Department of Computer Engineering

University of Engineering and Technology Taxila

MS Artificial Intelligence (AI)

Core Courses

Course Code	Course Name
AI-5001	Machine Learning
AI-5002	Programming for AI
AI-5003	Mathematics for AI

Elective Courses

Course Code	Course Name
AI-5004	Big Data Analytics
AI-5005	Image Processing and Analysis
AI-5006	Reinforcement Learning
AI-5007	Design and Analysis of Algorithms
AI-5008	Data Mining
AI-5009	Edge Computing and AI
AI-5010	Data Analytics for Business
AI-5011	Advanced Data Visualization
AI-5012	Advanced Neural Networks and Deep Learning
AI-5013	Natural Language Processing
AI-5014	Computer Vision and Video Analytics
AI-5015	Advanced Pattern Recognition
AI-5016	Ethical issues in Artificial Intelligence
AI-5017	Internet of Things
AI-5018	AI in Healthcare
AI-5019	Intelligent Robotic Systems
AI-5020	AI for Climate Change and Remote Sensing
AI-5021	MLOPs: Foundation and Applications
AI-5022	AI for Cyber Security
AI-5023	Speech and Audio Processing
AI-5024	AI for Agriculture
AI-5025	Smart Cities
AI-5026	Large Language Models
AI-5027	Digital Twins
AI-5028	Generative AI
AI-5029	Bioinformatics

Course Contents Details

1. **Machine Learning**

Introduction to machine learning; probability and statistics for machine learning, Version spaces Algorithm, Candidate elimination algorithm; Supervised Learning: decision trees, Naive Bayes, Artificial Neural Networks, Support Vector Machines; Overfitting, noisy data, and pruning, Measuring Classifier Accuracy; Linear and Logistic regression; Unsupervised Learning: Hierarchical Aglomerative Clustering. k-means partitional clustering; Self-Organizing Maps (SOM) k-Nearest-neighbor algorithm; Semi-supervised learning with EM using labeled and unlabeled data; Reinforcement Learning: Hidden Markov models, Monte Carlo inference Exploration vs. Exploitation Trade-off, Markov Decision Processes; Ensemble Learning: Using committees of multiple hypotheses. Bagging, boosting.

2. Programming for Artificial Intelligence

Paradigms of AI Programming, IDE for the programing, Introduction to Python Programming, Data Types and Data Structures in Python, Control Flow and Looping Constructs, Functions and Modules in Python, Object-Oriented Programming (OOP) in Python, File Handling and Input/Output Operations, Error Handling and Exception Handling in Python, Working with Libraries and Packages in Python, NumPy for Numerical Computing in Python, Pandas for Data Manipulation and Analysis, Matplotlib and Seaborn for Data Visualization, Introduction to Machine Learning with Scikit-learn, Building and Evaluating Machine Learning Models in Python, Introduction to Deep Learning with TensorFlow or PyTorch, Neural Network Architectures and Implementation in Python, Natural Language Processing (NLP) with NLTK or SpaCy, Computer Vision with OpenCV or TensorFlow, Reinforcement Learning Basics with OpenAI Gym.

3. Mathematics for AI

Linear algebra: matrices, vectors, eigenvalues and eigenvectors, diagonalization. Basic probability: Gaussian distribution, covariance, PCA. Calculus and optimization: Recall of the main ideas of calculus for functions of one variable: derivatives, maxima and minima, a few algorithms for determining maxima and minima. Partial derivatives of a function of more than one variable, gradient. Maxima and minima of a function of more than one variable. Constrained maxima and minima, Lagrange multipliers. Convex optimization: Convex sets, convex functions. Subderivatives and subgradients.

4. Big Data Analytics and AI

Introduction to Big Data and AI, Big data storage and processing frameworks (e.g., Hadoop, Spark), Scalable machine learning techniques, Feature engineering for big data, Dimensionality reduction techniques, Model evaluation and validation in big data environments, Real-world applications and case studies of big data analytics, Streaming data analytics, Graph analytics and network analysis, Distributed computing architectures for big data and AI, Anomaly detection in big data streams, Ensemble learning methods, Reinforcement learning for big data environments, Ethics and privacy considerations in big data and AI, Big data visualization techniques, Text mining and sentiment analysis, Predictive modeling with big data and AI, Clustering algorithms for large-scale datasets, Hyperparameter tuning and optimization techniques, Future trends and emerging technologies in big data analytics and AI, Data-driven decision-making, Cloud-based big data and AI services, Collaborative filtering and recommendation systems, Streaming analytics for Internet of Things (IoT) data, Social media analytics.

5. Image Processing and Analysis

Introduction to Image Processing and Analysis, Basics of Digital Image Representation and Enhancement, Image Filtering and Convolution Operations, Image Segmentation Techniques,

Feature Extraction and Representation in Images, Morphological Image Processing, Image Registration and Alignment, Object Detection and Recognition, Image Classification and Clustering Algorithms, Image Restoration and Denoising Methods, Medical Image Processing Techniques, Remote Sensing Image Analysis, Deep Learning for Image Processing, Image Matching, Mathematical Morphology, Image Watermarking and Steganography.

6. **Reinforcement Learning**

RL task formulation (action space, state space, environment definition), Tabular based solutions (dynamic programming, Monte Carlo, temporal-difference) , Function approximation solutions (Deep Q-networks), Policy gradient from basic (REINFORCE) towards advanced topics (proximal policy optimization, deep deterministic policy gradient, etc.) , Model-based reinforcement learning, Imitation learning (behavioral cloning, inverse RL, generative adversarial imitation learning), Meta-learning , Multi-agent learning, partial observable environments.

7. **Design and Analysis of Algorithms**

Introduction: what is an algorithm? Notation for programs, Proof techniques, Basics review: Sets - Functions - Limits - Simple series, Fundamentals: Instances and problems - Elementary operations, Efficiency, Average and worst-case analysis, Examples, Asymptotic notation: Introduction, A notation for "the order of", The omega notation, The theta notation, The conditional asymptotic notation, Analysis of algorithms: Analyzing control structures, Using a barometer, Amortized analysis, Solving recurrences, Data structures: Arrays, stacks and queues, Records and pointers, Lists, graphs, trees and associative tables, Heaps, Disjoint set structures, Greedy algorithms: Making change, General characteristics of Greedy algorithms, Graphs MST - Kruskal's and Prims's algorithms, Graphs: shortest paths, Knapsack problem, Scheduling, Divide - and – Conquer: Multiplying large integers, Binary search, Sorting by: merging and quicksort, Finding the median, Matrix multiplication, Exponentiation, Quick look at cryptography, Dynamic programming: Making change, Principles of optimality, The knapsack problem, Shortest paths - Floyd's algorithm, Chained matrix multiplication, Introduction to probabilistic algorithms - Parallel algorithms, Introduction to computational complexity.

8. **Data Mining**

Introduction to Data Mining, what is data mining? Related technologies - Machine Learning, DBMS, OLAP, Statistics, Data Mining Goals, Stages of the Data Mining Process, Data Mining Techniques, Knowledge Representation Methods, Applications, Example: weather data, Data Warehouse and OLAP, Data Warehouse and DBMS, Multidimensional data model, OLAP operations, Example: loan data set , Data preprocessing, Data cleaning, Data transformation, Data reduction, Discretization and generating concept hierarchies, Installing Weka 3 Data Mining System, Experiments with Weka - filters, discretization, Data mining knowledge representation, Task relevant data, Background knowledge, Interestingness measures, Representing input data and output knowledge, Visualization techniques, Experiments with Weka – visualization, Attribute-oriented analysis, Attribute generalization, Attribute relevance, Class comparison, Statistical measures, Experiments with Weka - using filters and statistics , Data mining algorithms: Association rules, Motivation and terminology , Example: mining weather data, Data mining algorithms: Classification , Data mining algorithms: Prediction, Evaluating what's been learned, Mining real data, Clustering , Advanced techniques, Data Mining software and applications

9. **Edge Computing and AI**

Architecture of edge networks, edge computing technologies, devices and sensors, network infrastructure, communication technologies, applications of edge networks such as IoT, Smart Cities, autonomous vehicles, healthcare, industrial automation, and retail and energy management, common security vulnerabilities as well as the challenges and fundamental concepts in security and privacy, AI-Edge Relational model, Supervised, Unsupervised, Reinforced, Gradient descent, Cross Entropy, Neural Networks, ANN|CNN, RNN, Arduino Nano BLE, Autonomous Vehicle, Load Prediction in Substation, ML for Reservoir Engineering, DL in Construction Industry, TOOLS: Python, TinyML, Scikit Learn, TensorFlow, TensorFlow lite, Weka, Thing speak

10. Data Analytics for Business

Introduction to Data Analytics and Business Intelligence, Basics of Data Collection and Data Cleaning, Exploratory Data Analysis (EDA) Techniques, Descriptive and Inferential Statistics for Business Analysis, Predictive Analytics and Forecasting Models, Machine Learning Algorithms for Business Applications, Customer Segmentation and Market Basket Analysis, Sentiment Analysis and Text Mining for Business Insights, Data Visualization Techniques and Tools, Business Performance Metrics and Key Performance Indicators (KPIs), A/B Testing and Experimental Design in Business, Ethical Considerations and Privacy Concerns in Data Analytics.

11. Advanced Data Visualization

Fundamental concepts of information visualizations, Introduction to visualization (importance, basic approaches, and existing tools), Abstract visualization concepts, human perception, Visualization methodology, Interaction, Design and development of interactive visualization systems, use of visualization software, tools, and systems, Use case studies from various domains, including biomedical engineering, scientific, and business applications, High-dimensional data: machine learning and visualization, Visualizing large graphs and networks. Topological abstraction and summarization for data visualization.

12. Advanced Neural Networks and Deep Learning

Fundamentals of Artificial Neural Networks (ANN), Regularization Techniques for Deep Learning, Overview of Optimization Algorithms for Training Deep Networks, Stochastic Gradient Descent, Adam, Dropout, Initialization; Different types of ANNs: Convolutional networks, RNNs, LSTM, GANs; Adversarial Approaches to ANN; Advanced Topics: Optimization, Hyper-Parameter, Advanced Optimization; simple examples and motivation for deep networks, use of tensorflow. Advantages of Deep Architectures; Architectures for Deep Neural Networks; Convolutional Neural Networks; Properties of CNN representations: Invertibility, Stability, invariance, Deep Unsupervised Learning: Auto-Encoders (Standard, Denoising, Contractive), Variational Auto-Encoders, Adversarial Generative Networks, Maximum Entropy Distributions, Optimization Methods for Deep Neural Networks, Non-convex optimization for deep networks, Deep Reinforcement Learning; Deep Belief Networks, Recent Applications of Deep Learning (case studies).

13. Natural Language Processing

Introduction to background knowledge; Various Applications of Natural Language Processing (NLP); Zipf'slaw; Linguistic Essentials; Corpus-Based Work; Collocations; N-grams Models over Sparse Data; Word Sense Disambiguation; Hidden Markov Model (Word Guessing and HMM. Decoding and Training); Maximum Entropy; Part-of-Speech Tagging; Fundamentals of Natural Language Parsing; Grammar and Parsing; Statistical approaches; Text Alignment and Machine Translation; Information Retrieval and Information Extraction; Deep Learning for NLP; Modern Trends in NLP and Speech Recognition.

14. Computer Vision and Video Analytics

Introduction to Computer Vision and Video Analytics, Image Formation and Preprocessing Techniques, Image Filtering and Enhancement, Image Segmentation and Object Detection, Feature Extraction and Representation, Deep Learning for Computer Vision, Convolutional Neural Networks (CNNs) for Image Classification, Object Recognition and Localization, Object Tracking and Motion Analysis, Video Representation and Processing, Optical Flow and Video Stabilization, Action Recognition and Event Detection, Video Summarization and Abstraction, Video Anomaly Detection and Surveillance, Face Recognition and Biometric Authentication, 3D Computer Vision and Depth Estimation, Multi-modal Fusion for Video Analytics, Applications of Computer Vision and Video Analytics in Robotics, Ethical and Privacy Considerations in Video Analytics, Advanced Topics in Computer Vision Research and Future Directions.

15. Advanced Pattern Recognition

Pattern recognition: basic concepts; Probability theory: conditional probability theory, bayes decision theory; Linear classifiers: the perceptron algorithm, least-squares methods; Nonlinear
classifiers: multilayer perceptron's, back propagation algorithm, decision trees, combinations of classifiers, boosting; Feature selection: data preprocessing, ROC curves, class separability measures,
feature subset selection, Bayesian information criterion; Dimensionality reduction: basis vectors, singular
value decomposition, independent component analysis, kernel PCA, wavelets; Additional features and
template matching: texture, shape and size characterization, fractals, features for audio, Context dependent classification; Clustering: sequential algorithms, hierarchical algorithms, functional optimization-based clustering, graph clustering, learning clustering, clustering high dimensional Data,
Cluster validity measures.

16. Ethical issues in Artificial Intelligence

Introduction to Ethics in Artificial Intelligence, Ethical Theories and Frameworks for AI, Bias and Fairness in AI Algorithms, Transparency and Explainability in AI Systems, Privacy and Data Protection in AI Applications, Accountability and Responsibility in AI Development and Deployment, Impacts of AI on Employment and Socioeconomic Inequality, Ethical Considerations in AI for Healthcare and Biomedicine, Autonomous Systems and Moral Decision Making, AI and Human Rights: Opportunities and Challenges, Cultural and Global Perspectives on AI Ethics, Governance and Regulation of AI Technologies, Ethical Design and Development Practices in AI Projects, Ethical implications of AI in surveillance and law enforcement, Ethical considerations in AI research and publication, Algorithmic accountability and transparency, Ethical implications and responsibilities, Fairness and equity in AI-powered decision-making systems, Ethical considerations in AI for education and learning technologies, Ethical implications of AI-enhanced human augmentation, AI and the future of work: Ethics of automation and job displacement.

17. Internet of Things

Background: Computer Internetworking, Devices: IoT Circuits, IoT in Practice: Automotive IoT, Devices: IoT Device Architecture, Devices: Arduino Programming, Protocols: Radio Frequency Modulation, Protocols: Media Access Control, IoT in Practice: Wireless IoT Infrastructure, Protocols:

Mesh Routing, Protocols: Service Discovery, Infrastructure: Enterprise Infrastructure, Infrastructure: Core Networking, IoT in Practice: Wired IoT Infrastructure, Infrastructure: Networking Devices, IoT in Practice: Cloud IoT Infrastructure , Infrastructure: Physical Infrastructure and Wiring

18. AI in Healthcare

The course "AI in Healthcare" delves into the dynamic intersection of artificial intelligence (AI) and the healthcare industry, exploring the transformative impact of AI technologies on various facets of healthcare delivery. Beginning with an introductory overview, students gain insight into the historical context and recent developments shaping the field. Fundamental concepts of healthcare data, encompassing both structured and unstructured data sources like Electronic Health Records (EHRs), are covered to establish a foundation for understanding AI applications. Machine learning techniques, including supervised, unsupervised, and reinforcement learning, are then examined in the context of medical diagnosis and prognosis. Deep learning methodologies, particularly Convolutional Neural Networks (CNNs), are explored for their applications in medical imaging analysis, such as radiology and pathology. Natural Language Processing (NLP) techniques are introduced for extracting valuable insights from clinical notes and text data, aiding clinical decision support and medical coding. The course further delves into AI's role in drug discovery and development, clinical decision support systems, and personalized medicine, addressing pertinent ethical and regulatory considerations along the way. Through case studies, practical applications, and a capstone project, students gain hands-on experience and insight into the real-world implementation of AI solutions in healthcare, preparing them to navigate future challenges and opportunities in this rapidly evolving field.

19. Intelligent Robotic Systems

Notions of robot and its behavior in a physical environment, Main issues in intelligent robotic systems design, Complex dynamical systems: basic definitions and concepts propaedeutic to robotic topics (notions of phase and state space, trajectory, attractor, bifurcation, phase transition, chaos), Kinematics of robot, Intelligent robotics algorithms, Robot planning, Navigation problems and main solution approaches, Fuzzy logic and fuzzy systems, Behavior trees, Experimental evaluation and parameter tuning of control software for robots: practical guidelines. Associative learning and Reinforcement learning for Robotics, Swarm Robotics.

20. AI for Climate Change and Remote Sensing

Overview of climate change science, Introduction to remote sensing technologies, Importance of AI in climate change research and environmental monitoring, Principles of remote sensing, Remote sensing platforms and sensors, Working with remote sensing data sets, AI Techniques for Climate Analysis, Ensemble learning approaches for climate prediction, change detection methods for monitoring environmental changes, Climate modeling and forecasting, Analysis of satellite imagery for environmental monitoring, Studying the impact of climate change on ecosystems and biodiversity, Analyzing climate data using AI tools and techniques, Predicting extreme weather events, Carbon footprint analysis using AI, Implementing AI algorithms for climate analysis.

21. MLOPs: Foundation and Applications

The "MLOps Foundations and Applications" course provides a comprehensive exploration of Machine Learning Operations (MLOps), focusing on the principles, tools, and best practices essential for deploying, managing, and scaling machine learning models effectively in production environments. Beginning with an introduction to MLOps and its significance in the machine learning lifecycle, students delve into key challenges bridging machine learning development and

operationalization, comparing MLOps practices with traditional software development and DevOps methodologies. Throughout the course, students learn critical concepts such as version control for machine learning artifacts using Git, containerization techniques with Docker and Kubernetes for seamless model deployment and implementing Continuous Integration and Continuous Deployment (CI/CD) pipelines tailored for machine learning workflows. Additionally, the course covers model monitoring and performance management strategies, governance, compliance considerations, and scalability techniques for resource management in distributed ML systems. Through hands-on labs, case studies, and practical projects, students gain valuable experience implementing MLOps pipelines using industry-standard tools and platforms, preparing them to address real-world challenges and capitalize on emerging trends in the field. The course concludes with a capstone project where students apply MLOps principles to solve a complex real-world problem, culminating in a presentation of their findings and insights gleaned throughout the course.

22. AI for Cyber Security

Advancements in Artificial Intelligence (AI) and Machine Learning (ML) have allowed for a surge in adoption of AI & ML solutions to address problems across numerous domains. With this rising reliance on AI & ML in many organizations, it is critical that such systems are protected from malicious activities. This course will discuss AI & ML cybersecurity issues, explore case studies of AI & ML cyber incidents, present AI & ML adversarial techniques, and demonstrate secure design approaches to protect AI & ML systems. With an emphasis on machine learning, the course will focus on secure machine learning systems development approaches and secure machine learning operations (MLOps). Students are expected to have knowledge of fundamental statistics and the ability to program in Python.

23. Speech and Audio Processing

Speech Production–human speech production mechanism, digital models for speech production, Speech perception, Speech Analysis–Time and frequency domain analysis of speech, Linear prediction, Speech compression, Audio processing–characteristics of audio signals, sampling, Audio compression techniques, Standards for audio compression in multimedia applications, MPEG audio encoding and decoding, audio databases and applications. Speech synthesis–text to speech synthesis, letter to sound rules, syntactic analysis, timing, and pitch segmental analysis. Speech recognition. Music Information Retrieval (MIR) Techniques, Speech and Audio Enhancement and Restoration, Acoustic Modeling for Automatic Speech Recognition (ASR).

24. AI for Agriculture

Introduction to AI for Agriculture, Fundamentals of Agriculture and Crop Science, Remote Sensing Techniques for Agricultural Monitoring, Data Collection and Analysis in Precision Agriculture, Crop Yield Prediction using Machine Learning, Disease Detection and Pest Management with AI, Smart Irrigation Systems and Water Management Techniques, Agricultural Robotics and Automation, Crop Health Monitoring and Nutrient Management with AI, Satellite Imaging and UAVs in Agriculture, Climate Prediction and Adaptation Strategies in Agriculture, Blockchain and IoT Applications in Agri-Tech, Ethical Considerations and Policy Implications in AI for Agriculture.

25. Smart Cities

Introduction to Smart Cities, Urbanization Trends and Challenges, Internet of Things (IoT) and Sensor Networks for Smart Cities, Big Data Analytics and Data Management in Urban Environments, Smart Transportation Systems and Traffic Management, Energy Management and Sustainable Infrastructure, Smart Buildings and Intelligent Infrastructure, Urban Planning and Design with AI and

Machine Learning, Citizen Engagement and Participatory Governance in Smart Cities, Environmental Monitoring and Pollution Control, Smart Healthcare and Public Safety Systems, Digital Twin Technologies for Urban Simulation and Planning, Ethical Considerations and Privacy Concerns in Smart City Development, Case Studies in Smart City Implementation.

26. Large Language Models

Understand the evolution of transformer architectures and the historical context behind ChatGPT: Decoder-only models (e.g., GPT-4), encoder-only models (e.g., BERT), encoder-decoder models (e.g., T5). Understand the use cases for working with different machine learning paradigms (supervised, self-supervised, in-context learning). Explain the lifecycle of LLMs: pre-training, fine-tuning, and inference. Know and work with different types of downstream tasks: Text classification, text similarity, search, question-answering, summarization, translation, and named entity recognition. Construct prompts and effective completions with OpenAI APIs. Understand the use cases for working with zero-shot and few-shot in-context learning approaches. Fine-tune and evaluate models in with the Hugging Face Transformers library. Understand the tradeoffs between zero-shot, k-shot, domain/knowledge transfer, in-context learning, and supervised fine-tuning. Choose the right architecture for the downstream task. Understand how to perform knowledge transfer with out-of-distribution datasets. Know and work with custom models with further pre-training. Know how to adapt case studies to new problems involving multilingual retrieval, QA, summarization, sequence labeling, machine translation, and other tasks. Understand how to use embeddings for dense retrieval, recommendations, and clustering. Explain synthetic data generation and negative mining pipelines. Know how to reduce model size via knowledge distillation and quantization.

27. Digital Twins

Introduction to the concept of digital twins, Types of digital twins, Applications of digital twins in manufacturing, Digital twins and IoT, Digital threads and interoperability, Challenges and solutions in implementing digital threads, Introduction to generative agents, Generative agents vs LLMs, Generative agents and agent-based modeling, Generative agent architectures, Frameworks and tools for generative agents (including Microsoft Autogen), Applications of generative agents in industry, Generative agents in interactive environments (NVIDIA Omniverse Digital Twins), Challenges and limitations, Building responsible agents, Integration of generative agents with IoT and digital threads

28. Generative AI

Understanding Generative AI, AI vs. ML vs. GenAI, Machine Learning: Learning by Example, Large Language Models, Generative AI: From Classification to Creation, From Recognition to Understanding, Areas of Greatest Impact: Product Design and Development, Customer Service, Marketing, Retail, Financial Services, Medical, Manufacturing, Movers and Shakers: OpenAI, Microsoft, Google, Amazon, Open-Source Solutions, Major Use Cases: Machine Learning: Object Detection, Labeling, Classification, Sentiment Analysis, Content Generation, Writing, Ideation, Summarization, Image Generation, Music Generation, Prompt Design, Prompts, Enterprise Services: Enterprise Search, Chatbots, Voicebots, Speech-to-Text, Text-to-Speech, Translation, Software Development, Code Generation, Codebots, Foundational Models: OpenAI GPT, Google PaLM, Meta AI Llama, Text Generation Tools: ChatGPT, Gemini, ChatGPT vs. Gemini, OpenAI Playground, Google AI Studio, Microsoft Copilot, Google Duet AI, Prompt Engineering: Asking the Right Questions,

Adding, Context, Image Generation, Image Generation Tools, OpenAI DALL-E, Midjourney, Stable, Diffusion 2, Speech and Music

29. Bioinformatics

Bioinformatics (the use of computing methods for the management and analysis of molecular biology data) has become an integral component of biomolecular sciences, especially genomics and proteomics. This course focuses on the principles and practical use of bioinformatics methods and resources for the analysis of DNA and protein sequences and structures, as well as results from microarray and proteomics, with emphasis on their evolutionary underpinnings and statistical foundations.