2019
2. Al-Turjman F. (ed.): "*Edge Computing: from hype to reality*", Springer, 2019.

Online Resources:

https://onlinecourses.nptel.ac.in/noc25_cs28/preview

CS36035	AI IN FINTECH	Credits 3-0-0: 3

Pre-Requisites: **Deep Learning**

Course Outcomes:

At the end of the course, the student will be able to

CO1	Understand the architecture of FinTech systems and the role of Artificial Intelligence in modern financial services
CO2	Apply machine learning techniques to analyze financial datasets for problems such as credit scoring, risk prediction, and customer analytics
CO3	Design and develop AI models including deep learning and NLP techniques for fraud detection, financial forecasting, and financial text analytics.
CO4	Evaluate ethical, regulatory, bias, and explainable AI aspects in AI-driven financial decision systems

Course Articulation Matrix:

	PO1	PO2	PO3	PO4	PO5	PO6
CO1	2	1	1	2	1	1
CO2	3	1	2	3	2	1
CO3	3	1	3	3	2	2
CO4	2	2	1	1	1	3

1 - Slightly; 2 - Moderately; 3 – Substantially

Syllabus:

Introduction to FinTech and AI : Evolution of financial technology, FinTech ecosystem, digital banking architecture, financial data sources, AI applications in banking and finance.

Machine Learning for Financial Analytics: Financial data characteristics, feature engineering, supervised and unsupervised learning models, evaluation metrics. Applications: credit scoring, loan default prediction, customer segmentation.

Fraud Detection and Financial Security: Types of financial fraud, anomaly detection, graph-based fraud detection, AI in Anti-Money Laundering (AML). Techniques include Random Forest, Isolation Forest, Autoencoders, and Graph Neural Networks.

Deep Learning for Financial Forecasting: Neural networks, RNN, LSTM, GRU, transformer models for financial time-series prediction. Applications in stock forecasting, portfolio optimization and risk prediction.

Natural Language Processing for Financial Analytics: Financial text mining, sentiment analysis, named entity recognition, transformer models like BERT and FinBERT.

Reinforcement Learning for Financial Decision Systems: Markov Decision Processes, Q-learning, Deep Q Networks, RL for trading and portfolio management.

Explainable AI and Responsible AI in Finance: Explainable AI techniques such as SHAP and LIME, fairness, bias mitigation, regulatory and ethical considerations.

Emerging Trends in AI-driven FinTech: AI for decentralized finance (DeFi), blockchain analytics, quantum computing in finance, and financial inclusion.

Learning Resources:

Text Books

1. Yves Hilpisch, "*Artificial Intelligence in Finance*", O'Reilly Media, 2020
2. Marcos López de Prado, "*Advances in Financial Machine Learning"*, Wiley, 2018.
3. Matthew F. Dixon, Igor Halperin, Paul Bilokon, "*Machine Learning in Finance: From Theory to Practice",* Springer, 2020
4. Stefan Jansen, "*Machine Learning for Algorithmic Trading: Predictive Models to Extract Signals from Market and Alternative Data for Systematic Trading Strategy",* Packt Publishing, 2020.
5. Yves Hilpisch, "*Python for Finance: Mastering Data-Driven Finance"*, Second Edition, O'Reilly Media, 2018.

CS36037	AI FOR SOFTWARE ENGINEERING	Credits 3-0-0: 3

Pre-Requisites: **Artificial Intelligence, Software Engineering**

Course Outcomes:

At the end of the course, the student will be able to

CO1	Describe the role of AI in different phases of the software development lifecycle.
CO2	Apply machine learning techniques for software requirement analysis and design.
CO3	Utilize AI tools and LLMs for code generation, refactoring, and debugging
CO4	Implement automated testing strategies using AI to enhance software quality.

Course Articulation Matrix:

	PO1	PO2	PO3	PO4	PO5	PO6
CO1	1	-	3	2	3	3
CO2	2	-	2	3	3	3
CO3	1	3	-	-	2	2
CO4	-	2	-	2	-	3

1 - Slightly; 2 - Moderately; 3 – Substantially

Syllabus:

Foundations of AI and Programming: Python for AI: Data structures, loops, functions, and libraries (NumPy, Pandas), Mathematics for AI: Linear algebra, calculus, and probability for understanding algorithms. Machine Learning (ML) Basics: Supervised/unsupervised learning, model evaluation.

Generative AI for Software Development: LLMs in Development: Using models like GPT-4 for code generation, debugging, and refactoring. Prompt Engineering for Developers: Designing prompts for generating accurate code, documentation, and test cases. AI-Powered IDEs: Practical experience with tools like GitHub Copilot, Cursor, and VS Code extensions.

AI-Integrated Development Workflow: Automated Code Review: Using AI to analyze, improve, and secure codebases. Software Design with AI: Assisting in architectural decisions and system design using GenAI. Version Control & DevOps: Integrating AI into CI/CD pipelines and GitHub workflows.

Advanced AI Engineering Topics: Large Language Models (LLMs): Understanding Transformer architectures. Retrieval Augmented Generation (RAG): Connecting LLMs to custom data sources for enhanced context. AI Agents: Building autonomous agents for software development tasks.

Quality, Security, and Ethics: Mismatch Detection: Identifying, evaluating, and fixing errors in AI-generated code. Security Risks: Addressing prompt injections, hallucinations, and insecure code generation. Responsible AI: Ensuring fairness and transparency in AI-driven tools.

Learning Resources:

Text Books

1. Yang, Yanming, Xin Xia, David Lo, and John Grundy. "*A survey on deep learning for software engineering.*" ACM Computing Surveys (CSUR) 54, no. 10s (2022): 1-73.
2. Watson, Cody, Nathan Cooper, David Nader Palacio, Kevin Moran, and Denys Poshyvanyk. "*A systematic literature review on the use of deep learning in software engineering research."* ACM Transactions on Software Engineering and Methodology (TOSEM) 31,

Reference Books

1. Zan, Daoguang, Bei Chen, Fengji Zhang, Dianjie Lu, Bingchao Wu, Bei Guan, Wang Yongji, and Jian-Guang Lou. "*Large language models meet NL2Code: A survey." In Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics*, 2023.
2. Wang, Junjie, Yuchao Huang, Chunyang Chen, Zhe Liu, Song Wang, and Qing Wang. "*Software testing with large language model: Survey, landscape, and vision*", 2023. https://arxiv.org/abs/2307.07221

CS36039	INTELLIGENT DATA NETWORKS	Credits 3-0-0: 3

Pre-Requisites: **Computer Networks, Data Structures, Probability and Statistics**

Course Outcomes:

At the end of the course, the student will be able to

CO1	Understand intelligent networking paradigms and data-driven network architectures.
CO2	Apply machine learning techniques for traffic analysis, congestion control, and anomaly detection.
CO3	Design software-defined and programmable networks using AI-based approaches.
CO4	Analyze security, optimization, and performance challenges in intelligent data networks.

Course Articulation Matrix:

	PO1	PO2	PO3	PO4	PO5	PO6
CO1	3	2	2	1	2	1
CO2	2	2	3	2	2	1
CO3	2	2	3	3	3	1
CO4	2	2	2	2	2	3

1 - Slightly; 2 - Moderately; 3 – Substantially

Syllabus:

Introduction to Intelligent Data Networks: Evolution from traditional networks to intelligent and data-driven networks, Network virtualization, Software Defined Networking (SDN), Network Function Virtualization (NFV), Programmable data planes.

Machine Learning for Networking: Supervised and unsupervised learning for traffic classification, congestion prediction, anomaly detection, reinforcement learning for routing and congestion control, Deep learning in network management.

Software Defined and Cognitive Networks: SDN architecture, OpenFlow, Control and data plane separation, Network slicing, Intent-based networking, Cognitive networking principles.

Network Traffic Analytics and Big Data: Network data collection, Flow monitoring, Network telemetry, Big data frameworks for network analytics, Real-time traffic analysis and QoS optimization.

Security in Intelligent Networks: AI-based intrusion detection systems, DDoS detection and mitigation, Blockchain for network security, Privacy-preserving networking, Trust management in intelligent networks.

Edge, IoT and 5G/6G Intelligent Networks: Edge computing in intelligent networks, AI-driven resource allocation in 5G/6G, IoT traffic management, Network automation and orchestration.

Emerging Trends: Self-organizing networks, Autonomous networks, Digital twins for networking, AI-native network architectures.

Learning Resources:

Text Books

1. William Stallings, "*Foundations of Modern Networking: SDN, NFV, QoE, IoT, and Cloud"*, Addison-Wesley, 2016.
2. Fei Hu, "*Machine Learning for Future Wireless Communications"*, Springer, 2020.
3. Mohamed G. Gouda, "*Elements of Network Protocol Design*", Wiley, 1998.

Reference Books

1. DN and NFV related RFC documents and Open Networking Foundation (ONF) publications.
2. Mehdi Bennis, Mérouane Debbah, H. Vincent Poor, *Ultra-Reliable and Low-Latency Wireless Communication*, Cambridge University Press, 2018.
3. Recent IEEE Communications Surveys & Tutorials and IEEE/ACM Transactions on Networking papers on AI-driven networking.

Online Resources:

1. MIT OpenCourseWare – Computer Networks

https://ocw.mit.edu/

2. Stanford CS144: Introduction to Computer Networking
   https://cs144.github.io/
3. CMU 15-441/641 Computer Networks Course Materials
   https://www.cs.cmu.edu/
4. Open Networking Foundation (ONF) – SDN Resources
   https://opennetworking.org/
5. Cisco DevNet – SDN and Network Automation Learning
   https://developer.cisco.com/
6. IEEE Xplore – AI in Networking Papers
   https://ieeexplore.ieee.org
7. ITU AI for 5G/6G and Intelligent Networking Reports
   https://www.itu.int/

CS36041	DATA LAKES AND DATA FABRIC	Credits 3-0-0: 3

Pre-Requisites: **Artificial Intelligence, Software Engineering**

Course Outcomes:

At the end of the course, the student will be able to

CO1	Describe the role of AI in different phases of the software development lifecycle.
CO2	Apply machine learning techniques for software requirement analysis and design.
CO3	Utilize AI tools and LLMs for code generation, refactoring, and debugging
CO4	Implement automated testing strategies using AI to enhance software quality.

Course Articulation Matrix:

	PO1	PO2	PO3	PO4	PO5	PO6
CO1	-	-	2	3	3	3
CO2	1	2	-	2	1	3
CO3	-	-	-	3	3	3
CO4	1	-	3	3	3	3

1 - Slightly; 2 - Moderately; 3 – Substantially

Syllabus:

Database Systems Foundations - Evolution of Database Technologies, Database Modeling and Design - Conceptual and Logical Data Models, ER Diagrams, Graph-Based Models, and Ontologies, Translating Models into Optimized Schema Designs, Normalization and Denormalization - Trade-offs Between Normalization and Performance, Denormalization for Analytics and Query Performance; Storage and Access Methods - Row-Based vs. Columnar Storage Models, Vectorized Storage and Processing, Indexing Techniques; Query Processing and Transaction Management and Logging - Data Flow Restart Strategies for Resilience; Schema-On-Read vs Schema-on-Write.

Distributed Database System – Distributed Database architecture - Design Considerations for Distributed Databases – Client Server Architecture, Peer-to-Peer, Federated, Multi-database, Shared-Nothing, Storage-Compute Separation; Levels of Transparency in Distributed Databases, Data Fragmentation.

Data Warehouse Design – Dimensional modeling, ETL pipelines, OLAP, Materialized Views; Data Lakes – Data Sources, Metadata Management, Metadata Models, Data Lakehouse – Optimising the processes. Data Storage/File formats and Table formats: HUDI, ICEBERG, PARQUET, ORC, AVRO, etc., Data Modeling for Data Lakes, Data Mesh- Data product.

Data Fabric - Components of Data Fabric Systems, Metadata Modeling and Management, Technologies and Systems, Semantic Layer, Data Pipelines – Ingest, transformation, self-healing, Quantum Pipelines, Data Lineage, Data Governance. Performance Optimization -Materialized, use cases, and applications.

AI for Data Fabric and Data Fabric for AI – Machine Learning pipelines, Agent orchestration, Automated Data Pipeline Creation using LLMs, Metadata Discovery.

Learning Resources:

Text Books:

1. E. Hechler, M. Weihrauch, Y. C. Wu*, "Data Fabric and Data Mesh Approaches with AI: A Guide to AI-based Data Cataloging, Governance, Integration, Orchestration, and Consumption",* 2023
2. R. Barik, *Data Fabric primer*, A white paper, 2022

CS36043	HPC SYSTEMS FOR AI	Credits 3-0-0: 3

Pre-Requisites: **Computer Organization and Architecture, basic knowledge of C/Python**

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

CO-1	Analyze for the performance of AI algorithms on HPC systems
CO-2	Generate parallel programs for real world problems using Cuda library
CO-3	Compare the performance of different algorithms for the numerical and data processing problems on GPGPUs and suggest methods for improving the performance.

Course Articulation Matrix:

	PO1	PO2	PO3	PO4	PO5	PO6
CO1	2	2	2	3	2	1
CO2	2	3	2	3	2	1
CO3	2	3	3	2	2	1
CO4	2	2	2	3	2	1

1 - Slightly; 2 - Moderately; 3 - Substantially

Syllabus:

Introduction: Foundations of Parallel Processing (SIMD, MIMD, SIMT, SPMD), HPC Hardware Overview: CPUs, GPU acceleration, high-speed interconnects (InfiniBand), The intersection of HPC and AI/ML/DL. HPC Ecosystem for AI – AI Frameworks: PyTorch, TensorFlow, MXNet on HPC, Job Scheduling: Batch processing, cluster management using SLURM, Storage Systems: High-performance parallel file systems and I/O for machine learning.

Advanced Parallel programming techniques: Accelerators, GPUs, CUDA, Overview of CUDA C; threads, blocks and grids, warps, different GPU memories, CUDA Kernels, Operations in Deep Learning and their implementation on CUDA, Designing High Performance Systems for Accelerated Machine Learning and Deep Learning Workloads. Deep Learning Software: Setting up Application Environment for Deep Learning and HPC workloads using Container, Platform like Docker. Precision scaling, in-memory computing, hyperdimensional computing, architectural modifications.

Learning Resources:

Text Books:

1. Jason Sanders, Edward Kandrot, "*CUDA By Example – An Introduction to General-Purpose GPU Programming",* Addison Wesley, 2011
2. Adrian Rosebrock, "*Deep Learning for Computer Vision with Python*", E-Book, 1 st ed. September 2017.
3. Online Web Materials

CS36045	AGENTIC AI	Credits 3-0-0: 3

Pre-Requisites: **Machine Learning**

Course Outcomes:

At the end of the course, the student will be able to

CO1	Understand the architecture of LLM-based agents.
CO2	Design autonomous AI agents with planning, memory, and tools
CO3	Implement agentic reinforcement learning and decision-making**.**
CO4	Evaluate agents using modern benchmarks and metrics**.**

Course Articulation Matrix:

	PO1	PO2	PO3	PO4	PO5	PO6
CO1	3	1	3	3	2	1
CO2	2	1	2	2	2	–
CO3	2	2	3	3	2	1
CO4	3	1	3	3	2	–

1 - Slightly; 2 - Moderately; 3 – Substantially

Syllabus:

Introduction, Evolution of AI agents, emergence of large language model agents, agentic AI paradigm, core components of agentic systems. Agent architectures, ReAct framework, plan-and-execute agents, Reflexion architecture, autonomous reasoning–action loops.

Reasoning mechanisms in agents, chain-of-thought reasoning, tree-of-thought reasoning, graph-of-thought reasoning. Planning and task decomposition, hierarchical planning, goal-directed agents, iterative decision-making.

Memory systems in agentic AI, short-term memory, long-term memory, episodic memory, semantic memory. Knowledge integration, retrieval-augmented generation, vector databases, knowledge graph integration. Tool usage in agents, API invocation, function calling, external tools, environment interaction.

Multi-agent systems, collaborative agents, agent communication protocols, coordination and negotiation.

Reinforcement learning for agentic systems, policy learning, reward optimization, decision making under uncertainty.

Agent frameworks and development platforms, orchestration pipelines, scalable agent workflows.

Evaluation of agentic AI systems, Safety and alignment in agentic AI, Applications and emerging directions, autonomous research agents, software engineering agents, multimodal and embodied agents.

Learning Resources

Text Books:

1. Stuart J. Russell and Peter Norvig, "*Artificial Intelligence: A Modern Approach*", 4th edition, 2020.
2. Micheal Lanham, "AI Agents in Action", Second edition, 2026.
3. Nicole Koenigstein, "*AI Agents: The Definitive Guide*", 1st Edition, 2026.
4. Nick Bostrom, "*Superintelligence: Paths, Dangers, Strategies",* 1st Edition, 2014.
5. Relevant state-of-the-art research and survey papers

CS36047	RESPONSIBLE AI	Credits 3-0-0: 3

Pre-Requisites: **Machine Learning**

Course Outcomes:

At the end of the course, the student will be able to

CO1	Understand the ethical, societal, and legal implications of artificial intelligence systems.
CO2	Analyze fairness, bias, transparency, accountability, and explainability in AI models.
CO3	Evaluate privacy, safety, and risk factors in real-world AI applications.
CO4	critically assess AI systems using case studies, governance mechanisms, and regulatory perspectives.

Course Articulation Matrix:

	PO1	PO2	PO3	PO4	PO5	PO6
CO1	1	2	1	–	–	3
CO2	1	1	2	–	–	3
CO3	1	2	2	1	1	3
CO4	1	1	3	1	1	3

1 - Slightly; 2 - Moderately; 3 – Substantially

Syllabus:

Introduction to Responsible Artificial Intelligence, motivation and necessity of Responsible AI, historical failures and real-world harms caused by AI systems, ethical challenges specific to data-driven and autonomous systems, distinction between AI ethics, Responsible AI, and AI governance.

Ethical foundations for Artificial Intelligence, including major ethical theories such as utilitarianism, deontological ethics, virtue ethics, and human-rights-based approaches, moral agency and autonomy of AI systems, responsibility and accountability in automated decision-making, human-in-the-loop, human-on-the-loop, and human-out-of-the-loop AI systems.

Principles and frameworks for Responsible AI, encompassing fairness, accountability, transparency, explainability, privacy, safety, robustness, and sustainability, value-sensitive and ethics-by-design approaches, lifecycle perspective of AI systems, and the Accountability–Responsibility–Transparency (ART) framework for responsible AI development and deployment.

Bias, fairness, and discrimination in AI systems, sources of algorithmic bias arising from data, models, and deployment contexts, societal and historical biases reflected in datasets, group fairness and individual fairness notions, fairness metrics and evaluation criteria, bias detection, mitigation strategies, and trade-offs between model performance and fairness.

Transparency, interpretability, and explainability in AI, including black-box versus interpretable models, intrinsic and post-hoc explanation techniques, explainability in machine learning and deep learning systems, trust, user understanding, and limitations of explainability methods in high-stakes applications.

Privacy, security, and accountability in AI systems, privacy risks in data collection and machine learning pipelines, data governance, consent, anonymization, and secure model deployment, accountability mechanisms across the AI lifecycle, responsibility gaps, auditing, documentation, and traceability of AI systems.

Responsible AI in practice, covering ethical and societal implications of AI applications in domains such as healthcare, finance, surveillance, education, and public services, AI risk assessment and impact analysis, governance structures, regulatory considerations, case studies of AI failures and best practices, and future challenges and research directions in Responsible Artificial Intelligence.

Learning Resources

Text Books:

1. Paula Boddington, *“AI Ethics: A Textbook”,* First edition, Springer, 2023
2. Virginia Dignum, *“Responsible Artificial Intelligence”,* First edition, Springer, 2019
3. Sray Agarwal & Shashin Mishra, “*Responsible AI: Implementing Ethical and Unbiased Algorithms”,* Packt Publishing, first edition, 2024.

CS 36022	QUANTUM MACHINE LEARNING	Credits 3-0-0: 3

Pre-Requisites: **Machine Learning**

Course Outcomes:

At the end of the course, the student will be able to

CO1	Understand the fundamentals of quantum computing and quantum information
CO2	Analyse principles of machine learning and their adaptation to quantum systems.
CO3	Explore quantum algorithms for machine learning tasks**.**
CO4	Implement quantum machine learning models using quantum frameworks

Course Articulation Matrix:

	PO1	PO2	PO3	PO4	PO5	PO6
CO1	-	-	3	3	3	3
CO2	1	2	3	3	3	1
CO3	2	-	1	1	3	3
CO4	-	1	3	3	3	3

1 - Slightly; 2 - Moderately; 3 – Substantially

Syllabus:

Fundamentals of Quantum Computing: Introduction to Quantum Computing, Classical vs Quantum Computing, Qubits and Quantum States, Superposition and Quantum Entanglement, Quantum Gates and Quantum Circuits, Quantum Measurement and Quantum Algorithms basics.

Concept of Quantum Machine Learning, Advantages and limitations of QML, Quantum Data Encoding Methods, Variational Quantum Circuits, Hybrid Quantum-Classical Models, Quantum Kernel Methods

Quantum Support Vector Machines, Quantum Neural Networks, Quantum Principal Component Analysis (QPCA), Quantum k-means clustering, Quantum recommendation systems, Variational Quantum Eigensolver and QAOA in ML tasks.

QML tools and platforms (Qiskit, Cirq, PennyLane), Implementation of quantum circuits for ML, Quantum datasets and benchmarking, Applications in optimization, drug discovery, finance, cybersecurity, Challenges and future directions of QML.

Learning Resources:

Text Books:

1. Peter Wittek, “Quantum Machine Learning: What Quantum Computing Means to Data Mining”, Elsevier, 2025
2. Maria Schuld & Francesco Petruccione, “*Machine Learning with Quantum Computers”,* Second edition, Springer, 2021.

CS 36024	CLOUD AI	CREDITS 3-0-0: 3

Pre-Requisites: **Data Structures & Algorithms, Basic Machine Learning, Basic Networking Programming in Python**

Course Outcomes:

At the end of the course, the student will be able to

CO1	Analyze cloud computing architecture and services.
CO2	Design scalable cloud-native systems.
CO3	Implement containerized and serverless applications.
CO4	Develop and deploy AI applications on cloud platforms.

Course Articulation Matrix:

	PO1	PO2	PO3	PO4	PO5	PO6
CO1	3	-	1	2	-	1
CO2	-	3	2	-	2	2
CO3	1	2	-	3	1	2
CO4	-	1	3	1	1	1

1 - Slightly; 2 - Moderately; 3- Substaintially;

Syllabus:

Fundamentals of Cloud Computing: Evolution of Distributed Systems to Cloud Computing, Cloud Characteristics (NIST model), Service Models: IaaS, PaaS, SaaS, Deployment Models: Public, Private, Hybrid, Multi-cloud, Cloud Architecture and Components, Virtualization Concepts and Hypervisors, Data Center Architecture, Introduction to Cloud Providers (AWS, Azure, GCP).

Cloud Infrastructure & Storage Systems: Compute Services in Cloud, Virtual Machines and Resource Provisioning, Storage Systems: Block, File, Object Storage, Distributed Storage Systems (HDFS, GFS concepts), Cloud Databases (SQL & NoSQL), Networking in Cloud (SDN basics, Load balancing), Cloud Security Fundamentals (IAM, Encryption basics), SLA and QoS in Cloud.

Cloud-Native Technologies & Scalability**:** Containerization (Docker architecture), Container Orchestration (Kubernetes fundamentals), Microservices Architecture, Serverless Computing (FaaS concepts), Auto-scaling and Elasticity, Cloud Monitoring & Logging, DevOps & CI/CD in Cloud, Performance metrics in cloud systems.

AI in Cloud – Scalable Machine Learning Systems: AI/ML lifecycle in cloud, Managed AI Services (AutoML, ML APIs), GPU/TPU provisioning in cloud, Distributed Machine Learning, Data Parallelism vs Model Parallelism, Parameter Servers & AllReduce, ML Frameworks on Cloud (TensorFlow, PyTorch), Big Data processing for AI (Spark, Ray overview), Workflow orchestration (Kubeflow basics), MLOps fundamentals.

Advanced Cloud AI & Emerging Trends: Model Deployment & Serving (TF Serving, TorchServe), AI Microservices architecture, Serverless inference, Generative AI & LLM deployment in cloud, Edge-Cloud AI integration, Privacy-preserving AI (Federated Learning basics), AI Security in Cloud, Cost optimization & Green AI.

Learning Resources:

Text Books

1. Thomas Erl, Ricardo Puttini, Zaigham Mahmood, *Cloud Computing: Concepts, Technology & Architecture* by Cloud Native Patterns, Pearson digital enterprise series, 2013
2. *Cornelia Davis, "Cloud Native Patterns: Designing change-tolerant software", Manning, 2019*

Reference Books

1. Chip Huyen, "Machine Learning Systems Design", O'Reilly, 2022

Online Resources:

Cloud Computing & Cloud-Native Tech

1. Cloud Computing Specialization, University of Illinois / Coursera
   https://www.coursera.org/specializations/cloud-computing
2. Architecting with Google Cloud: Specialization, Coursera
   https://www.coursera.org/specializations/gcp-architecture
3. AWS Cloud Practitioner + Solutions Architect (Associate), AWS Skill Builder / ACloudGuru
   https://www.aws.training
4. Kubernetes Fundamentals, Linux Foundation / edX
   https://www.edx.org/course/introduction-to-kubernetes
5. Docker and Containers, Pluralsight / Udemy
   https://www.udemy.com/course/docker-and-kubernetes-the-practical-guide

AI, Scalable Machine Learning, MLOps

1. Machine Learning Engineering for Production (MLOps) Specialization, Coursera
   by DeepLearning.AI
   https://www.coursera.org/specializations/machine-learning-engineering-for-production-mlops
2. AI for Everyone, Andrew Ng / Coursera
   https://www.coursera.org/learn/ai-for-everyone
3. TensorFlow in Practice Specialization, Coursera
   https://www.coursera.org/specializations/tensorflow-in-practice
4. Scalable Machine Learning on Big Data using Apache Spark, Coursera
   https://www.coursera.org/learn/scalable-machine-learning
5. Practical Deep Learning for Cloud & Edge, NVIDIA DLI
   https://developer.nvidia.com/dli

CS 36026	AUTONOMOUS VEHICLES AND DRONE TECHNOLOGIES	Credits 3-0-0: 3

Pre-Requisites: **Computer Networks**

Course Outcomes:

At the end of the course, the student will be able to

CO1	Understand the knowledge of Wireless Vehicular AdHoc Networks
CO2	Analyse Sensor Technology deployment in AV and Drones
CO3	Explore AI/ML techniques in autonomy
CO4	Implement Routing algorithms used in AV and Drones

Course Articulation Matrix:

	PO1	PO2	PO3	PO4	PO5	PO6
CO1	1	-	-	3	3	3
CO2	2	2	1	-	3	3
CO3	-	-	3	3	3	3
CO4	-	2	-	3	3	3

1 - Slightly; 2 - Moderately; 3 – Substantially

Syllabus:

Introduction to Autonomous System: Evolution of Autonomous Vehicles, Levels of Automation (SAE Levels 0–5), Components of Autonomous Systems, Comparison: Ground Vehicles vs UAVs, Applications: Smart Cities, Defense, Agriculture, Logistics.

Sensors and Perception: Sensors: LiDAR, Radar, Cameras, Ultrasonic, IMU, GPS, GNSS, Sensor fusion concepts, Object detection (YOLO, CNN-based models), Lane detection and obstacle detection, Point cloud processing. Localization and Mapping: GPS-based localization, Dead reckoning, Kalman Filter, Extended Kalman Filter, Particle Filters, SLAM (Simultaneous Localization and Mapping).HD Maps.

Path Planning and Decision Making: Graph-based planning (Dijkstra, A*), Sampling-based planning (RRT, RRT*), Trajectory generation, Behavioral decision-making, Collision avoidance algorithms.

Control Systems for AVs and Drones: Vehicle kinematics and dynamics, Drone dynamics (Quadrotor model), PID control, Model Predictive Control (MPC), Adaptive and Robust control, Stability analysis.

Learning Resources:

Text Books:

1. Weisong Shi , Yuankai He, *Introduction to Autonomous Driving*, Springer, 2026
2. Kimon P. Valavanis, George J. Vachtsevanos, *"Handbook of Unmanned Aerial Vehicles"*, Springer, 2025.

CS36028	AI FOR CYBER SECURITY	Credits 3-0-0: 3

Pre-Requisites: **Cryptography and Network security, Machine and Deep Learning**

Course Outcomes:

At the end of the course, the student will be able to

CO1	Understand AI foundations applied to cybersecurity
CO2	Analyze malware and network traffic using AI
CO3	Apply deep learning for intrusion detection
CO4	Analyse intrusion detection model to distinguish attack types and generate alarms

Course Articulation Matrix:

	PO1	PO2	PO3	PO4	PO5	PO6
CO1	1	-	1	3	3	3
CO2	-	1	2	3	3	3
CO3	2	2	2	-	2	2
CO4	1	2	-	-	1	1

1 - Slightly; 2 - Moderately; 3 – Substantially

Syllabus:

Introduction to Cybersecurity: CIA triad : Confidentiality, Integrity, Availability, Security Concern: TCP/IP and security concerns, DNS and security concerns, Mail server and security concerns, Web Server and security concern, firewall, Types of Intrusion. Vulnerabilities and Threats: malware, phishing, DDoS, ransomware, Trojan Remote Access Trojan RAT, Virus, Worms and Malwares, Data poisoning.

Statistical pattern recognition for detection and classification of attacks: Data Collection Mechanism: Data Sampling, Packet Sampling, Flow Sampling, Network Traffic Monitoring and Analysis: Examining Packet Header Fields, normal and abnormal values in IP, TCP, UDP, and ICMP header fields, buffer overflows and race conditions, malicious codes.

Tools: Tcpdump windump, Wireshark tool, Writing Tcpdump/Windump Filters, libcap/winpcap libraries, pcap file, sniffing and spoofing tools.

AI for Threat Intelligence: Artificial Intelligence-Based Intrusion Detection Techniques, Behavioral Analytics, User and Entity Behavior Analytics, Insider threat detection, anomaly detection, signature-based analysis, Semantic-aware signature, policy-based analysis, and host-based analysis. Evaluation and tuning of IDS, Reduction of false negatives and positives.

Intrusion Detection Systems (IDS): IDS and IPS Definition, and IDS Architecture, IDS/IPS Management, Issues with regard to deploying IDS/IPS systems.

Types IDS: host-based IDS, network-based IDS, Methods used in intrusion detection: signature-based IDS, Anomaly based IDS, Hybrid IDS, Honeypot, Shadow Honeypot

IDS tools: Snort and Bro IDS tools, NIDS Evasion, Insertion, and Checksums to confuse NID systems, Snort Fundamentals and Configuration, Snort GUIs & Sensor Management, Snort Performance, Active Response & Tagging, Snort Rules, Stimulus Response.

Federated AI Models for Privacy-Preserving Intelligent Systems: Introduction, Federated Averaging (FedAvg), Federated SGD, Personalized FL models, Meta-learning (MAML-based FL), Few-shot learning in FL, Transfer learning in FL, Federated Deep Neural Networks (CNN, LSTM, GNN).

Learning Resources:

Text Books:

1. Endorf, C., Schultz, E., and Mellander, J., "*Intrusion Detection and Prevention*," McGraw-Hill. 2003.
2. Wenliang Du, "*Computer Security: A Hands-on Approach",* Third edition, SEED Labs, 2022
3. McDonald, Malcolm, "*Web security for developers: Real threats, practical defence*", No Starch Press, 2020.
4. Yaworski, Peter, "*Real-world bug hunting: a field guide to web hacking*", No Starch Press, 2019.
5. Stephen Northcutt and Judy Novak, "Network Intrusion Detection", Sams Publishing, 3rd edition, 2002.
6. Bejtlich, "*Extrusion Detection: Security Monitoring for Internal Intrusions" ,* 1st Edition, Pearson Education, 2005.

CS36030	AI FOR DIGITAL FORENSICS	Credits 3-0-0: 3

Pre-Requisites: **Machine Learning**

Course Outcomes:

At the end of the course, the student will be able to

CO1	Understand the fundamentals of digital forensics and cybercrime investigation
CO2	Apply AI methods to detect cyber attacks, malware, and digital evidence patterns
CO3	Use AI-based tools for automated evidence collection and analysis
CO4	Analyse how Artificial Intelligence and Machine Learning techniques assist forensic analysis.

Course Articulation Matrix:

	PO1	PO2	PO3	PO4	PO5	PO6
CO1	1	1	3	3	3	3
CO2	2	-	3	3	3	3
CO3	-	2	1	3	3	3
CO4	-	-	3	3	3	3

1 - Slightly; 2 - Moderately; 3 – Substantially

Syllabus:

Fundamentals of Digital Forensics: Introduction to Digital Forensics, Cybercrime and Digital Evidence, Forensic Investigation Process, Chain of Custody, Legal and Ethical Issues in Digital Forensics, Types of Digital Evidence (volatile and non-volatile).

AI Techniques in Digital Forensics: Pattern Recognition in Digital Evidence, AI-based Log Analysis, Natural Language Processing for forensic text analysis, Image and Video Forensics using AI, AI-based anomaly detection, Timeline reconstruction using machine learning.

Malware and Network Forensics using AI: Malware detection using machine learning, AI-based intrusion detection systems (IDS), Network traffic analysis, Behavioral analysis of cyber attacks, Threat intelligence using AI.

Tools and Case Studies: Digital forensic tools (Autopsy, FTK, EnCase), AI tools for cyber investigation, Evidence extraction and automated analysis, Case studies of cybercrime investigations, Forensic reporting and presentation.

Learning Resources:

Text Books:

1. Nelson, Philips, steuar, " *Guide to Computer Forensics and Investigations*", Cengage, Sixth edition, 2020
2. Darren Hayes, *"A Practical Guide to Digital Forensics Investigations*", Second edition, Pearson, 2022.

CS36032	GENERATIVE AI	Credits 3-0-0: 3

Pre-Requisites: **Machine Learning**

Course Outcomes:

At the end of the course, the student will be able to

CO1	Explain generative AI models and foundations.
CO2	Apply prompt engineering and RAG techniques.
CO3	Analyze multimodal generative AI systems.
CO4	Develop generative AI applications.

Course Articulation Matrix:

	PO1	PO2	PO3	PO4	PO5	PO6
CO1	1	1	3	3	3	3
CO2	-	-	2	3	3	3
CO3	1	2	2	3	3	3
CO4	-	-	-	3	3	3

1 - Slightly; 2 - Moderately; 3 – Substantially

Syllabus:

Overview of Generative AI: Concepts, Scope, and Applications in Various Domains. Mathematical & Computational Foundations: Linear Algebra (Vectors, Matrices, Eigenvalues), Probability, Statistics, and Optimization basics for Generative modelling.

Generative Model Categories: Autoregressive, Energy-based, Latent Variable Models. Variational Autoencoders (VAEs): Architecture, Loss functions, Applications. Generative Adversarial Networks (GANs): Architecture, Loss functions, Applications, and Variants (DCGAN, CycleGAN, StyleGAN).

Introduction to Language Models: Statistical, Neural and Transformer-based. Transformer Architecture (Attention Is All You Need).

Prompt Engineering Fundamentals: Zero-shot, Few-shot, Chain-of-thought Prompting. Strategies for Controlling Model Behaviour & Output Quality. Lang Chain Framework: Architecture, Components, Chains, Memory, Tool integration.

Retrieval-Augmented Generation (RAG): Embeddings, Indexing, Vector databases, and Workflow. Practical Example: Building a Domain-specific Chatbot.

Code Generation: Codex, Copilot. Image Generation: DALL-E, Stable Diffusion, Midjourney. Audio generation: Text-to-Speech (TTS), Music Synthesis Models (e.g., MusicLM).

Video Generation: Current State-of-the-art Models (Runway Gen-2, Pika Labs, Sora by OpenAI). Multimodal AI: CLIP, Flamingo, GPT-4-Vision.

Industrial Applications: Summarization, Creative Design, Education, Healthcare.

Learning Resources:

Text Books

1. Alammar, J., & Grootendorst, M., “Hands-On Large Language Models.” O’Reilly Media, 2023
2. Auffarth, B., & Kuligin, A., “Generative AI with LangChain: Second edition, Packt Publishing, 2024
3. Goodfellow, Y., Bengio, Y., & Courville, A., “*Deep Learning.”* MIT Press, 2016

Reference Books

1. Phoenix, J., “*Prompt Engineering for Generative AI”*. Packt Publishing, 2023.
2. Vurukonda, V, & Srivastava, A., “*Prompt Engineering for AI Systems”.* Manning Publications, 2024.

Online Resources:

https://study.iitm.ac.in/ds/course_pages/BSDA5002.html?utm_source=chatgpt.com

CS36034	AI FOR GRAPH OPTIMIZATION	Credits 3-0-0: 3

Pre-Requisites: **Data Structures and Algorithms, Discrete Mathematics, Basic Machine Learning**

Course Outcomes:

At the end of the course, the student will be able to

CO1	Model real-world problems as graph optimization tasks.
CO2	Apply classical and AI-based approaches for graph optimization.
CO3	Design and implement learning-based graph solvers.
CO4	Analyze performance and scalability of AI-driven graph algorithms.

Course Articulation Matrix:

	PO1	PO2	PO3	PO4	PO5	PO6
CO1	-	2	3	3	3	-
CO2	1	-	2	3	3	3
CO3	2	-	1	3	-	2
CO4	-	3	3	-	3	3

1 - Slightly; 2 - Moderately; 3 – Substantially

Syllabus:

Graph Theory Fundamentals: Shortest path, MST, Max-flow, Traveling Salesman Problem (TSP), Graph colouring, matching, cuts, NP-hardness in graph problems, Approximation and heuristic algorithms

Classical Optimization Techniques: Linear programming formulations, Integer linear programming for graphs, Branch and bound , Dynamic programming, Greedy and local search, Meta-heuristics (Genetic Algorithms, Simulated Annealing).

Machine Learning for Graph Problems: Learning to optimize, Supervised learning for combinatorial optimization, Reinforcement learning for routing, Neural combinatorial optimization, Pointer networks, Policy gradient methods.

Graph Neural Networks (GNNs) : Graph representation learning, Message passing neural networks, Spectral vs spatial GNNs, Graph attention networks, Learning heuristics using GNNs, Applications in routing and scheduling.

AI for Large-Scale Network Optimization: AI in transportation networks, Wireless network optimization, VANET routing optimization, Congestion prediction using ML , AI-driven resource allocation

Advanced Topics : Federated learning on graph data, Adversarial attacks on graph models, Robust graph optimization, Explainable AI for graph decisions, and Scalable distributed graph learning.

Learning Resources:

Text Books :

1. J A Bondy & Murty, "*Graph Theory with Applications*", Revised edition, Elsevier Science Publishing, 1997.
2. Sutton & Barto, "*Reinforcement Learning",* Second edition, MIT Press, 2018
3. William L. Hamilton, "*Graph Representation Learning*" , First edition, Springer nature, 2020.

CS 36036	AI FOR HEALTHCARE	Credits 3-0-0: 3

Pre-Requisites: **None**

Course Outcomes:

At the end of the course, the student will be able to

CO1	Apply machine learning and deep learning techniques to healthcare problems.
CO2	Analyze medical images, clinical text, and physiological signals using AI.
CO3	Design AI-based clinical decision support systems.
CO4	Understand ethical, regulatory, and deployment challenges in healthcare AI.

Course Articulation Matrix:

	PO1	PO2	PO3	PO4	PO5	PO6
CO1	2	3	-	-	3	3
CO2	-	-	3	3	3	3
CO3	2	-	3	3	3	3
CO4	-	1	-	3	3	3

1 - Slightly; 2 - Moderately; 3 – Substantially

Syllabus:

Overview of healthcare systems and workflows. Role and scope of AI in healthcare. Applications in diagnosis, prognosis, treatment planning, and monitoring. Types of healthcare data: Electronic Health Records (EHR), clinical text, medical imaging (X-ray, CT, MRI), and physiological signals (ECG, EEG). Healthcare datasets and sources. Data preprocessing techniques including missing data handling, normalization, noise removal, annotation, and class imbalance handling.

Supervised and unsupervised learning methods. Classification and regression techniques. Logistic regression, decision trees, random forests, and support vector machines. Model evaluation metrics: accuracy, precision, recall, F1 score, ROC-AUC. Cross-validation, overfitting, and bias-variance tradeoff. Clinical prediction models and disease risk prediction systems.

Neural network fundamentals and training using backpropagation. Convolutional Neural Networks (CNNs) for medical image analysis. Image preprocessing, feature extraction, image classification, and segmentation. Advanced architectures such as ResNet and U-Net. Applications in tumor detection, organ segmentation, and disease diagnosis.

Natural Language Processing for clinical text analysis. Named entity recognition, clinical coding, and information extraction. Transformer models and clinical language models. Time-series analysis of physiological signals such as ECG and EEG. Recurrent neural networks (RNN), LSTM, and GRU for patient monitoring and disease prediction.

Explainable AI techniques including SHAP and LIME for model interpretability. Generative AI for synthetic medical data generation and augmentation. Clinical decision support systems and personalized medicine. Deployment of AI models in healthcare environments, cloud and edge deployment. Ethical, privacy, and regulatory considerations in healthcare AI. Emerging trends including multimodal AI and wearable healthcare systems.

Learning Resources:

Text Books

1. Bikesh Kumar Singh, G.R. Sinha., "*Machine Learning in Healthcare Fundamentals and Recent Applications*", CRC Press, 2022.
2. Gyorgy J. Simon, Constantin Aliferis., "*Artificial Intelligence and Machine Learning in Health Care and Medical Sciences"*, Springer, 2024.
3. Bhanu Chander, Koppala Guravaiah, B. Anoop, G. Kumaravelan, "*Handbook of AI-Based Models in Healthcare and Medicine*", CRC Press, 2024.

Reference Books

1. Sandeep Reddy, "Translational Application of Artificial Intelligence in Healthcare", CRC Press, 2024.

Online Resources:

- https://www.coursera.org/learn/ai-for-healthcare
- https://www.coursera.org/learn/ai-technologies-in-healthcare

CS36038	SOCIAL NETWORK ANALYSIS	Credits 3-0-0: 3

Pre-Requisites: **None**

Course Outcomes:

At the end of the course, the student will be able to

CO1	Understand the importance of social networks and social graphs.
CO2	Enhance analytical skills for analyzing social networking data.
CO3	Develop skills to leverage extended enterprise data
CO4	Create real-life case studies using social networks.

Course Articulation Matrix:

	PO1	PO2	PO3	PO4	PO5	PO6
CO1	-	1	-	2	3	3
CO2	1	2	2	-	3	3
CO3	2	-	3	3	-	-
CO4	-	3	-	-	-	3

1 - Slightly; 2 - Moderately; 3 – Substantially

Syllabus:

Introduction to social network analysis: Graphs – nodes, edges, direct and indirect friends/neighbors, degree and degree distribution, shortest path, cycle, tree, complete graph, bipartite graphs, directed graphs, weighted graphs, adjacency matrix, social interactions and connected components. Technological networks (internet, telephone network, power grids, transportation networks), social networks (facebook, movie collaboration, paper collaboration), information networks (web), biological networks (neural networks, ecological networks).

Network Centrality Measures and Models: Properties of real-world network – degree distribution, clustering coefficient, average path length; Random Graphs – Evolution of random graphs, properties of random graphs, modeling real-world networks with random graphs, Erdos-Renyi model of random graph; Small-world Model – Properties of the Small-world model, modeling real-world networks with the small-world model; Preferential attachment model – Properties of the preferential attachment model, modeling real-world networks with the preferential attachment model.

Random walk-based proximity measures, other graph-based proximity measures. Clustering with random-walk based measures.

Influence and Homophily: Measuring Assortativity, Measuring and modeling Influence, Measuring and modeling Homophily, Distinguishing influence and homophily – shuffle test, edge-reversal test, randomization test; Spread of influence through a network, influence maximization in networks, spread of disease on networks.

Games on networks, game theory strategies, dominant strategies, dominated strategies, pure strategies and mixed strategies, Nash equilibrium, multiple equilibria-coordination games, multiple equilibria-the Hawk-Dove game, mixed strategies, Modeling social network traffic using game theory.

Learning Resources:

Text Books

1. Reza Zafarani, Mohammad Ali Abbasi, Huan Liu, "*Social Media Mining – An Introduction*", Cambridge University Press, 2014.
2. David Easley and Jon Kleinberg, "*Networks, Crowds, and Markets: Reasoning About a Highly Connected World*", Cambridge University Press, 2010.

Reference Books:

1. Mark Newman, "*Networks: An Introduction*", Oxford University Press,2010.
2. Hansen, Derek, Ben Sheiderman, Marc Smith, "*Analyzing Social Media Networks with NodeXL: Insights from a Connected World*", Morgan Kaufmann,2011.
3. Avinash Kaushik, "Web Analytics 2.0: *The Art of Online Accountability*", Sybex,2009.

CS36040	AI FOR CYBER-PHYSICAL SYSTEMS	Credits 3-0-0: 3

Pre-Requisites: **Discrete Mathematics, Theory of Computation, Artificial Intelligence, Machine Learning**

Course Outcomes:

At the end of the course, the student will be able to

CO1	To introduce the fundamental architecture of Cyber-Physical Systems (CPS) and the role of AI in their evolution.
CO2	To master the modeling of continuous and discrete dynamics using hybrid system frameworks
CO3	To apply advanced AI techniques (RL, DL, Anomaly Detection) for the control, optimization, and monitoring of physical processes
CO4	To address the unique security, safety, and reliability challenges in AI-driven physical systems

Course Articulation Matrix:

	PO1	PO2	PO3	PO4	PO5	PO6
CO1	3	-	-	3	3	3
CO2	2	2	2	3	3	3
CO3	-	1	2	3	3	3
CO4	2	3	-	-	3	3

1 - Slightly; 2 - Moderately; 3 – Substantially

Syllabus:

Introduction to Cyber-Physical Systems: Key Features-Synchronous Models-Safety Requirements-Asynchronous Models-Liveness Requirements-Dynamical Systems-Timed Models-Real Time Scheduling-Hybrid Systems.

Adaptive Control in CPS and Essence of AI: Types of CPS: Linear vs Nonlinear-Why Adaptive Control-Adaptive Control Versus Conventional Feedback Control-Basic Adaptive Control Schemes.

AI: From rule-based systems to machine learning, deep learning, and neural networks.

AI for Control and Optimization: Reinforcement Learning (RL) in CPS: Markov Decision Processes (MDP), Q-Learning, and Policy Gradients for robotic control. Model Predictive Control (MPC) with AI: Integrating neural networks into MPC frameworks for better trajectory planning. Edge AI: Deploying lightweight AI models on embedded hardware (FPGA, ARM Cortex-M). Anomaly Detection.

AI for Secure and Resilient CPS:ML-based Intrusion detection, AI-supported risk assessment, AIML for Advanced Threat Detection, Adaptive defense strategies.

Case Study of Smart Manufacturing/Industry 5.0.

Learning Resources:

Text Books

1. Rajeev Alur, "*Principles of Cyber-Physical Systems*", MIT Press Books, 2nd edition, 2023
2. Landau, Lazo, Saasd and Karimi, *"Adaptive Control: Algorithms, Analysis and Applications*", 3rd Edition, 2022.
3. Pushan Kumar Dutta, Pethuru Raj, "*Artificial Intelligence Solutions for Cyber-Physical Systems"*, First edition, 2025.

CS36042	EXPLANIABLE AI	Credits 3-0-0: 3

Pre-Requisites: **Artificial Intelligence, Machine Learning**

Course Outcomes:

At the end of the course, the student will be able to

CO1	To apply interpretability techniques, such as feature importance analysis, rule-based models, and model-agnostic methods, to enhance the transparency of machine learning models.
CO2	To evaluate and compare different interpretability methods.
CO3	To develop a critical understanding of the trade-offs between model complexity and interpretability in diverse scenarios.
CO4	To demonstrating practical skills, students will apply explainability techniques to real-world datasets and challenges.

Course Articulation Matrix:

	PO1	PO2	PO3	PO4	PO5	PO6
CO1	2	2	3	2	2	3
CO2	2	2	3	2	2	3
CO3	2	2	3	2	2	3
CO4	2	2	3	2	2	3

1 - Slightly; 2 - Moderately; 3 – Substantially

Syllabus:

Introduction to Explainable AI : Science of Interpretable Machine Learning, Motivation, Challenges, and Mythos of Model Interpretability, Human Factors in Explainability, Interpreting Interpretability.

Post hoc Explanations:

Explaining the Predictions of Any Classifier, Pitfalls, Challenges, and Evaluation of Feature Attributions, LIME and SHAP, OpenXAI, The Disagreement Problem in Explainable Machine Learning, Counterfactual Explanations (or) Algorithmic Recourse, Learning Model-Agnostic Counterfactual Explanations for Tabular Data.

Attention and Concept Based Explanations: Quantifying Interpretability of Deep Visual Representations, Interpretability Beyond Feature Attribution, Data Attribution and Interactive Explanation, Equitable Valuation of Data, Explainable Active Learning (XAL), Theory of Explainability and Interpreting Generative Models.

Explainability for Fair Machine Learning: Connections with Robustness, Privacy, Fairness, and Unlearning, Right to Explanation and the Right to be Forgotten, Fairness via Explanation Quality, Mechanistic Interpretability and Compiled Transformers, Understanding and Reasoning in Large Language Models.

Learning Resources:

Text Books :

1. Michael Munn, David Pitman, "*Explainable AI for Practitioners*", O'Reilly Media, Inc. , 2022.
2. Leonida Gianfagna, Antonio Di Cecco, "*Explainable AI with Python*", Springer , 2021.

Reference Books:

1. Christoph Molnar, Interpretable Machine Learning, "*A Guide for Making Black Box Models Explainable*," Leanpub , 2023.
2. Denis Rothman, "Hands-On Explainable AI (XAI) with Python: Interpret, visualize, explain, and integrate reliable AI for fair, secure, and trustworthy AI apps", Packt Publishing Limited , 2020.

Online Resources:

1. AAAI Conference on Artificial Intelligence.
2. Explainable Artificial Intelligence, ELSEVIER.

CS36044	HARDWARE ACCELERATION AND OPTIMIZATION FOR AI	Credits 3-0-0: 3

Pre-Requisites: Basics of linear algebra, Computer Organization and Architecture, basic knowledge of C/Python

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

CO-1	Analyze for the performance of AI algorithms
CO-2	Generate parallel programs for real world problems using Cuda library
CO-3	Compare the performance of different algorithms for the numerical and data processing problems on FPGA/GPGPUs and suggest methods for improving the performance.
CO-4	Explore Real time deployment of complex problems

Course Articulation Matrix:

	PO1	PO2	PO3	PO4	PO5	PO6
CO1	2	2	2	3	2	1
CO2	2	3	2	3	2	1
CO3	2	3	3	2	2	1
CO4	2	2	2	2	2	3

1 - Slightly; 2 - Moderately; 3 - Substantially

Syllabus:

Introduction: Overview of AI and machine learning hardware systems, motivation and trends, fundamentals of digital hardware – FPGA, power and speed estimation, accelerating linear algebra,

AI system concepts – SVM and Deep Learning Neural Networks, feature extraction (PCA, filtering), inference engine, matrix vector multiplication (sparsity), non-linearity and pooling, resolution-performance trade-off, training optimization engines (cost function, regularization), online and stochastic training, forward-backward propagation, emerging hardware architectures, memristor based designs, spiking architectures.

AI Acceleration Hardware: Precision scaling, in-memory computing, hyperdimensional computing, architectural modifications, GPUs, and vector architectures, quantum computing as well as recent, hardware programming tools such as Xilinx AI Vitis, Xilinx, HLS, and IBM Qiskit.

Learning Resources:

Text Books:

1. Jason Sanders, Edward Kandrot, "*CUDA By Example – An Introduction to General-Purpose GPU Programming*", Addison Wesley, 2011

CS36046	AI FOR DATA PRIVACY	Credits 3-0-0: 3

Pre-Requisites: **Machine Learning, Deep Learning**

Course Outcomes:

At the end of the course, the student will be able to

CO1	Understand the role of AI for data privacy.
CO2	Design mechanisms for query release problem and response learning algorithms.
CO3	Design of noise addition techniques to achieve data privacy.
CO4	Analyze computational complexity of differentially private mechanisms

Course Articulation Matrix:

	PO1	PO2	PO3	PO4	PO5	PO6
CO1	1	1	1	3	3	3
CO2	-	2	2	3	3	3
CO3	-	-	3	3	3	3
CO4	1	-	-	3	-	3

1 - Slightly; 2 - Moderately; 3 – Substantially

Syllabus:

Introduction to AI and Data Privacy: Definition of data privacy, Data collection and data publishing concept, Concept of sensitive data, Risk of privacy leaks, Privacy preserving data publishing, Role of AI in privacy preserving data publishing.

Data Privacy Challenges: Data leaks, identification, record linkage, attribute linkage, Re-identification attacks, Homogenic attack, Skewness attack, Data leakage in machine learning model, Issues in automated decision making, Reconstruction attacks, Sensitivity in differential privacy.

Fundamentals of Data Protection: Definition: Categorial data, Statistical data, identical dataset, adjacent dataset, Privacy preserving data publishing Techniques- Non-Interactive technique: K-anonymization, diversity, t-closeness.Interactive techniques : Curator model, Randomized response, Differential privacy.

Application of AI/ML Techniques in Data Privacy: Generating Random Noise, Laplace mechanism, Gaussian Mechanism, Exponential mechanism, (ε,0)-Differential Privacy, (ε,δ)- differential privacy, Sampling model, Histogram, Gradient descent, Modern machine learning and adaptive data analysis, Approximate Differential Privacy, Advanced Composition.

Applications of AI with Differential Privacy: Healthcare data privacy, Privacy-preserving recommendation systems, Smart city data analytics, Financial fraud detection, Privacy in social media analytics, AI for secure data sharing.

Learning Resources:

Text Books:

1. C. Dwork and A. Roth, *"The Algorithmic Foundations of Differential Privacy"*, First edition, Publishers, 2014.

CS36048	FEDERATED LEARNING	Credits 3-0-0: 3

Pre-Requisites: **Machine Learning**

Course Outcomes:

At the end of the course, the student will be able to

CO1	Understand popular methods used in federated learning
CO2	Be able to construct and scale a simple federated system
CO3	Have gained an appreciation of the core limitations to existing methods, and the approaches available to cope with these issues
CO4	Developed an intuition for related technologies like differential privacy and secure aggregation, and are able to use them within typical federated settings

Course Articulation Matrix:

	PO1	PO2	PO3	PO4	PO5	PO6
CO1	1	-	2	3	3	3
CO2	2	2	3	3	-	3
CO3	-	1	-	2	3	1
CO4	-	2	-	3	3	3

1 - Slightly; 2 - Moderately; 3 – Substantially

Syllabus:

Introduction to Federated Learning, Decentralized Optimization, Statistical and Systems Heterogeneity, Variations of Federated Aggregation, Secure Aggregation, Differential Privacy within Federated Systems, Extensions to Federated Analytics, Applications to Speech, Video, Images and Robotics, Federating a Centralized ML Classifier, Behaviour under Heterogeneity, Scaling a Federated Implementation, Exploring Privacy with Federated Settings.

Learning Resources:

Text Books

1. Qiang Yang, Yang Liu, Yong Cheng, Yan Kang, Tianjian Chen, Han Yu, "*Federated Learning",* Morgan & Claypoo lPublishers, 2019.
2. Kim-Kwang Raymond Choo, Ali Dehghantanha, "*Handbook of Big Data Privacy"*, Springer Nature Switzerland, 2020.

Reference Books

1. Ludwig, H., & Baracaldo, N, "*Federated learning: A comprehensive overview of methods and applications"*, Springer, 2022
2. Jin, Y., Zhu, H., Xu, J., & Chen, Y., "*Federated Learning"* Springer Nature Singapore,2023.

Online Resources:

1. https://onlinecourses.nptel.ac.in/noc24_cs86/preview
2. https://github.com/alexjungaalto/FederatedLearning

CS36050	ALGORITHMIC GAME THEORY	Credits 3-0-0: 3

Pre-Requisites: **Data Structures and Algorithms, Probability and Statistics**

Course Outcomes:

At the end of the course, the student will be able to

CO1	Analyze strategic interactions computationally by modeling games and computing or approximating Nash equilibria.
CO2	Design and evaluate incentive-compatible mechanisms, including auctions and resource allocation algorithms.
CO3	Assess efficiency loss due to selfish behavior using concepts such as the Price of Anarchy.
CO4	Apply game-theoretic algorithms to real-world systems, such as networks, online markets, and distributed platforms.

Course Articulation Matrix:

	PO1	PO2	PO3	PO4	PO5	PO6
CO1	1	-	-	3	3	3
CO2	-	2	2	-	3	3
CO3	-	-	3	3	3	3
CO4	1	1	3	3	3	3

1 - Slightly; 2 - Moderately; 3 – Substantially

Syllabus:

Introduction to Game Theory : Games in Normal Form (strategic games), Zero-sum and General-sum games. Prisoner's dilemmaMatching pennies, Nash Equilibrium (NE) and Best Response Functions.
Computational Complexity of Games: Finding Nash equilibria, complexity classes: FNP, TFNP, PPAD. Mixed Strategies and Randomized Algorithms. Potential Games and Local Search.
Inefficiency of Equilibria : Price of Anarchy (PoA) and Price of Stability (PoS). Selfish Routing, Congestion Games, and Wardrop Equilibrium. Braess's Paradox.
Mechanism Design (Algorithm Design for Incentives): Introduction to Mechanism Design (inverse game theory).Quasi-linear utility, Social Welfare maximization.VCG Mechanism (Vickrey-Clarke-Groves). Gibbard-Satterthwaite Theorem.
Auctions and Markets: Single-item auctions, Sealed-bid auctions. Combinatorial Auctions, Sponsored Search Auctions (Generalized Second Price). Market Equilibrium and Stability.
Learning and Dynamics: Regret Minimization, No-Regret Dynamics. Correlated Equilibrium. Stable Matching (Gale-Shapley algorithm).

Learning Resources:

Text Books

1. Nisan, Roughgarden, Tardos, and Vazirani, "*Algorithmic Game Theory"* Cambridge University Press, First edition, 2007.
2. Michael Maschler, Eilon Solan, and Shmuel Zamir, "*Game Theory*", *2nd edition, 2020.*
3. Y. Narahari, *Game Theory and Mechanism Design*

Reference Books

1. Martin J. Osborne, "An Introduction to Game Theory",
2. **Game Theory** – Mas-Colell, Whinston, Green

Online Resources:

1. *Course Notes from NPTEL/Swayam - Algorithmic Game Theory*.
2. **Twenty Lectures on Algorithmic Game Theory** – Tim Roughgarden
3. https://youtu.be/YVz4APWm1FE

CS 36052	TESTING AI MODELS	Credits 3-0-0: 3

Pre-Requisites: **Deep Learning**

Course Outcomes:

At the end of the course, the student will be able to

CO1	Understand the fundamental differences between traditional software testing and AI/ML model testing
CO2	Analyze and validate data used in AI systems by performing data quality testing
CO3	Apply appropriate evaluation metrics and statistical testing techniques—such as confusion matrices, ROC–AUC, precision–recall, cross-validation
CO4	Evaluate the robustness, fairness, and explainability of AI models by conducting adversarial testing, stress and scalability testing

Course Articulation Matrix:

	PO1	PO2	PO3	PO4	PO5	PO6
CO1	2	2	3	2		3
CO2	2	2	3	2		3
CO3	2	2	3	2		3
CO4	2	2	3	2		3

1 - Slightly; 2 - Moderately; 3 – Substantially

Syllabus:

Foundations of AI Testing: Introduction to AI Systems and Testing Challenges, Traditional Software Testing vs. AI Testing, AI Lifecycle and Failure Points, Types of Testing in AI. Data Quality Testing, Data Distribution Analysis and Drift Detection, Bias and Fairness Testing in Datasets, Data Leakage, Synthetic Data and Test Case Generation for ML. Evaluation Metrics for Classification, Regression, Confusion Matrix, ROC–AUC, Precision–Recall, Cross-Validation and Statistical Significance Testing, Error Analysis and Model Debugging, Calibration and Uncertainty Estimation. Robustness to Noise and Perturbations, Adversarial Attacks, Stress Testing and Scalability Testing. Fairness Metrics and Bias Auditing, Explainable AI: LIME, SHAP, Feature Importance, Interpretability vs Performance Trade-offs, Ethical Testing and Responsible AI Practices. Regression Testing for ML Models, A/B Testing and Online Evaluation, Monitoring in Production and Model Drift Detection, Automated Testing Pipelines in MLOps and Case Studies.

Learning Resources:

Text Books

1. Paul Ammann, Jeff Offutt, "*Introduction to Software Testing*", 2008
2. Kush R. Varshney, "*Trustworthy Machine Learning*" 2022
3. Christopher Bishop, "*Pattern Recognition and Machine Learning*", 2006

Reference Books

1. Ian Goodfellow, Yoshua Bengio, Aaron Courville, "Deep Learning", 2015

CS 36054	ROBOTICS	Credits 3-0-0: 3

Pre-Requisites: None

Course Outcomes:

CO1	Apply mathematical modeling, coordinate transformations, and kinematic analysis to solve robotic motion and workspace problems.
CO2	Design and implement perception, localization, and path planning algorithms using probabilistic and graph-based approaches.
CO3	Develop intelligent robotic control systems using reinforcement learning and deep learning techniques.
CO4	Analyze and evaluate the safety, robustness, and ethical implications of AI-driven autonomous robotic systems.

Course Articulation Matrix:

	PO1	PO2	PO3	PO4	PO5	PO6
CO1	1	-	2	3	3	3
CO2	2	2	3	3	3	3
CO3	-	2	-	3	3	3
CO4	3	3	-	3	3	3

1 - Slightly; 2 - Moderately; 3 – Substantially

Syllabus:

Introduction to Robotics: History, growth and AI-driven applications; Types and classification of robots; Ethical and societal considerations in AI-driven robotic systems; Robot architectures and computational frameworks; Coordinate frames and rigid body transformations; Homogeneous transformation matrices; Denavit–Hartenberg (DH) parameters; Forward kinematics; Inverse kinematics; Jacobian matrix and differential motion; Workspace analysis; Trajectory planning; Overview of robot dynamics; Planning and control architecture in intelligent robotic systems. Actuators: DC motors, BLDC motors, Servo motors, Mathematical modeling of actuators, PWM control, Joint motion control, Feedback control systems, PID control, Introduction to model-based and computed torque control; Sensors: IMU, LIDAR, Vision sensors, Sensor integration and sensor fusion, Computer vision basics for robotics, Feature extraction and object detection, Deep learning for robotic perception, Probabilistic robotics framework, Bayes filters, Kalman Filter and Extended Kalman Filter, Particle Filter, Localization techniques. Configuration space, Graph-based path planning: BFS, DFS, Dijkstra, A*, D*, Sampling-based methods, Voronoi diagrams, Potential field methods, Hybrid approaches, Simultaneous Localization and Mapping (SLAM), EKF-SLAM, Graph-based SLAM, Practical considerations in real systems Reinforcement Learning for robotics: Markov Decision Process (MDP), Q-learning, Policy gradients, RL for navigation and control, Imitation learning, Deep Reinforcement Learning, Multi-agent robotics, Safe and ethical AI in robotics, Introduction to secure and robust robotic systems.

Learning Resources:

Text Books

1. John J. Craig, Introduction to Robotics: Mechanics and Control, Pearson, 2005
2. Kevin M. Lynch & Frank C. Park Modern Robotics: Mechanics, Planning, and Control, Cambridge University Press, 2013

Reference Books

1. Forsyth & Ponce Computer Vision: A Modern Approach, Pearson, 2015
2. Spong, Hutchinson & Vidyasagar Robot Modeling and Control, Wiley, Second edition, 2013

CS36056	MULTI-MODAL DATA ANALYTICS	Credits 3-0-0: 3

Pre-Requisites: **Deep Learning**

Course Outcomes:

At the end of the course, the student will be able to

CO1	Demonstrate a comprehensive understanding of the different types of multi-modal data
CO2	Critically evaluate and measure the effectiveness and efficiency of different data integration, analysis, and machine learning techniques on multi-modal datasets
CO3	Apply advanced data fusion techniques and machine learning models to integrate multi-modal data, enhancing prediction, classification, and decision-making capabilities
CO4	Select and adapt appropriate algorithms and models to address specific challenges

Course Articulation Matrix:

	PO1	PO2	PO3	PO4	PO5	PO6
CO1	2	2	3	2	-	3
CO2	2	2	3	2	-	3
CO3	2	2	3	2	-	3
CO4	2	2	3	2	-	3

1 - Slightly; 2 - Moderately; 3 – Substantially

Syllabus:

Introduction to multi-modal data, Types of modalities, Challenges in multi-modal learning, Data preprocessing and normalization, Feature extraction for text, images and audio/video, Dimensionality reduction and representation learning, Introduction to multi-modal fusion strategies, Classical multi-modal and multi-view learning approaches, Deep learning foundations for multi-modal analysis, Multi-modal neural network architectures, Joint embedding and cross-modal representation learning, Transformer-based multi-modal models, Multi-modal sequence and time-series data, Event detection and sequential decision making, Loss functions and optimization for multi-modal models, Evaluation metrics and benchmarking datasets, Handling missing, noisy, and imbalanced modalities, Self-supervised and contrastive multi-modal learning, Zero-shot and few-shot multi-modal learning, Robustness, fairness, bias, and explainability, Applications in various domains.

Learning Resources:

Text Books :

1. Christopher Bishop, "*Pattern Recognition and Machine Learning*", Springer-Verlag New York Inc, 2006
2. Kevin P. Murphy, "*Machine Learning: A Probabilistic Perspective*", The MIT Press, 2012
3. Daniel Jurafsky and James H. Martin, "*Speech and Language Processing*", Second edition, Pearson, 2026

Reference Books:

1. Ian Goodfellow, Yoshua Bengio, Aaron Courville , "Deep Learning", The MIT Press, 2016

Online Resources:

1. https://ieeexplore.ieee.org/document/8269806
2. https://www.cs.cmu.edu/~morency/MMML-Tutorial-ACL2017.pdf

CS 36058	eXTENDED REALITY (XR)	Credits 3-0-0: 3

Pre-Requisites: **None**

Course Outcomes:

At the end of the course, the student will be able to

CO1	Explain concepts of Augmented, Virtual, and Extended Reality.
CO2	Apply perception, tracking, and interaction techniques in XR systems.
CO3	Design XR solutions for real-world applications.
CO4	Analyze human-centered and usability aspects of XR.
CO5	Evaluate emerging trends, including AI-enabled XR.

Course Articulation Matrix:

	PO1	PO2	PO3	PO4	PO5	PO6
CO1	1	1	1	3	3	3
CO2	-	2	1	3	3	3
CO3	1	-	-	3	-	3
CO4	2	-	-	3	3	3
CO5	-	-	1	3	3	3

1 - Slightly; 2 - Moderately; 3 – Substantially

Syllabus:

Foundations of Extended Reality (XR): Reality–Virtuality continuum; Overview of AR, VR, MR and XR; XR system architecture and components; Displays: optical and video see-through; Sensors and tracking fundamentals; Immersion, presence, and embodiment; Overview of AI in XR.

Computer Vision, Tracking and Registration for AR: Coordinate systems and transformations; Camera models and calibration; Marker-based tracking; Markerless tracking and SLAM; Sensor fusion; Registration and alignment; Object recognition for AR.

Human Perception and Interaction in XR: Visual, auditory, and haptic perception; Depth, motion, and spatial cognition; Latency and cybersickness; Interaction techniques: gesture, gaze, voice; User interface design in XR; Human-centered and iterative design.

Virtual Reality Systems and Immersive Environments: VR rendering pipeline; Tracking systems; Locomotion techniques; Haptic feedback; Social and collaborative VR; Intelligent virtual environments.

XR Development, Applications and Emerging Trends: XR development frameworks and OpenXR; Mobile and wearable AR; Visualization and annotation; Applications: healthcare, manufacturing, education, entertainment.

Learning Resources:

Text Books:

1. Jason Jerald, "*The VR Book: Human-Centered Design for Virtual Reality"*, ACM and Morgan & Claypool Publishers, 2015, 1st Edition.
2. Dieter Schmalstieg, Tobias Hollerer, "*Augmented Reality, Principles and Practice*", Addison Wesley, 2016, 1st Edition.
3. Ralf Doerner, Wolfgang Broll, Paul Grimm, Bernhard Jung, "*Virtual and Augmented Reality (VR/AR): Foundations and Methods of Extended Realities (XR)",* Springer, 2022.

Reference Books:

1. Steven M. LaValle, "*Virtual Reality"*, Cambridge University Press, 2023.
2. Jesse Glover and Jonathan Linowes, *"Complete Virtual Reality and Augmented Reality Development with Unity"*, Packt Publishers, 2019.

Online Resources:

1. Unity VR Curricular Framework Resources, https://learn.unity.com/course/vr-curricular-framework-resources
2. Unity XR Manual, https://docs.unity3d.com/6000.3/Documentation/Manual/XR.html

CS36060	AI FOR MALWARE ANALYSIS	Credits 3-0-0: 3

Pre-Requisites: None

Course Outcomes:

At the end of the course, the student will be able to

CO1	Analyse malware behaviour using static and dynamic analysis techniques.
CO2	Extract and engineer meaningful features from malware binaries, memory, and network traces.
CO3	Design and implement machine learning and deep learning models for malware detection and classification.
CO4	Evaluate, interpret, and enhance AI-based malware detection systems considering adversarial behaviour and explainability.

Course Articulation Matrix:

	PO1	PO2	PO3	PO4	PO5	PO6
CO1	2	-	3	1	2	1
CO2	3	-	1	2	1	2
CO3	-	2	3	1	-	1
CO4	-	3	-	1	3	1

1 - Slightly; 2 - Moderately; 3 – Substantially

Syllabus:

Malware Fundamentals and Threat Landscape: Malware taxonomy: virus, worm, trojan, ransomware, spyware, rootkits, Malware life cycle and infection vectors, Modern malware ecosystems (APT, botnets, fileless malware), Static vs dynamic vs hybrid malware analysis, Malware datasets and ethical considerations

Static Malware Analysis Techniques**:** PE file format and ELF binaries, Disassembly and decompilation, Opcode, API call, and string analysis, Control flow graphs (CFG) and call graphs, Feature extraction for AI models, Obfuscation and packing techniques,

Dynamic Malware Analysis and Sandboxing: Behavioral analysis concepts, API tracing, system calls, and registry monitoring, Memory forensics and process injection detection, Network traffic analysis for malware, Sandboxing frameworks and automation, Anti-sandbox and evasion techniques,

Machine Learning for Malware Detection: Malware detection as a classification problem, Feature engineering for malware ML, Supervised learning: SVM, Random Forest, XGBoost, Unsupervised learning: clustering, anomaly detection, Model evaluation metrics (ROC, F1, precision-recall), Dataset imbalance and concept drift.

Deep Learning and Advanced AI Techniques: CNN-based malware classification, RNN/LSTM/Transformer models for sequence-based malware, Graph neural networks (GNN) for CFG-based detection, Representation learning for binaries, Transfer learning for malware analysis, Federated learning for privacy-preserving malware detection.

Adversarial Malware and XAI**:** Adversarial attacks on malware classifiers, Evasion, poisoning, and mimicry attacks, Explainable AI (SHAP, LIME) for malware decisions, Trustworthy AI in cybersecurity, Zero-day malware detection.

Learning Resources:

Text Books:

1. M. Stamp, M. Alazab, and A. Shalaginov, "*Malware Analysis Using Artificial Intelligence and Deep Learning*, "Eds. Cham, Switzerland: Springer, 2021.
2. M. Stamp and M. Jureček, Eds. "*Machine Learning, Deep Learning and AI for Cybersecurity*, "Cham, Switzerland: Springer, 2025.
3. Mohanta and A. Saldanha. "*Malware Analysis and Detection Engineering*", Berkeley, CA, USA: Apress, 2021.

4. M. Yang et al., Eds. Cham, "*Demystifying AI and Machine Learning for Cyber Threat Intelligence,*" Switzerland: Springer, 2023. doi: 10.1007/978-3-031-90723-4.
5. B. Goldmann, *AI Can Hacked: Malware Analysis and Reverse Engineering with AI,* Independently published, 2024.

Reference Books:

1. M. Sikorski and A. Honig. "*Practical Malware Analysis*: The Hands-On Guide to Dissecting Malicious Software,"San Francisco, CA, USA: No Starch Press, 2012.
2. Birmingham, *Data Science for Malware Analysis,* Packt Publishing. , U.K.: Packt, 2025.
3. S. Russell and P. Norvig. "*Artificial Intelligence: A Modern Approach*," Upper Saddle River, NJ, USA: Pearson, 4th ed., 2021.