Electives for Executive M.Tech AI & DSE

Sl. No.	Subject Code	Elective-I	L	T	P	C
1	EAI 6101	Computational Data Analysis	3	0	0	3
2	EAI 6102	Pattern Recognition	3	0	0	3
3	EAI 6103	Advanced Machine Learning	3	0	0	3

Sl. No.	Subject Code	Elective-II	L	T	P	C
1	EAI 6201	Deep Learning	3	0	0	3
2	EAI 6202	Physics of Neural Network	3	0	0	3
3	EAI 6203	Advanced Cloud Computing	3	0	0	3

Sl. No.	Subject Code	Elective-III	L	T	P	C
1	EAI 6301	Artificial Internet of Things	3	0	0	3
2	EAI 6302	Natural Language Processing	3	0	0	3
3	EAI 6303	Blockchain Technologies: Platforms & Applications	3	0	0	3

Sl. No.	Subject Code	Elective-IV	L	T	P	C
1	EAI 6401	Reinforcement Learning	3	0	0	3
2	EAI 6402	Meta Learning	3	0	0	3
3	EAI 6403	Selective Topics in Generative AI	3	0	0	3

Course Number	EAI 6101
Course Credit	L-T-P-C: 3-0-0-3
Course Title	**Computational Data Analytics**
Learning Mode	Online
Learning Objectives	In this subject, the students will be trained with the knowledge of various data analytics techniques encountered in real life.
Course Description	Current Physical systems/devices are highly complex and fast and operate with high data acquisition and generation capabilities. Data generated from such systems require advanced level of analytics for apprehension and further usage. This course aims to give a broad understanding what is advanced level data analytics techniques and how they play a critical role in analyzing modern day physical systems acquired data.
Course Outline	Introduction, Operation of physical systems and data generation, Complexity, Drawbacks and Challenges in data generation from physical devices. Requirement of advanced data analytics. Foundations of advanced data analytics principles, mathematical models, probabilistic models, optimization models, deep learning and machine learning models. Role of advanced data analytics in data apprehension and compression, curve-based approximation techniques, interpolation techniques, machine learning models for data interpretation. Statistical models to advanced data analytics, data analytics for 2D and 3D data processing and data manipulation, application of advanced data analytics to real life cases, problem solving.
Learning Outcome	• Gain understanding on data generation systems and the role of advanced data analytics. • Apply the Mathematical models of advanced data analytics to real time • Understand the utilities of statistical models and ML models for advanced data analytics.
Assessment Method	Quiz / Assignment / ESE

Suggested Reading

- Probability & Statistics for Engineers & Scientists (9th Edn.), Ronald E. Walpole, Raymond H. Myers, Sharon L. Myers and Keying Ye, Prentice Hall Inc.
- Advanced Data Analysis from an Elementary Point of View Cosma Rohilla Shalizi
- Mining Massive Data Sets, A. Rajaraman and J. Ullman, Cambridge University Press, 2012
- Advances in Complex Data Modeling and Computational Methods in Statistics, Anna Maria Paganoni and Piercesare Secchi, Springer, 2013

Course Number	EAI 6102
Course Credit	L-T-P-C: 3-0-0-3
Course Title	**Pattern Recognition**
Learning Mode	Online
Learning Objectives	This course aims to help the students (a) Understand the advanced topics of pattern recognition, including classification and clustering methods. (b) To understand the advanced topics of feature selection, multi-label classification. (c) Apply advanced pattern recognition algorithms to practical applications in image processing, speech recognition, and data mining.
Course Description	This course on pattern recognition aims to equip students with the advanced topics of classification, clustering, and feature selection. By focusing on advanced topics, students will develop the ability to implement and evaluate various pattern recognition algorithms. Students will enhance their understanding of advanced topics of classification, clustering, statistical methods, and data preprocessing techniques through interactive lectures, exercises, and projects. Upon completion, students will be proficient in designing and applying advanced pattern recognition systems for applications such as image processing, text mining, speech recognition, and data mining, thereby enhancing their analytical and problem-solving capabilities in diverse domains.
Course Outline	Introduction and motivation of advanced pattern recognition Modern Classification Methods, Random fields, Pattern recognition based on multidimensional models Contextual classification, Hidden Markov models, Multi-classifier systems Advanced parameter estimation methods, Advanced Unsupervised classification, Modern methods of feature selection. Data normalization and invariants, Benchmarking. Analysis and synthesis of image information. Applications od pattern recognition in Text Processing and Healthcare.

Learning Outcome	• Mastery of advanced concepts in pattern recognition. • In-depth understanding of various advanced algorithms across different pattern recognition paradigms. • Comprehensive knowledge of advanced aspects of classification, clustering, feature selection, feature extraction, and projection techniques. • Ability to apply advanced pattern recognition algorithms to real-world projects
Assessment Method	Quiz / Assignment / ESE

Textbooks:

1. "Pattern Recognition and Machine Learning" by Christopher M. Bishop, Springer, 2006.
2. "Pattern Classification" by Richard O. Duda, Peter E. Hart, and David G. Stork, Wiley, 2001.
3. "Machine Learning: A Probabilistic Perspective" by Kevin P. Murphy, MIT Press, 2012.
4. "Deep Learning" by Ian Goodfellow, Yoshua Bengio, and Aaron Courville, MIT Press, 2016.
5. "Introduction to Statistical Pattern Recognition" by Keinosuke Fukunaga, Academic Press, 1990.

Course Number	EAI 6103
Course Credit	L-T-P-C: 3-0-0-3
Course Title	**Advanced Machine Learning**
Learning Mode	Online
Learning Objectives	This course aims to help the students to understand the advanced machine learning techniques and its application in various dimensions.
Course Description	This course will concentrate on some advanced topics of machine learning like graphical models, auto-encoders, GANs, time series forecasting, advanced unsupervised classification algorithms, neural architectures for sequence and graph-structured predictions. When appropriate, the techniques will be linked to applications in translation, conversation modeling, and information retrieval.
Course Outline	Mathematics of machine learning, Overview of supervised, unsupervised learning and Multi-task learning Undirected graphical models: Undirected graphical models: overview, representation of probability distribution and conditional independence statement, Factorization, CRFs, Applications to NLP Deep Networks for Sequence Prediction: Encoder-decoder models (case study translation), Attention models, LSTM, Memory Networks Deep Network for Generation: Sequence to Sequence Models, Variational Auto-encoders, Generative Adversarial Networks (GANs), Pointer Generator Networks, Transformer Networks, Learning Representations, Learning representations for text Models for continuous variables: Time series forecasting Modern clustering techniques Recent topics for solving various problems of natural language processing, bioinformatics information retrieval.
Learning Outcome	Students can design and implement advanced machine learning models, such as deep learning, and transfer learning. Students can apply advanced techniques, such as attention mechanisms, generative adversarial networks (GANs), and transformers. Students can analyze and solve challenging machine learning problems, including those involving large datasets, high-dimensional spaces, and complex relationships.
Assessment Method	Quiz / Assignment / ESE

Textbooks:

- Kevin P. Murphy. Machine Learning: A Probabilistic Perspective. MIT Press 2012
- Ian Goodfellow, YoshuaBengio and Aaron Courville. Deep Learning. MIT Press 2016

Course Number	EAI 6201
Course Credit	L-T-P-C: 3-0-0-3
Course Title	**Deep Learning**
Learning Mode	Online
Learning Objectives	This course aims to provide an introductory overview of deep learning and its application varied domains. The course will provide basic understanding of neural networks, mathematical description of it and finally applications of it in multiple domains. A few open source tools will be demonstrated during the course to provide hands-on experience.
Course Description	This course provides an introduction to big data problems and linear algebra, covering essential machine learning techniques like linear regression and classification. It delves into neural networks, gradient-based learning, regularization methods, optimization strategies, and advanced topics such as CNNs, RNNs, and deep reinforcement learning.
Course Outline	● Introduction: Introduction to bigdata problem, overview of linear algebra ● Feature engineering: Basics of machine learning (linear regression, classification) ● Neural network: Deep feed forward network, cost function, activation functions, overfitting, underfitting, Universal approximation theorem ● Gradient based learning: DG, SGD, Backpropagation ● Regularization: L2, L1, L\infinity, drop-out, early stopping, data augmentation, etc. ● Optimization: Multivariable taylor series, momentum, adaptive learning rate, ADAM, Nesternov, AdaGrad, etc. ● CNN: CNN and its application in computer vision ● RNN: RNN, LSTM, GRU and their applications in NLP ● Advanced topics: Autoencoder, Transformer, Deep reinforcement learning
Learning Outcome	● Basic understanding of deep learning and neural networks. ● Problem modeling skill ● Usage of different open source tools / libraries. ● Analysis of large volume of data
Assessment Method	Quiz / Assignment / ESE

Textbooks:

- Ian Goodfellow, Yoshua Bengio and Aaron Courville, “Deep Learning”, Book in preparation for MIT Press, 2016.

Reference books:

- Jerome H. Friedman, Robert Tibshirani, and Trevor Hastie, “The elements of statistical learning”, Springer Series in Statistics, 2009.
- Charu C Aggarwal, “Neural Networks and Deep Learning”, Springer.
- Aston Zhang, Zachary C. Lipton, Mu Li, Alexander J. Smola, "Dive into Deep Learning"
- Iddo Drori, "The Science of Deep Learning", Cambridge University Press
- Simon O. Haykin, "Neural Networks and Learning Machines", Pearson Education India
- Richard S. Sutton, Andrew G. Barto, "Reinforcement Learning: An Introduction", MIT Press
- Christopher M. Bishop, Hugh Bishop, "Deep Learning: Foundations and Concepts", Springer, 2022.
- Simon J. D. Prince, "Understanding Deep Learning", MIT Press 2023.

Course Code	EAI 6202
Course Credit	L-T-P-C: 3-0-0-3
Course Title	**Physics of Neural Network**
Learning Mode	Online
Learning Objectives	This course aims to help the students (a) Comprehend the basic models and structures of neural networks, including the structure and function of the central nervous system, associative memory, and information storage and recall principles. (b) Gain detailed knowledge of various neuron types, such as stochastic and cybernetic neurons, and different network architectures, including layered and perceptron networks.(c) Learn to apply neural network models to practical applications such as time series prediction, game playing (e.g., Backgammon), and protein structure prediction, as well as exploring their use in biomedicine and economics.(d) Delve into advanced topics like pattern recognition, unsupervised learning, evolutionary algorithms, combinatorial optimization, VLSI, specialized networks (e.g., Hopfield networks, Kohonen maps), and advanced learning techniques like back-propagation and solving optimization problems.
Course Description	This course offers a comprehensive exploration of neural networks, encompassing their fundamental models, the structure of the central nervous system, and a brief historical overview. Students will delve into the core principles of associative memory, information storage and recall, and learning mechanisms such as Hebb's rule. The curriculum covers a variety of neuron types, including stochastic and cybernetic neurons, and introduces layered and perceptron network architectures. Throughout the course, students will investigate practical applications of neural networks, ranging from time series prediction to strategic game playing (e.g., Backgammon) and protein structure prediction. The course also highlights the role of neural networks in biomedicine and economics, showcasing their versatility and impact. Advanced topics are thoroughly explored, including pattern recognition, unsupervised learning, and evolutionary algorithms. Students will engage with combinatorial optimization, VLSI design, and specialized network models such as Hopfield networks and Kohonen maps. The course emphasizes the significance of back-propagation, learning functions, and optimization problem-solving techniques. By the end of the course, students will have a deep understanding of neural networks' theoretical foundations and practical applications, equipping them with the skills to leverage these powerful tools in various scientific and industrial domains.

Course Outline	Models of Neural Networks, A Brief History of Neural Network Models, Prediction of the Secondary Structure of Proteins, Associative Memory for Time Sequences.
Learning Outcome	1. Understand the basic concept of PINN 2. Apply the concept of Partial Differential in PINN 3. Analysis of Optimization techniques for PINNs. 4. Demonstrate the practical utility of PINNs in handling complex, real-time applications that require efficient and accurate simulations.
Assessment Method	Quiz / Assignment / ESE

Text book:

1. Müller, B., Reinhardt, J. and Strickland, M.T., 2012. *Neural networks: an introduction*.
2. Peretto, P., 1992. An introduction to the modeling of neural networks (Vol. 2).

Course Number	EAI 6203
Course Credit	L-T-P-C: 3-0-0-3
Course Title	**Advanced Cloud Computing**
Learning Mode	Online
Learning Objectives	This course aims to help the students understand **(a)** how and why cloud systems work and the cloud technologies that manifest these concepts, such as those from Amazon AWS and Microsoft Azure; **(b)** distributed systems concepts like virtualisation, data parallelism, CAP theorem, and performance analysis at scale; **(c)** Big Data programming patterns such as Map-Reduce (Hadoop), Vertex-centric graphs (Giraph), Continuous Dataflows (Storm), and NoSQL storage systems to build Cloud applications; **(d)** Cloud native computing and micro-services
Course Description	This course provides an in-depth understanding of cloud computing, virtualisation, and distributed systems. It covers foundational concepts, advanced techniques, and real-world applications. Students will explore various aspects of cloud infrastructure, virtualisation technologies, distributed algorithms, and cloud-native computing. By the end of the course, students will be equipped with the knowledge and skills to design, implement, and manage cloud-based solutions and distributed systems effectively.
Course Outline	Introduction to Clouds, Virtualization, and Virtual Machines. Network Virtualization and Geo-distributed Clouds. Leader Election in Cloud, Distributed Systems, and Industry Systems. Classical Distributed Algorithms and Industry Systems. Consensus, Paxos, and Recovery in Clouds. Cloud Storage: Key-value Stores/NoSQL Systems and their Use in Industry Systems. Cloud Applications: MapReduce, Spark, and Apache Kafka. Cloud Native Computing and Micro-services.

Learning Outcome	Cloud Computing as a Distributed Systems: Explain and contrast the role of Cloud computing within this space. Cloud Virtualization, Abstractions and Enabling Technologies: Explain virtualisation and their role in elastic computing. Characterise the distinctions between Infrastructure, Platform and Software as a Service (IaaS, PaaS, SaaS) abstractions, and Public and Private Clouds, and analyse their advantages and disadvantages. Programming Patterns for "Big Data" Applications on Cloud: Demonstrate using Map-Reduce, Vertex-Centric and Continuous Dataflow programming models. Application Execution Models on Clouds: Compare synchronous and asynchronous execution patterns. Design and implement Cloud applications that can scale up on a VM and out across multiple VMs. Illustrate the use of NoSQL Cloud storage for information storage. Performance, scalability and consistency on Clouds: Explain the distinctions between Consistency, Availability and Partitioning (CAP theorem), and discuss the types of Cloud applications that exhibit these features.
Assessment Method	Quiz / Assignment / ESE

Textbooks:

- Distributed and Cloud Computing From Parallel Processing to the Internet of Things; Kai Hwang, Jack Dongarra, Geoffrey Fox Publisher: Morgan Kaufmann, Elsevier, 2013.
- Cloud Computing: Principles and Paradigms; Rajkumar Buyya, James Broberg, and Andrzej M. Goscinski Publisher: Wiley, 2011.
- Distributed Algorithms Nancy Lynch Publisher: Morgan Kaufmann, Elsevier, 1996.
- Cloud Computing Bible Barrie Sosinsky Publisher: Wiley, 2011.
- Cloud Computing: Principles, Systems and Applications, Nikos Antonopoulos, Lee Gillam Publisher: Springer, 2012.

Course Number	EAI 6301
Course Credit	3-0-0-3
Course Title	**Artificial Internet of Things**
Learning Mode	Online
Learning Objectives	• Gain a comprehensive understanding of the convergence of Artificial Intelligence (AI) and Internet of Things (IoT), including basic concepts, architectures, and applications. • Learn various AI techniques and their applications in IoT, including machine learning, deep learning, and data analytics. • Develop skills in designing and implementing IoT systems, integrating sensors, and managing data flow. • Understand the processes for collecting, storing, processing, and analyzing IoT data using AI techniques. • Identify and mitigate security risks and privacy concerns in AIoT systems. • Analyze various real-world applications of AIoT in industries such as healthcare, smart cities, agriculture, and manufacturing. • Understand the regulatory and ethical considerations related to AIoT technologies and their deployment.
Course Description	This course provides an in-depth exploration of the convergence of Artificial Intelligence (AI) and the Internet of Things (IoT), known as AIoT. It covers the fundamental principles and technologies of both AI and IoT, demonstrating how they can be integrated to create intelligent, autonomous systems. Students will learn about IoT architecture, AI algorithms, machine learning, data analytics, and the implementation of AI-driven IoT solutions. Through hands-on projects and real-world case studies, students will gain practical experience in developing smart applications for various domains such as smart cities, healthcare, industrial automation, and smart homes.

Course Outline	• Introduction to AIoT, Intersection of AI and IoT, Benefits and challenges of AIoT • Fundamentals of IoT, IoT Architecture and Protocols, Layers of IoT architecture, Communication protocols and standards, IoT Devices and Sensors • Fundamentals of Artificial Intelligence, Machine Learning and Deep Learning, Overview of AI tools and frameworks • AIoT System Architecture, Components and Designing AIoT, Edge Computing in AIoT, Edge vs. cloud computing, AI Models for IoT • Data Management in AIoT, Data Processing and Analysis, Handling large-scale IoT data, Big data technologies and platforms • AIoT Applications and Use Cases: Smart Homes and Buildings, Healthcare and Wearables, Industrial IoT (IIoT), Smart Cities and Transportation • AIoT Platforms and Tools: AI Development Tools, Case Studies of AIoT Solutions, AIoT Project Development, Future Trends and Innovations in AIoT
Learning Outcome	At the end of course, students will have achieved the following learning objectives. • Students should grasp the foundational concepts of AI and IoT, including machine learning algorithms, data analytics, sensor technologies, and network protocols. • Ability to integrate AI algorithms with IoT devices and platforms to create intelligent systems capable of data collection, analysis, and decision-making in real-time. • Proficiency in developing AI-driven IoT applications, including sensor data processing, predictive analytics, anomaly detection, and automation. • Awareness of security challenges and solutions in AIoT systems, including data privacy, authentication, encryption, and intrusion detection. • Knowledge of optimization techniques for AIoT systems to enhance performance, scalability, and energy efficiency.
Assessment Method	Quiz / Assignment / ESE

<u>Suggested Reading</u>

- Olivier Hersent, David Boswarthick, and Omar Elloumi, The Internet of Things: Key Applications and Protocols, Wiley
- Maciej Kranz, Building the Internet of Things: Implement New Business Models, Disrupt Competitors, Transform Your Industry, Wiley
- John Paul Mueller and Luca Massaron, Machine Learning for the Internet of Things: Practical Guide, Packt.

Course Number	EAI 6302
Course Credit	L-T-P-C: 3-0-0-3
Course Title	**Natural Language Processing**
Learning Mode	Online
Learning Objectives	• Define and describe fundamental concepts in NLP, including syntax, semantics, and pragmatics, and their relevance to text aalysis and language modeling. • Apply preprocessing techniques to clean and prepare text data for analysis, such as tokenization, lemmatization, stemming, and stop-word removal. • Utilize methods for feature extraction and representation from text data, including bag-of-words, TF-IDF, and various types of word embeddings (e.g., Word2Vec, GloVe). • Construct and evaluate machine learning and deep learning models for various NLP tasks, such as classification, regression, and sequence labeling, using techniques like Naive Bayes, SVM, RNNs, and Transformers. • Design and implement solutions for practical NLP problems, including sentiment analysis, named entity recognition, text summarization, and machine translation.
Course Description	• Understand foundational concepts and challenges in NLP, including language modeling, parsing, and semantic analysis. • Apply text preprocessing techniques to prepare data for analysis, including tokenization, stemming, and lemmatization. • Utilize various feature extraction and representation methods such as bag-of-words, TF-IDF, and word embeddings. • Develop and evaluate both traditional machine learning models and advanced deep learning models for a range of NLP tasks. • Implement practical solutions for applications such as sentiment analysis, named entity recognition, and text summarization.
Course Outline	Introduction to NLP Simple Word Vector representations: word2vec, GloVe:Distributed Representations of Words and Phrases and their Compositionality, [Efficient Estimation of Word Representations in Vector Space, Advanced word vector representations: language models, GloVe: Global Vectors for Word Representation, PoS tagging and named entity recognition, Language modeling and other tasks, Opinion Mining Parsing, Sentence classification, Machine Translation, Dynamic Memory

	Networks, Question Answering, Natural Language Generation and Summarization, Contextual Word Representations: BERT
Learning Outcome	• Demonstrate a solid understanding of key concepts in NLP, including tokenization, stemming, lemmatization, and part-of-speech tagging. • Apply techniques for text preprocessing and cleaning, including removing stop words, normalizing text, and handling noisy data. • Implement and evaluate methods for feature extraction and representation, such as bag-of-words, TF-IDF, and word embeddings (e.g., Word2Vec, GloVe). • Develop and train various NLP models, including traditional machine learning models (e.g., Naive Bayes, SVM) and deep learning models (e.g., RNNs, LSTMs, Transformers). • Apply techniques for natural language understanding (e.g., named entity recognition, sentiment analysis) and natural language generation (e.g., text summarization, machine translation).
Assessment Method	Quiz / Assignment / ESE

Textbook:

- Dan Jurafsky and James H. Martin.Speech and Language Processing (3rd ed. draft)
- Jacob Eisenstein.Natural Language Processing
- Yoav Goldberg.A Primer on Neural Network Models for Natural Language Processing
- Ian Goodfellow, YoshuaBengio, and Aaron Courville.Deep Learning
- Delip Rao and Brian McMahan.Natural Language Processing with PyTorch (requires Stanford login).
- Michael A. Nielsen.Neural Networks and Deep Learning
- Eugene Charniak. Introduction to Deep Learning

Course Number	EAI 6303
Course Credit	L-T-P-C: 3-0-0-3
Course Title	**Blockchain Technologies: Platforms and Applications**
Learning Mode	Online
Learning Objectives	• Articulate blockchain platforms that show promise in solving complex business problems. • Examine the life cycle of a chain code and its components. • Implement various blockchain-based enterprise applications.
Course Description	This course explores various blockchain platforms and their business applications, focusing on developing and deploying smart contracts on Ethereum using Solidity.
Course Outline	**Module 1 - INTRODUCTION TO BLOCKCHAIN TECHNOLOGIES** • Introduction to Blockchain Technologies • Overview of Blockchain Platforms: Ethereum, Hyperledger Project, IBM Blockchain, Multichain, Hydrachain, Ripple, R3 Corda, BigChainDB, IPFS **Module 2 - ETHEREUM SMART CONTRACTS** • Introduction to Smart Contracts • Solidity Programming Language • Contract Creation and Deployment • Web3js and RPC Protocols • Miners, Transactions, and Blocks in Ethereum • Front-End Development with React and Web3 **Module 3 - HYPERLEDGER FABRIC** • Introduction to Hyperledger Fabric • Fabric Model • Identity Management in Fabric: Membership Service Provider (MSP) • Policies in Fabric • Ledgers in Fabric: World State and Transaction Log • Chaincode in Fabric: Writing and Deploying Smart Contracts • Endorsement Peers and Endorsement Policies in Fabric Module 4 - ADVANCED TOPICS IN BLOCKCHAIN TECHNOLOGIES • Ordering Nodes in Hyperledger Fabric: Solo Ordering Service, Kafka

	• Committing Peers and Anchor Peers in Hyperledger Fabric • Private Data Sharing in Hyperledger Fabric: Sharing Private Data, Private Data Sharing Patterns • Key-level Transaction Access Control and Endorsement in Hyperledger Fabric. • Setting up a Production Network on Hyperledger Fabric.
Learning Outcome	• Demonstrate an understanding of various blockchain platforms and their potential use cases in business. • Develop and deploy smart contracts on the Ethereum platform using Solidity programming language. • Configure and deploy a production network on the Hyperledger Fabric platform.
Assessment Method	Quiz / Assignment / ESE

TEXT BOOKS

- Tom Serres, Bill Wagner and Bettina Warburg, Basics of Blockchain, 2019. ISBN 9781089919441.
- Gaur and Nitin, Hands-On Blockchain with Hyperledger: Building decentralized applications with Hyperledger Fabric, Packt Publishing Ltd, 2018. ISBN 978-17889945.

Course Number	EAI 6401
Course Credit	L-T-P-C: 3-0-0-3
Course Title	**Reinforcement Learning**
Learning Mode	Online
Learning Objectives	This course aims to help the students **(a)** Understand the foundational concepts and mathematical frameworks of reinforcement learning. **(b)** Gain proficiency in key reinforcement learning algorithms, including dynamic programming, Monte Carlo methods, and temporal-difference learning **(c)** Apply deep reinforcement learning techniques to solve complex problems using methods such as deep Q-networks and policy gradient algorithms. **(d)** Explore recent advancements and applications of reinforcement learning, including multi-agent systems and ethical considerations.
Course Description	This specialized course on reinforcement learning aims to give students a deep understanding of the algorithms and methodologies used to train agents to make decisions through trial and error. Students will learn to develop and implement reinforcement learning models by focusing on foundational theories and practical applications. Students will explore key concepts such as Markov decision processes, policy gradients, Q-learning, and deep reinforcement learning through a mix of theoretical lectures, coding exercises, and project-based learning. Upon completion, students will be equipped to design and apply reinforcement learning solutions to complex problems in fields such as robotics, game development, and autonomous systems, enhancing their expertise in this dynamic area of artificial intelligence.

Course Outline	Foundations: Basics of machine learning and reinforcement learning (RL) terminology. Probability Concepts: Axioms of probability, random variables, distributions, and correlation. Markov Decision Process: Introduction to MDPs, Markov property, and Bellman equations. State and Action Value Functions: Concepts of MDP, state, and action value functions. Tabular Methods and Q-networks: Dynamic programming, Monte Carlo, TD learning, and deep Q-networks. Policy Optimization: Policy-based methods, REINFORCE algorithm, and actor-critic methods. Recent Advances and Applications: Meta-learning, multi-agent RL, ethics in RL, and real-world applications.
Learning Outcome	1. Mastery of fundamental principles and mathematical frameworks of reinforcement learning. 2. Proficiency in implementing key reinforcement learning algorithms and techniques. 3. Ability to apply deep reinforcement learning methods to complex, real-world problems. 4. Understanding of recent advancements in reinforcement learning and their ethical implications.
Assessment Method	Quiz / Assignment / ESE

Suggested Reading

- Reinforcement Learning: An Introduction by Richard S. Sutton and Andrew G. Barto, The MIT Press (1 January 1998).
- Deep Reinforcement Learning Hands-On by Maxim Lapan, Packt Publishing Limited (21 June 2018).
- Algorithms for Reinforcement Learning by Csaba Szepesvari, Morgan and Claypool Publishers (2010)
- Deep Reinforcement Learning: Fundamentals, Research and Applications by Hao Dong, Springer Verlag (2020)

Course Number	EAI 6402
Course Credit	L-T-P-C: 3-0-0-3
Course Title	**Meta Learning**
Learning Mode	Online
Learning Objectives	This course aims to help the students (a) Gain a solid understanding of the foundational principles of meta-learning, including model evaluation, basic machine learning concepts, and their limitations. (b) Delve into advanced techniques such as deep learning, transfer learning, and multitask learning, and understand how these methodologies enhance meta-learning capabilities. (c) Develop proficiency in key meta-learning strategies, including model-based, metric-based, and optimization-based approaches, and familiarize yourself with advanced architectures like memory-augmented networks and conditional sequential neural networks (CSNNs). (d) Apply meta-learning techniques to practical applications in various domains, such as computer vision, natural language processing (NLP), reinforcement learning, healthcare, recommendation systems, and climate science, demonstrating the ability to solve complex real-world problems.
Course Description	This comprehensive course provides an in-depth overview of meta-learning, guiding students from foundational principles to advanced techniques. The curriculum begins with the basics of model evaluation, machine learning concepts, and their inherent limitations. Students will then explore advanced topics such as deep learning, transfer learning, and multitask learning, gaining a robust understanding of how these methodologies enhance the capabilities of meta-learning systems.Key meta-learning strategies are thoroughly examined, including model-based, metric-based, and optimization-based approaches. The course features advanced architectures like memory-augmented networks and conditional sequential neural networks (CSNNs), showcasing their roles in improving learning efficiency and effectiveness.Practical applications of meta-learning are highlighted across various fields, including computer vision, natural language processing (NLP), reinforcement learning, healthcare, recommendation systems, and climate science. These examples demonstrate the versatility and power of meta-learning in addressing complex, real-world problems.By the end of the course, students will be equipped with a robust understanding of meta-learning principles and techniques, enabling them to leverage these advanced methodologies to solve intricate problems across diverse domains.

Course Outline	Meta-Learning Basics and Background, Evaluation of Meta learning, Model-Based Meta-Learning Approaches, Metric-Based Meta-Learning Approaches, Optimization-Based Meta-Learning Approaches
Learning Outcome	1. Understand and articulate the foundational principles of meta-learning 2. Apply probabilistic modeling and Bayesian inference to quantify uncertainty and improve model robustness in decision-making processes. 3. Analysis of Optimization-Based Meta-Learning Approaches. 4. Explore and address new challenges in emerging applications
Assessment Method	Quiz / Assignment / ESE

Textbook:

1. Zou, L., 2022. *Meta-learning: Theory, algorithms and applications*.
2. Brazdil, P., Van Rijn, J.N., Soares, C. and Vanschoren, J., 2022. *Metalearning: applications to automated machine learning and data mining* (p. 346).

Course Number	EAI 6403
Course Credit	L-T-P-C: 3-0-0-3
Course Title	**Selective Topics in Generative AI**
Learning Mode	Online
Learning Objectives	• To gain a comprehensive understanding of advanced AI architectures, particularly in the context of Generative AI. • To develop proficiency in implementing and evaluating a variety of Generative AI techniques and models. • To understand the principles and applications of Generative Pre-trained Transformers and other application-specific architectures. • To explore and address ethical considerations and biases in Generative AI, emphasizing the importance of explainability. • To engage with advanced topics and apply knowledge through hands-on projects.
Course Description	This course provides an in-depth exploration of Generative AI (GenAI), focusing on advanced AI architectures such as Generative Adversarial Networks (GANs), Variational Autoencoders (VAEs), and Generative Pre-trained Transformers (GPT). Students will learn about hybrid and emerging models, application-specific architectures, and the ethical considerations and biases in Generative AI. The course includes hands-on projects to design, implement, and evaluate sophisticated generative AI models, emphasizing innovation and practical problem-solving skills.

Course Outline	Introduction to advanced AI, overview of advanced AI architectures and Generative AI (GenAI) Generative Adversarial Network (GAN): various GAN architectures, DCGAN Advanced Variational AutoEncoder (VAE): hierarchical VAEs, Semi-supervised VAE Hybrid and emerging models: Energy-based models, diffusion models, autoregressive and flow-based models, attention mechanism in generative models Generative Pre-trained Transformers (GPT): architectural details and variations. Advanced application-specific architecture: Models for Image-to-Text generation, Text-to-Image generation, Prompt engineering, Multimodality Ethical consideration and bias in Generative AI, Explainability Some advanced topics and project
Learning Outcome	By the end of the course, students will be able to: ● Master various Generative AI architectures, including GANs, VAEs, and emerging models. ● Demonstrate proficiency in implementing and evaluating advanced Generative AI techniques, such as hierarchical VAEs and energy-based models. ● Understand the design principles and applications of Generative Pre-trained Transformers (GPT) and application-specific architectures. ● Analyze and address ethical considerations and biases in Generative AI, emphasizing the importance of explainability. ● Explore advanced topics in Generative AI and apply acquired knowledge through hands-on projects, fostering innovation and practical problem-solving skills.
Assessment Method	Quiz / Assignment / ESE