Gates Institute of Technology
(Autonomous)

Affiliated to

Jawaharlal Nehru Technological University Anantapur
(Established by Govt. of A.P., ACT No.30 of 2008)
Ananthapuramu – 515 002 (A.P) India

==========

B. Tech (Regular-Full time)

(Effective for the students admitted into I B.Tech from the Academic Year 2023-24 onwards and LE admitted students of 2024-25 onwards)

CSE-Artificial Intelligence

Course Structure
&
Syllabus

IV B.Tech I Semester

S.No	Course Code	Title	L	T	P	Credits
1	23AIT13	Generative AI & Prompt Engineering	3	0	0	3
2	23MBTU4a 23MBTU4b 23MBTU4c	**Management Course- II** 1.Business Ethics and Corporate Governance 2.E-Business 3.Management Science	2	0	0	2
3	23AIT14a 23AIT14b 23AIT14c 23AIT14d 23CST23e	**Professional Elective-IV** 1. Explainable AI & Model Interpretability 2. AI for Robotics 3. AI in Cybersecurity 4. AI-driven Software Engineering & DevOps 5. Agentic AI Engineering	3	0	0	3
4	23AIT15a 23AIT15b 23AIT15c 23AIT15d 23CST30d	**Professional Elective-V** 1. AI for Smart Cities & IoT Systems 2. MLOps & AI Model Deployment 3. Data Wrangling 4. Healthcare AI 5. AI Powered Application Development	3	0	0	3
5	23BST33	**Open Elective-III** Employability Skills	3	0	0	3
6	23BST40	**Open Elective-IV** Advanced Aptitude and Critical Thinking	3	0	0	3
7	23AIP06	**Skill Enhancement Course** Prompt Engineering	0	1	2	2
8	23BST39	**Audit Course** Gender Sensitization	2	0	0	-
9	23AIEII	Evaluation of Industry Internship	-	-	-	2
Total			**19**	**1**	**02**	**21**

Open Elective – III

S.No	Course Code	Course Name	Offered by the Dept.
1	23CET24	Building Materials and Services	CIVIL
2	23CET25	Environmental Impact Assessment	
3	23EET22	Smart Grid Technologies	EEE
4	23MET23	3D Printing Technologies	ME
5	23ECT12	Microprocessors and Microcontrollers	ECE
6	23BST30	Wavelet transforms and its applications	Mathematics
7	23BST31	Smart Materials and Devices	Physics
8	23BST34	Introduction to Quantum Mechanics	
9	23BST32	Green Chemistry and Catalysis for Sustainable Environment	Chemistry
10	23BST33	Employability Skills	Humanities

Open Elective – IV

S.No	Course Code	Course Name	Offered by the Dept.
1	23CET26	Geo-Spatial Technologies	CIVIL
2	23CET27	Solid Waste Management	
3	23EET25	Electric Vehicles	EEE
4	23MET24	Total Quality Management	ME
5	23ECT29	Transducers and Sensors	ECE
6	23BST35	Financial Mathematics	Mathematics
7	23BST40	Advanced Aptitude and Critical Thinking	
8	23BST36	Sensors and Actuators for Engineering Applications	Physics
9	23BST37	Chemistry of Nanomaterials and Applications	Chemistry
10	23BST38	Literary Vibes	Humanities

IV B.Tech II Semester

S.No.	Course code	Title	Category	L	T	P	Credits
1	23AIIAP	Internship and Project		-	-	24	4+8=12

IV B.Tech I Semester

23AIT13	GENERATIVE AI & PROMPT ENGINEERING (Professional Core)	L	T	P	C
		3	0	0	3

Course Objectives

- Understand the foundations and working of Generative AI models.
- Explore various generative models like GANs, VAEs, and LLMs.
- Learn prompt engineering techniques to effectively interact with language models.
- Design applications using LLMs with precise control through prompting.
- Understand ethical and societal implications of using Generative AI.

Course Outcomes (COs)

After completing this course, students will be able to:

CO1. Explain the fundamentals of Generative AI and describe how GANs, VAEs, and Transformers are used to create text, images, and other digital content.

CO2. Apply prompt engineering techniques to control and generate accurate outputs from Generative AI models.

CO3. Analyze the architecture and capabilities of Large Language Models (LLMs) and develop NLP applications using prompt engineering and fine-tuning techniques.

CO4. Design complex multi-step prompting workflows using tools like LangChain and LlamaIndex, and generate structured or multimodal outputs safely and effectively.

CO5. Assess the ethical and legal issues of Generative AI and evaluate its responsible use in healthcare, education, and legal applications.

Unit I: Introduction to Generative AI

Overview of Generative AI and Applications, Generative vs Discriminative Models, Latent Space and Data Generation Concepts, Architectures: GANs, VAEs, Autoregressive Models, Generative AI in Text, Image, Audio, and Video, LLMs: Pretrained Transformers as, Generators, Training Challenges and Evaluation of Generative Models.

Unit II: Prompt Engineering Fundamentals

Introduction to Prompt Engineering, Prompt Formats: Zero-shot, One-shot, Few-shot, Prompt Tuning vs Prompt Programming, In-Context Learning & Chain-of-Thought Prompting, Role of Instructions and Examples in Prompts, Controlling Output Style, Tone, and Format, Prompt Failure Cases and Debugging, Prompt Engineering for Coding, Text Completion, Q&A

Unit III: Generative Models in NLP

Transformer Architecture Recap (BERT, GPT), GPT-3/4, PaLM, Claude, and LLaMA Architectures, Text Generation Pipelines and APIs (OpenAI, HuggingFace), Prompt Engineering with GPT Models, Fine-tuning vs Instruct Tuning, Retrieval-Augmented Generation (RAG), Evaluation Metrics: BLEU, ROUGE, Perplexity, Building LLM-based Apps with LangChain.

Unit IV: Advanced Prompt Engineering & Tools

Role of Temperature, Top-k, Top-p Sampling, Structured Outputs: Tables, JSON, Function Calls, Agentic Prompting and Multi-step Reasoning, Prompt Chaining and Memory Handling, Prompt Templates for Automation (LangChain, LlamaIndex), Prompt Engineering for Multimodal Models (DALL·E, Gemini, Sora), Safety Layers & Guardrails in Prompting, AutoGPT, BabyAGI, and Agentic Workflow Building.

Unit V: Ethics, Risks, and Applications of Generative AI

Risks: Hallucination, Toxicity, Bias, Deepfakes and Misinformation Challenges, Copyright, IP, and Data Privacy in Generated Content, Evaluation of Responsible AI Outputs, Red Teaming and Safety Testing, Applications in Education, Medicine, Art, and Law, Regulatory Landscape for Generative AI, Future Trends and Research Directions

Text books

1. "Deep Learning with Python", François Chollet, Manning, 2nd Edition
2. "Generative Deep Learning", David Foster, O'Reilly, 2nd Edition
3. "Building Systems with ChatGPT", Emmanuel Ameisen (O'Reilly Short Reads)
4. "The Art of Prompt Engineering", Nathan Hunter (Free online eBook)

Reference Material

1. Vaswani et al., Attention is All You Need
2. OpenAI Technical Reports on GPT Models
3. Papers from NeurIPS, ACL, ICML related to XAI and LLMs
4. LangChain Documentation

IV B.Tech I Semester

23AIT14a	EXPLAINABLE AI &MODEL INTERPRETABILITY (Professional Elective-IV)	L	T	P	C
		3	0	0	3

Course Objectives:

- To introduce the principles of interpretability and explainability in AI/ML models.
- To analyze the trade-offs between model accuracy and interpretability.
- To explore popular post-hoc and intrinsic explainability techniques.
- To examine fairness, accountability, and transparency in AI systems.
- To develop hands-on skills with interpretability tools and libraries.

Course Outcomes:

After completing this course, students will be able to:

CO1: Understand the need for explainability in modern AI systems.
CO2: Differentiate between black-box and white-box models.
CO3: Apply interpretability techniques such as SHAP, LIME, and PDPs.
CO4: Evaluate the fairness and transparency of AI systems.
CO5: Use explainability tools for model auditing and deployment in high-stakes domains.

UNIT I: Foundations of Explainable AI

Introduction to Explainability and Interpretability, Importance of XAI in Healthcare, Finance, and Law , White-box vs Black-box Models, Desiderata: Fairness, Accountability, Transparency, Human-Centered AI and Trust ,Taxonomy of XAI Techniques (Global vs Local, Post-hoc vs Intrinsic), Regulatory and Ethical Implications (GDPR, AI Bill of Rights), Model Simplicity vs Predictive Power.

UNIT II: Model-Specific Explainability Techniques

Decision Trees and Rule-based Models, Linear Models and Feature Importance, Generalized Additive Models (GAMs), Visualization of Weights and Coefficients, Logistic Regression Coefficient Interpretation, Case Study: Credit Scoring using Transparent Models,Comparison of Interpretable ML Models, Use Cases and Trade-offs.

UNIT III: Model-Agnostic Explainability Techniques

Local Interpretable Model-agnostic Explanations (LIME), SHAP Values (SHapley Additive exPlanations), Partial Dependence Plots (PDPs), Individual Conditional Expectation (ICE) Plots,Anchors and Counterfactual Explanations, Feature Interaction and Permutation Importance, Comparative Analysis of SHAP, LIME, PDP, Model Debugging with XAI.

UNIT IV: Deep Learning Explainability

Visualizing CNNs: Filters, Feature Maps, Saliency Maps and Grad-CAM, Integrated Gradients, Explaining RNNs and LSTM Outputs, Concept Activation Vectors (TCAV), Attention-based Interpretability in Transformers, Explaining Language Models (BERT, GPT) Evaluation of Deep Model Explanations.

UNIT V: Fairness, Bias & Tools for XAI

Fairness Metrics: Demographic Parity, Equal Opportunity, Sources of Bias in Data and Models, Discrimination Detection and Mitigation Strategies, Introduction to AIF360, What-If Tool, Fairlearn, Case Study: Bias in Hiring Algorithms, Explainability in ML Pipelines (MLFlow, Skater), XAI in Federated and Privacy-Preserving AI, Designing Interpretable AI Systems from Scratch.

Text books:

1. Christoph Molnar, ―Interpretable Machine Learning‖, Leanpub.
2. Sameer Singh et al., ―Explainable AI: Interpreting, Explaining and Visualizing Deep Learning‖, Springer.
3. Dan Roth, Zachary Lipton, and Been Kim, ―Explainable AI: Foundations, Developments, Prospects‖, MIT Press (Online forthcoming).

Reference Books:

1. Marco Tulio Ribeiro et al., ―Why Should I Trust You?‖ (LIME) – Research Paper
2. Scott Lundberg et al., ―A Unified Approach to Interpreting Model Predictions‖ (SHAP) – NIPS
3. A. Barredo Arrieta et al., ―Explainable Artificial Intelligence (XAI): Concepts, Taxonomies, Opportunities and Challenges‖, Information Fusion Journal.
4. Zachary C. Lipton, ―The Mythos of Model Interpretability‖ – Communications of the ACM

Online Learning Resources:

- Coursera – Explainable AI with Google Cloud
- Udacity – AI for Everyone by Andrew Ng
- HarvardX – Data Science: Machine Learning Interpretability

IV B.Tech I Semester

23AIT14b	AI for Robotics (Professional Elective-IV)	L	T	P	C
		3	0	0	3

Course Objectives

- Introduce the fundamental principles of robotics and artificial intelligence integration.
- Understand robot perception, localization, mapping, motion planning, and control using AI algorithms.
- Apply machine learning and deep learning techniques in robotic environments.
- Explore the use of reinforcement learning, behavior-based AI, and neural networks in autonomous robots.
- Enable students to build intelligent robots that can perceive, learn, and adapt to dynamic environments.

Course Outcomes

After completing this course, students will be able to:

CO1: Demonstrate an understanding of how AI techniques are applied in robot control and autonomy.
CO2: Apply vision, perception, and sensor fusion techniques for real-time robotic applications.
CO3: Implement path planning and navigation algorithms in dynamic environments.
CO4: Analyze and apply learning-based models such as reinforcement learning and deep neural networks in robotics.
CO5: Evaluate AI-enabled robotic systems based on their performance, efficiency, and adaptability.

UNIT I – Fundamentals of Robotics and AI

Introduction to Robotics: Types and Components, Overview of Artificial Intelligence and Machine Learning, Relationship between Robotics and AI, Sensors and Actuators in Robotics, Embedded Systems and Microcontrollers in Robotics, Architecture of Autonomous Robots, Programming Tools: ROS (Robot Operating System), Python, C++, Applications of AI in Robotics – Overview

UNIT II – Perception and Sensor Fusion

Computer Vision for Robotics: Basics and Techniques, Depth Sensing, RGB-D Cameras, LIDAR, and Ultrasonic Sensors, Feature Extraction and Object Recognition, Kalman Filter and Extended Kalman Filter, Particle Filter and Sensor Fusion Techniques, SLAM (Simultaneous Localization and Mapping) – Concepts, Visual SLAM and LiDAR-based SLAM, 3D Mapping and Scene Reconstruction.

UNIT III – Motion Planning and Navigation

Path Planning Algorithms – Dijkstra, A*, RRT, Obstacle Detection and Avoidance, Robot Kinematics and Dynamics, Trajectory Generation and Optimization, Localization Techniques – GPS, Wi-Fi, Odometry, Autonomous Navigation in Indoor and Outdoor Environments, Multi-Robot Coordination and Swarm Intelligence, Integration of Perception and Planning Systems

UNIT IV – AI Techniques in Robotics

Supervised and Unsupervised Learning for Robotics, Neural Networks and Deep Learning Models for Robot Vision, Reinforcement Learning – Q-Learning and Deep Q Networks, Policy

Gradient and Actor-Critic Methods, Behavior-Based Robotics and Finite State Machines, Imitation Learning and Learning from Demonstration, Transfer Learning for Robotic Tasks, Safety and Generalization in AI Models for Robots

UNIT V – Advanced Applications and Ethical Considerations

Humanoid and Service Robots with AI, AI in Industrial, Healthcare, and Assistive Robotics, Edge AI and Real-Time Inference in Robots, Human-Robot Interaction and Social Intelligence, Autonomous Vehicles and Delivery Drones, AI in Robotics Competitions (RoboCup, DARPA), Ethical Issues in AI-Enabled Robotics, Future Trends: Neuromorphic and Quantum Robotics

Text books

1. "Artificial Intelligence for Robotics" by Robin R. Murphy
2. "Probabilistic Robotics" by Sebastian Thrun, Wolfram Burgard, Dieter Fox
3. "Introduction to Autonomous Robots" by Nikolaus Correll, Bradley Hayes, et al.

Reference Books

1. "Robot Operating System (ROS) for Absolute Beginners" by Lentin Joseph
2. "Modern Robotics: Mechanics, Planning, and Control" by Kevin M. Lynch and Frank C. Park
3. "Learning for Adaptive and Reactive Robot Control" by Aude Billard, Jean-Jacques Slotine
4. IEEE Transactions and Springer Journals on Robotics and Intelligent Systems

Online Courses

1. AI for Robotics – Udacity (by Sebastian Thrun)
2. Modern Robotics: Mechanics, Planning, and Control – Coursera (Northwestern University)
3. Deep Learning for Robotics – edX

IV B.Tech I Semester

23AIT14c	AI IN CYBERSECURITY (Professional Elective-IV)	L	T	P	C
		3	0	0	3

Course Objectives:

- To introduce the fundamental concepts of AI and their applications in cybersecurity.
- To understand AI-driven techniques for threat detection, classification, and mitigation.
- To explore machine learning and deep learning methods used for malware and intrusion detection.
- To equip students with skills in building intelligent security systems.
- To examine ethical, legal, and privacy aspects in AI-driven cybersecurity.

Course Outcomes:

After completing this course, students will be able to:

CO1: Understand AI principles and their relevance in cybersecurity.
CO2: Apply machine learning techniques to detect and respond to threats.
CO3: Analyze security incidents using intelligent tools and models.
CO4: Evaluate and implement AI models for malware detection and anomaly analysis.
CO5: Design AI-based cybersecurity frameworks for real-world scenarios.

UNIT I: Introduction to AI in Cybersecurity

Role of AI in Modern Cybersecurity, Overview of Cyber Threats and Attack Vectors, Fundamentals of Machine Learning for Security, AI vs Traditional Security Techniques, AI-Based Cyber Defense Lifecycle, Threat Intelligence with AI, Cybersecurity Data Types and Challenges, Case Studies of AI-Driven Attacks and Defenses.

UNIT II: Machine Learning for Cyber Threat Detection

Supervised Learning for Intrusion Detection, Unsupervised Learning for Anomaly Detection, Feature Engineering from Network Traffic, Classification Algorithms: SVM, Decision Trees, Random Forests, Clustering Techniques: K-Means, DBSCAN, Ensemble Models and Model Evaluation Metrics, Real-Time Threat Detection Pipelines, Data Imbalance and Adversarial Sampling.

UNIT III: Deep Learning in Cybersecurity

Neural Networks for Threat Classification, CNNs for Malware Detection from Binary Files, RNNs/LSTMs for Sequential Log Analysis, Autoencoders for Anomaly Detection,GANs in Malware Evasion and Defense,Transfer Learning for Threat Signature Extraction, Deep Learning vs Traditional Models: A Comparative Study,Real-World Use Cases and Limitations

UNIT IV: AI for Specific Security Domains

AI for Phishing and Spam Detection, AI in Cloud Security and Edge Devices, Botnet and DDoS Attack Detection, AI-Driven Endpoint Security, Natural Language Processing for Threat Intelligence, Behavioral Biometrics and Fraud Detection, AI in Social Engineering Attack Prevention, Security Information and Event Management (SIEM) with AI.

UNIT V: Challenges, Ethics & Future of AI in Cybersecurity

Explainable AI (XAI) in Cybersecurity, Adversarial Attacks and Defenses in AI Systems, Data Privacy and Federated Learning, Legal and Ethical Issues in AI Security Solutions, AI Model Bias and Fairness in Security Decisions, Securing AI Models Against Manipulation, Building Scalable AI-Powered SOCs, Future Trends: Autonomous Security, AI-Augmented Threat Hunting

Text books:

1. Clarence Chio & David Freeman, ―Machine Learning and Security‖, O'Reilly Media.
2. Sumeet Dua & Xian Du, ―Data Mining and Machine Learning in Cybersecurity‖, CRC Press.
3. Mark Stamp, ―Information Security: Principles and Practice‖, Wiley.

Reference Books:

1. Xiaofeng Chen et al., ―Artificial Intelligence and Big Data Analytics for Cybersecurity‖, Springer.
2. Shai Shalev-Shwartz & Shai Ben-David, ―Understanding Machine Learning‖, Cambridge University Press.
3. Zhiwei Lin & Yang Xiang, ―Cyber Security Intelligence and Analytics‖, Springer.
4. Bhavani Thuraisingham, ―Data Mining for Malware Detection‖, CRC Press.

Online Learning Resources:

- Coursera – ―AI for Cybersecurity‖ by University of Colorado
- Udemy – ―Machine Learning for Cybersecurity‖
- edX – ―Cybersecurity MicroMasters‖ by RIT

IV B.Tech I Semester

23AIT14d	AI-DRIVEN SOFTWARE ENGINEERING & DEVOPS (Professional Elective-IV)	L	T	P	C
		3	0	0	3

Course Objectives:

- To introduce the principles and practices of AI-driven software engineering and DevOps.
- To explore how AI techniques can automate and optimize software development workflows.
- To study intelligent tools for code generation, testing, monitoring, and deployment.
- To equip students with skills in AI-powered CI/CD pipelines and operations.
- To foster an understanding of ethical implications and reliability in intelligent software systems.

Course Outcomes:

After completing this course, students will be able to:

CO1: Understand AI's role in modern software development and operations.
CO2: Apply machine learning techniques to automate software engineering tasks.
CO3: Design and manage intelligent CI/CD and DevOps workflows.
CO4: Evaluate AI tools in software testing, refactoring, and monitoring.
CO5: Implement AI-based solutions for predictive maintenance and decision support in DevOps.

UNIT I: Foundations of AI in Software Engineering

Overview of Traditional vs AI-driven Software Development, AI Opportunities in Software Lifecycle Phases, Introduction to ML/DL Models in Engineering Tasks, Code Representation and Learning from Code, NLP for Source Code Understanding, Software Knowledge Graphs and Reasoning, Datasets and Benchmarks for Software Engineering AI, Case Studies of AI-Enhanced Development Tools

UNIT II: AI in Code Generation and Refactoring

Program Synthesis and Code Completion Models, Large Language Models (e.g., Codex, CodeBERT) in IDEs, Code Clone Detection and Automated Refactoring, Learning-Based Bug Detection and Code Smell Identification, AI in Software Architecture Recommendations, Embedding Techniques for Source Code, Prompt Engineering for Software Tasks, Reliability and Safety in Generated Code

UNIT III: Intelligent Testing, QA, and Debugging

Test Case Generation Using AI, Automated Unit Testing, Regression Testing with ML, Learning Bug Patterns from Repositories, AI-Based Static and Dynamic Code Analysis, Fault Localization and Automated Debugging, Quality Assurance Metrics Enhanced by AI, Reinforcement Learning for Test Prioritization, Ethics and Bias in AI-Driven Testing – (E)

UNIT IV: AI in DevOps Automation and CI/CD

DevOps Fundamentals and Integration with AI, Intelligent CI/CD Pipeline Design, Predictive Build Failure and Log Analysis, AI in Infrastructure-as-Code and Deployment Orchestration, Self-Healing Systems and AIOps Concepts, Log Analytics and Anomaly Detection in Production, AI in Monitoring, Tracing, and Feedback Loops, DevSecOps and AI for Security Automation

UNIT V: Advanced Topics and Ethical Considerations

Explainability and Transparency in AI-Driven Tools, Ethical and Legal Aspects in Automated Engineering, Human-AI Collaboration in Software Teams, Risk Management in Autonomous Code Deployment, AI for Technical Debt Prediction and Management, AI for Developer Productivity Analytics, Research Trends and Challenges in AI for SE, Capstone: Building a Smart DevOps Workflow

Text books:

1. Tim Menzies, Diomidis Spinellis, and Thomas Zimmermann, ―Perspectives on Data Science for Software Engineering‖, Morgan Kaufmann.
2. André van der Hoek, Reid Holmes, ―Software Engineering for Machine Learning‖, Springer.
3. Len Bass, Ingo Weber, Liming Zhu, ―DevOps: A Software Architect's Perspective‖, Addison-Wesley.

Reference Books:

1. Carlos Eduardo Parnin et al., ―AI for Software Engineering: Foundations, Advances, and Trends‖, Springer.
2. Luciano Baresi et al., ―Machine Learning Techniques for Software Quality Evaluation‖, Springer.
3. Gene Kim, Jez Humble, and Nicole Forsgren, ―Accelerate: The Science of Lean Software and DevOps‖, IT Revolution.

Online Learning Resources:

- Coursera – ―AI for Software Engineering‖ by DeepLearning.AI
- edX – ―DevOps for Developers‖ by Microsoft
- GitHub Copilot and OpenAI Codex documentation
- PapersWithCode – AI for Software Engineering benchmarks
- MIT OCW – ―Software Systems‖ and ―DevOps and CI/CD‖
- Udemy – ―AI-Powered DevOps Pipelines and Automation‖
- Google Cloud – AIOps and MLOps tutorials

IV B.Tech I Semester

23CST23e	Agentic AI Engineering	L	T	Pr	C
		3	-	-	3

Prerequisite: Knowledge of Python programming and Artificial Intelligence concepts.

Course Objectives:

- **To provide foundational knowledge of intelligent AI agents, their architectures, and operational principles.**
- **To equip students with the skills to design, develop, and orchestrate AI agents using modern frameworks and platforms.**
- **To enable students to build, deploy, and evaluate agentic AI solutions for real-world applications.**

Course Outcomes:

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

CO1: Explain the principles, architectures, and types of intelligent AI agents.

CO2: Develop AI agents using memory management, tool calling, and agent development platforms.

CO3: Build LLM-powered agents using planning, orchestration, RAG, and workflow automation techniques.

CO4: Analyze the coordination, collaboration, and workflow design of multi-agent systems.

CO5: Analyze agent deployment architectures and emerging trends in Agentic AI.

Unit I: Foundations of AI Agents **Lecture 9 Hrs**

Evolution of AI Agents and Agentic AI, Characteristics and Architecture of Intelligent Agents, Types of AI Agents, Agent Environment and Lifecycle, Reasoning and Decision-Making Concepts, Applications of AI Agents.

Unit II: Agent Design and Development **Lecture 8 Hrs**

Goal-Oriented and Task-Oriented Agents, Agent Memory and Context Management, Tool Calling and Function Calling, Coding Agents using Frameworks and Developing Agents using Low-Code/No-Code Platforms (n8n, Zapier, Make), Designing Domain-Specific Agents, Agent Testing and Evaluation.

Unit III: LLM-Powered Agents and Agent Workflows **Lecture 9 Hrs**

Large Language Models as Agent Brains, Agent Planning and Task Decomposition, Agent Orchestration and Workflow Management, Retrieval-Augmented Generation (RAG) for Agents, Workflow Automation using AI Agents, AI Coding and Knowledge Agents, Human-in-the-Loop Systems.

Unit IV: Multi-Agent Systems and Collaboration **Lecture 8 Hrs**

Fundamentals of Multi-Agent Systems, Agent Coordination and Collaboration, Multi-Agent Workflow Design, Enterprise Workflow Automation, Enterprise AI Agent Applications, Case Studies of Multi-Agent Solutions.

Unit V: Agent Deployment and Emerging Trends **Lecture 8 Hrs**

Agent Deployment Architectures, Monitoring and Maintenance of Agent Systems, Emerging Trends in Agentic AI.

Design and Development of Multi-Agent Systems for Domain-Specific Applications such as Banking, Agriculture, Healthcare.

Text Books

1. Karthik Ramgopal, Designing Autonomous AI Agents, O'Reilly Media, 2024.
2. Paul Iusztin and Maxime Labonne, AI Engineering: Building Applications with Foundation Models, O'Reilly Media, 2025.

Reference Books

1. Valentino Zocca and Ram Sriharsha, Building LLM Powered Applications, Packt Publishing, 2024.
2. Jay Alammar and Maarten Grootendorst, Hands-On Large Language Models, O'Reilly Media, 2024.
3. Michael Wooldridge, An Introduction to MultiAgent Systems, 2nd Edition, Wiley, 2009.
4. Noah Gift and Alfredo Deza, Practical MLOps: Operationalizing Machine Learning Models, 2nd Edition, O'Reilly Media, 2024.

IV B.Tech I Semester

23AIT15a	AI for Smart Cities & IoT Systems (Professional Elective-V)	L	T	P	C
		3	0	0	3

Course Objectives:

1. To understand the foundational concepts of smart cities and IoT architectures integrated with AI technologies.
2. To explore AI-driven solutions for urban mobility, transportation, and traffic management systems.
3. To apply AI and IoT techniques for efficient energy, waste, and water management in smart urban environments.
4. To examine AI applications in smart healthcare, surveillance, and public safety systems.
5. To design, deploy, and evaluate AIoT systems with an understanding of real-time processing, governance, and future challenges.

Course Outcomes:

After completing this course, students will be able to:

CO1: Describe the architecture and components of smart cities and explain how AI and IoT integrate to optimize urban planning and services.

CO2: Apply AI models in transportation systems to improve traffic flow, public mobility, and autonomous vehicle operations.

CO3: Develop AI and IoT solutions for sustainable energy, waste, and water management in smart city ecosystems.

CO4: Analyze AI-based healthcare, surveillance, and emergency response applications, considering privacy and ethical aspects.

CO5: Design and deploy AIoT systems using edge/cloud platforms and evaluate them using appropriate governance and performance metrics.

Unit I: Introduction to AI in Smart Cities and IoT Systems

Smart City Concepts: Components, Infrastructure, and Urban Needs, Overview of IoT and AI Integration, Smart City Frameworks (India, Singapore, EU, etc.), IoT Architecture: Sensing, Network, Processing, and Application Layers, Role of AI in Urban Planning and Resource Optimization, Case Studies on AI in Smart Cities, Edge, Fog, and Cloud Computing Concepts for Smart Systems

Unit II: AI Applications in Smart Transportation and Mobility

Traffic Monitoring and Congestion Prediction using AI, Intelligent Traffic Signal Control using Reinforcement Learning, Autonomous Vehicles and AI Algorithms, Vehicle Detection and License Plate Recognition using CV, Public Transport Optimization using Predictive Analytics, Smart Parking and Navigation Systems, Use of Drones and AI for Traffic Surveillance

Unit III: AI and IoT for Smart Energy, Waste, and Water Management

AI for Smart Grids and Energy Consumption Prediction, Load Balancing and Demand Forecasting using ML, Waste Segregation and Collection Automation using CV, Water Quality Monitoring Systems using IoT Sensors, Leak Detection and Anomaly Detection Models, Smart Metering and Energy Theft Detection, Sustainability and Carbon Monitoring AI Models

Unit IV: Smart Healthcare, Surveillance, and Public Safety
IoT-based Health Monitoring and Alert Systems, Predictive Healthcare and Disease Outbreak Detection, AI for CCTV Surveillance, Crowd Monitoring, and Violence Detection, NLP for Emergency Response and Chatbot Assistance, Smart Ambulance Routing and Response Optimization, COVID-19 Contact Tracing and Monitoring via AI & IoT, Data Privacy, Security & Ethical Issues in Surveillance Systems

Unit V: AIoT System Design, Deployment, and Governance
AI Model Deployment on Edge Devices (Raspberry Pi, Jetson Nano), Smart City Dashboards and Data Visualization, Real-time Streaming and Analytics Platforms (Apache Kafka, Spark), Cloud Integration: AWS IoT, Google Cloud AI, Azure IoT Suite, Governance Frameworks, Data Privacy, and Policy Standards, Evaluation Metrics for Smart City Projects, Future Trends in AIoT and Smart Urban Living

Text Books:

1. Pethuru Raj & Anupama C. Raman, The Internet of Things: Enabling Technologies, Platforms, and Use Cases, CRC Press.
2. Janaka Ekanayake, Smart Grid: Technology and Applications, Wiley.
3. Rajkumar Buyya, Fog and Edge Computing: Principles and Paradigms, Wiley.
4. Adrian McEwen, Hakim Cassimally, Designing the Internet of Things, Wiley.

Reference Books:

1. Mahalik N. P., Sensor Networks and Applications, McGraw Hill.
2. Kim F. Taylor, Urban Artificial Intelligence and Governance, Springer.
3. Dastbaz, J. & Pattinson, C., Smart Cities: Innovation and Sustainability, Springer.
4. Research papers from IEEE Smart Cities, AIoT Journal, and Springer Urban Tech.

Online Courses:

1. Coursera – Smart Cities: Management of Smart Urban Infrastructures (EPFL)
2. edX – Internet of Things (IoT) Program – Curtin University

IV B.Tech I Semester

23AIT15b	ML Ops & AI Model Deployment (Professional Elective-V)	L	T	P	C
		3	0	0	3

Course Objectives:

- To understand the principles and best practices of operationalizing machine learning models in production environments.
- To explore the life cycle of AI model development, deployment, monitoring, and maintenance using modern MLOps frameworks.
- To develop skills in CI/CD for ML, reproducibility, model versioning, and containerization using Docker and Kubernetes.
- To deploy machine learning models using cloud-native services and track their performance using real-time metrics.
- To address scalability, reliability, and ethical considerations in ML model deployment.

Course Outcomes:
After completing this course, students will be able to:

CO1: Illustrate the lifecycle and pipeline components of MLOps and implement basic version control and orchestration for ML workflows.
CO2: Package ML models using appropriate tools and deploy them using Docker and Kubernetes environments with effective resource management.
CO3: Develop and deploy machine learning models as APIs using FastAPI/Flask and configure for real-time or batch inference scenarios.
CO4: Monitor and log ML systems using modern tools and detect data/model drift with strategies for continuous evaluation and feedback.
CO5: Implement end-to-end MLOps solutions using cloud platforms and CI/CD tools, and analyze deployment challenges in real-world use cases.

UNIT I: Introduction to MLOps and Deployment Pipelines

Definition and need of MLOps, ML system lifecycle and pipeline components, DevOps vs. MLOps: key differences, CI/CD for ML projects, Data versioning and model lineage, Introduction to DVC, Git, and MLFlow, Workflow orchestration using Apache Airflow, Automated testing in ML pipelines,

UNIT II: Model Packaging and Environment Management.

Packaging ML models using Pickle, Joblib, ONNX, Python virtual environments, Conda, Pipenv, Introduction to Docker for ML workloads, Building Dockerfiles for ML apps, Using Kubernetes for orchestration, Security, logging, and resource management, Docker Compose and Helm charts for deployment, Hands-on: Containerize and deploy a scikit-learn model

UNIT III: Model Serving and APIs

RESTful API design for ML models, Model deployment using FastAPI and Flask, TensorFlow Serving, TorchServe basics, Introduction to gRPC for ML deployment, Asynchronous inference and batch vs real-time serving, Load testing and benchmarking, Authentication and authorization in model APIs, Deploying models on edge devices

UNIT IV: Monitoring, Logging, and Continuous Evaluation

Importance of monitoring and alerting in MLOps, Data drift and model drift detection, Logging prediction results and metadata, Prometheus, Grafana, and ELK Stack, A/B testing and canary deployments, Shadow deployments and rollback strategies, Feedback loops for continuous learning, Integration with external monitoring tools

UNIT V: Cloud-native MLOps and Case Studies

ML deployment on AWS SageMaker, Azure ML, Google AI Platform, CI/CD using GitHub Actions, Jenkins, and GitLab CI, AutoML and model registry, Real-world case study: End-to-end MLOps pipeline, Challenges and limitations in enterprise ML deployment, Responsible AI in production systems, Future trends in MLOps, Capstone Project Planning

Text Books:

1. Introducing MLOps: How to Scale Machine Learning Projects with DevOps Tools – Mark Treveil, Alok Shukla, O'Reilly Media.
2. Machine Learning Engineering – Andriy Burkov, TrueShelf Publishing.
3. Designing Machine Learning Systems – Chip Huyen, O‘Reilly Media.

Reference Books:

1. Practical MLOps – Noah Gift, O‘Reilly Media
2. Kubeflow for Machine Learning – Trevor Grant et al., O'Reilly
3. Hands-On MLOps: Implement Machine Learning in Production – Munn, Meza, Vohra, Packt Publishing
4. Research papers from arXiv, MLSys Conference, and ICML Industry Track

Online Courses:

1. Coursera – MLOps Specialization by DeepLearning.AI
2. Google Cloud – MLOps: Continuous Delivery and Automation Pipelines
3. Udemy – MLOps: ML Pipelines, CI/CD, and Model Deployment

IV B.Tech I Semester

23AIT15c	Data Wrangling (Professional Elective-V)	L	T	P	C
		3	0	0	3

Course Objectives:

- To introduce the fundamental techniques for acquiring, cleaning, transforming, and manipulating data.
- To enable students to handle real-world messy data for analysis and machine learning.
- To teach efficient use of libraries like Pandas, NumPy, and SQL for data wrangling.
- To promote understanding of handling missing values, outliers, and inconsistent formats.
- To expose students to automation, reproducibility, and workflow design in data preprocessing.

Course Outcomes:

After completing this course, students will be able to:

CO1: Understand and apply core data wrangling techniques.
CO2: Clean, transform, and reshape data using Python and SQL.
CO3: Handle missing values, data inconsistencies, and outliers.
CO4: Merge and join multiple datasets from different sources.
CO5: Automate data pipelines and preprocessing workflows for analytics and ML.

UNIT I: Introduction to Data Wrangling and Data Acquisition

Introduction to Data Wrangling: Importance and Use Cases, Types of Data: Structured, Semi-Structured, Unstructured, Data Acquisition Techniques: APIs, Web Scraping, Reading Data from CSV, Excel, JSON, XML, Using Python libraries: pandas, requests, BeautifulSoup, Working with Databases using SQLAlchemy and pandas, Loading Large Datasets and Chunking, Exploratory Analysis Before Cleaning.

UNIT II: Handling Missing, Noisy, and Inconsistent Data

Identifying and Understanding Missing Data, Techniques for Imputing Missing Values, Handling Inconsistent Data: Dates, Texts, Units, Removing Duplicates and Irrelevant Data, Detecting and Treating Outliers, Normalization and Standardization Techniques, Regular Expressions for Text Cleaning, Visualizing Missing/Outlier Data.

UNIT III: Data Transformation and Feature Engineering

Data Type Conversion and Parsing , Feature Extraction from Text, Dates, and Strings, One-Hot Encoding, Label Encoding, Binning and Discretization, Data Aggregation and Grouping, Pivoting, Melting, and Reshaping Data, Handling Imbalanced Data, Creating Derived Features and Feature Selection.

UNIT IV: Data Integration, Joining, and Workflows

Merging and Joining Datasets (Inner, Outer, Left, Right), Concatenation and Appending DataFrames, Data Consistency and Referential Integrity, Resolving Schema Mismatches, Designing Reusable Data Wrangling Functions, Automating Workflows with Functions and Pipelines, Data Lineage and Documentation, Case Study: End-to-End Data Wrangling Pipeline.

UNIT V: Tools, Libraries, and Case Studies in Data Wrangling

Pandas and NumPy Advanced Techniques, Pyjanitor, Dask, and Polars for Efficient Wrangling, Using OpenRefine for Data Cleaning, SQL vs NoSQL in Data Wrangling, Real-world Wrangling Case Studies (Finance, Healthcare, Retail), Best Practices and Common Pitfalls in Data Wrangling, Reproducibility and Versioning in Data Pipelines, Final Capstone: Build and Evaluate a Clean Dataset for ML/

Text books:

1. M. Heydt – Data Wrangling with pandas, O'Reilly Media.
2. Hadley Wickham – R for Data Science (Data Wrangling Chapters), O'Reilly.
3. J. VanderPlas – Python Data Science Handbook, O'Reilly Media.

Reference Books:

1. Wes McKinney – Python for Data Analysis, O'Reilly.
2. Cathy O'Neil and Rachel Schutt – Doing Data Science, O'Reilly.
3. David Mertz – Cleaning Data for Effective Data Science, Packt.

Online Learning Resources:

1. Data Wrangling with pandas (Datacamp): https://www.datacamp.com/courses/data-manipulation-with-pandas
2. Coursera: Data Wrangling, Analysis and AB Testing with SQL – *https://www.coursera.org/learn/data-wrangling-analysis-abtesting*
3. edX: Data Wrangling with R – https://online.rice.edu/courses/data-wrangling-r

IV B.Tech I Semester

23AIT15d	Healthcare AI (Professional Elective-V)	L	T	P	C
		3	0	0	3

Course Objectives:

- To provide a foundational understanding of AI applications in healthcare.
- To familiarize students with medical data types, preprocessing, and ethical considerations.
- To explore ML and DL algorithms tailored for diagnosis, prognosis, and treatment recommendations.
- To expose students to real-world healthcare systems and AI solutions like predictive modeling, EHRs, and medical imaging.
- To enable students to design, evaluate, and deploy AI-driven healthcare applications.

Course Outcomes:
After completing this course, students will be able to:

CO1: Understand the scope, challenges, and benefits of AI in healthcare.
CO2: Apply data preprocessing and modelling techniques specific to biomedical data.
CO3: Analyze the performance of ML/DL models in clinical contexts.
CO4: Develop AI-driven applications for tasks like disease diagnosis, drug discovery, and patient monitoring.
CO5: Evaluate ethical, legal, and societal implications of AI in healthcare.

UNIT I: Introduction to AI in Healthcare

Overview of Healthcare Systems and Data Ecosystem, AI in Clinical Decision Support Systems (CDSS), Types of Medical Data: EHRs, Imaging, Genomic, Sensor Data, Applications of AI in Diagnosis, Prognosis, and Monitoring, Use Cases: Radiology, Pathology, Oncology, Cardiology, Limitations and Challenges of AI in Healthcare, AI for Telemedicine and Remote Patient Monitoring.

UNIT II: Medical Data Preprocessing and Feature Engineering

Data Cleaning, Imputation, and Normalization for Clinical Data, Handling Missing Values, Outliers, and Bias, Feature Engineering from EHRs and Time-Series Data, Text Mining for Medical Notes using NLP, Encoding Diagnosis and Procedure Codes (ICD, CPT), Temporal Pattern Extraction from Clinical Sequences, Data Privacy, Anonymization, and HIPAA Compliance.

UNIT III: Machine Learning & Deep Learning in Healthcare

Supervised Learning for Risk Prediction and Classification, Unsupervised Learning for Patient Segmentation, Deep Learning for Medical Imaging: CNNs, Transfer Learning, Recurrent Neural Networks for Time-series Clinical Data, Survival Analysis and Time-to-Event Prediction, Model Evaluation Metrics: Sensitivity, Specificity, AUC, Handling Imbalanced Datasets in Healthcare, Interpretability in Medical ML Models (LIME, SHAP)

UNIT IV: Specialized Healthcare AI Applications

AI for Disease Diagnosis: Diabetes, Cancer, Heart Disease, AI in Medical Imaging: X-ray, MRI, CT Scan Analysis, Predictive Modeling for ICU Admission & Mortality Risk, AI in Genomics and Personalized Medicine, Drug Discovery and Repurposing with AI, Chatbots and Virtual Health Assistants, Remote Monitoring using IoT & Wearables + AI.

UNIT V: Ethics, Regulation, and Future Directions in Healthcare AI

Ethical AI in Healthcare: Bias, Fairness, and Accountability, Regulatory Landscape: FDA Approval, CE Marking, Explainable AI and Clinical Trust, Federated Learning for Privacy-Preserving AI, Clinical Trials and AI Decision-Support Tools, Case Studies: Google DeepMind, IBM Watson Health, PathAI, Responsible Deployment of AI in Healthcare Settings.

Text books:

1. Jiang, Fei et al. – Artificial Intelligence in Healthcare: Past, Present and Future.
2. Kevin Frick – Introduction to Healthcare AI.
3. Eric Topol – Deep Medicine: How Artificial Intelligence Can Make Healthcare Human Again.

Reference Books:

1. Mathias Goyen – AI in Medical Imaging.
2. Bertalan Meskó – The Guide to the Future of Medicine: Technology and The Human Touch.
3. Peter Szolovits – Artificial Intelligence in Medicine (Morgan Kaufmann).

Online Learning Resources:

1. Coursera: AI for Medicine Specialization (offered by DeepLearning.AI) https://www.coursera.org/specializations/ai-for-medicine
2. HarvardX: Data Science in Healthcare (edX) **https://online-learning.harvard.edu/course/data-science-healthcare**

IV B.Tech I Semester

23CST30d	AI Powered Application Development	L	T	P	C
		3	-	-	3

Course Objectives:

- **To provide knowledge of AI-powered application development using modern software engineering and Python programming concepts.**
- **To enable students to design and develop full-stack applications integrated with AI models, APIs, and databases.**
- **To equip students with skills to build, deploy, and evaluate AI-powered applications for real-world scenarios.**

Prerequisite: Knowledge of Python programming, databases, and Artificial Intelligence concepts.

Course Outcomes:

After successful completion of the course, students will be able to: CO1: Explain the fundamentals of AI-powered application development.

CO2: Develop full-stack AI applications using frontend, backend, and database technologies.

CO3: Integrate LLMs and SLMs into applications using AI services, APIs, and SDKs.

CO4: Analyze the design and implementation of AI-powered application features.

CO5: Analyze deployment, scaling, and monitoring strategies for AI-powered applications.

Unit I: Application Development (9 Hrs)

Introduction to AI-Powered Applications, Software Development Life Cycle (SDLC), Agile Methodology and Scrum Basics, AI Application Development Lifecycle, Lambda Functions, map() and filter(), Iterators and Generators, Object-Oriented Programming Concepts, Exception Handling, Modules and Packages, Working with JSON Data.

Unit-II: Full-Stack AI Application Development (9 Hrs)

Web Application Architecture, MVC and MVT Design Patterns, Frontend Development using React.js, RESTful API Development, Database Integration using Python and ORM, CRUD Operations, Integration of Frontend and Backend Applications.

Unit III: AI Model Integration for Applications (9 Hrs)

Introduction to Large Language Models (LLMs) and Small Language Models (SLMs), AI Service Integration using APIs, Working with AI SDKs and Frameworks, Integrating LLMs and SLMs into Python Applications, Prompt Engineering Fundamentals, AI Response Processing and Validation.

Unit IV: Building AI-Powered Applications (9 Hrs)

Design of AI-Powered Application Features, Conversational AI and Chatbots, Document Analysis and Knowledge Applications, AI Content Generation Applications AI Agents and Productivity Applications, Case Studies of AI-Powered Applications.

Unit V: Deployment and Scaling (8 Hrs)

Application Deployment Fundamentals, Containerization Concepts, Cloud Deployment Basics, Scaling AI Applications, Logging and Monitoring, Capstone Project Design and Implementation.

Text Books

1. Eric Matthes, Python Crash Course, 3rd Edition, No Starch Press, 2023.
2. William S. Vincent, Django for APIs, 5th Edition, WelcomeToCode, 2024.
3. Paul Iusztin and Maxime Labonne, AI Engineering: Building Applications with Foundation Models, O'Reilly Media, 2025.

Reference Books

1. Valentino Zocca and Ram Sriharsha, Building LLM Powered Applications, Packt Publishing, 2024.
2. Jay Alammar and Maarten Grootendorst, Hands-On Large Language Models, O'Reilly Media, 2024.
3. Simon Holmes, Clive Harber, Tomomi Imura, Noel Hurst, Christian Kaindl, The React Workshop, 2nd Edition, Packt Publishing, 2024.
4. Noah Gift and Alfredo Deza, Practical MLOps: Operationalizing Machine Learning Models, 2nd Edition, O'Reilly Media, 2024.

IV B.Tech I Sem

23AIP06	PROMPT ENGINEERING Skill Enhancement Course	L	T	P	C
		0	1	2	2

Course Objective:

This course delves into prompt engineering principles, strategies, and best practices, a crucial aspect in shaping AI models' behaviour and performance. Understanding Prompt Engineering is a comprehensive course designed to equip learners with the knowledge and skills to effectively generate and utilize prompts in natural language processing (NLP) and machine learning (ML) applications. This course delves into prompt engineering principles, strategies, and best practices, a crucial aspect in shaping AI models' behaviour and performance.

Course Outcomes:

After completing this course, students will be able to:

CO1: Understanding the fundamentals and evolution of prompt engineering.
CO2: Demonstrate proficiency in prompt design by creating closed-ended, open-ended, and role-based prompts
CO3: Learning to probe and stress-test AI models for bias and robustness.
CO4: Applying prompt optimization techniques and performance evaluation methods.
CO5: Mitigating bias and promoting ethical prompting practices in NLP/ML systems.

Module 1: Introduction to Prompt Engineering

Lesson 1: Foundations of Prompt Engineering
- Overview of prompt engineering and its significance in NLP and ML.
- Historical context and evolution of prompt-based approaches.

Module 2: Types of Prompts and Their Applications

Lesson 2: Closed-Ended Prompts
- Under standing and creating prompts for specific answers.
- Applications in question- answering systems.

Lesson 3: Open-Ended Prompts
- Crafting prompts for creative responses.
- Applications in language generation models.

Module 3: Strategies for Effective Prompting

Lesson 4: Probing Prompts
- Designing prompts to reveal model biases.
- Ethical considerations in using probing prompts.

Lesson 5: Adversarial Prompts
- Creating prompts to stress-test models.
- Enhancing robustness through adversarial prompting.

Module 4: Fine-Tuning and Optimizing with Prompts

Lesson 6: Fine-Tuning Models with Prompts
- Techniques for incorporating prompts during model training.
- Balancing prompt influence and generalization.

Lesson 7: Optimizing Prompt Selection
- Methods for selecting optimal prompts for specific tasks.
- Customizing prompts based on model behavior.

Module 5: Evaluation and Bias Mitigation

Lesson 8: Evaluating Prompt Performance
- Metrics and methodologies for assessing model performance with prompts.
- Interpreting and analyzing results.

Lesson 9: Bias Mitigation in Prompt Engineering
- Strategies to identify and address biases introduced by prompts.
- Ensuring fairness and inclusivity in prompt-based models.

Module 6: Real-World Applications and Case Studies

Lesson 10: Case Studies in Prompt Engineering
- *Exploration of successful implementations and challenges in real-world scenarios.*
- *Guest lectures from industry experts sharing their experiences.*

Text books:

1. **"**Prompt Engineering in Action" – *Danny D. Sullivan*
2. "The Art of Prompt Engineering with Chat GPT: A Hands-On Guide" – *Nathan Hunter.*

Reference Books:

1. **"**Prompt Engineering in Practice" – *Michael F. Lewis*
2. "Mastering AI Prompt Engineering: The Ultimate Guide for Chat GPT Users" – *Adriano Damiao*
3. "Writing AI Prompts For Dummies" – *Stephanie Diamond and Jeffrey Allan*
4. "Prompt Engineering Guide" (Online Resource) – *promptingguide.ai*

Online Resources:

https://www.udemy.com/course/understanding-prompt-engineering/?couponCode=NVDINCTA35TRT