MARRI LAXMAN REDDY
INSTITUTE OF TECHNOLOGY AND MANAGEMENT
(AN AUTONOMOUS INSTITUTION)
(Approved by AICTE, New Delhi & Affiliated to JNTUH, Hyderabad)
Accredited by NAAC with 'A' Grade & Recognized Under Section 2(f) & 12(B) of the UGC act,1956

Department of Computer Science and Engineering (Data Science)
Minor in Machine Learning
MLRS-R24 Regulations

Course Structure and Syllabus

Applicable From 2024-25 Admitted Batch

V Semester

Course Code	Course Name	L	T	P	Credits
245MDS6701	Foundations of Artificial Intelligence	3	0	0	3
245MDS6702	Principles of Machine Learning	2	0	0	2
245MDS6771	Principles of Machine Learning Laboratory	0	0	2	1

VI Semester

Course Code	Course Name	L	T	P	Credits
246MDS6703	Principles of Deep Learning	3	0	0	3
246MDS6772	Principles of Deep Learning Laboratory	0	0	2	1

VII Semester

Course Code	Course Name	L	T	P	Credits
247MDS6704	Natural Language Processing	3	0	0	3
247MDS6705	Computer Vision and Robotics	3	0	0	3
247MDS6706	Expert Systems				
247MDS6707	Generative AI				
247MDS6708	Speech and Video Processing				

VIII Semester

Course Code	Course Name	L	T	P	Credits
248MDS6773	Mini-Project	0	0	4	2

245MDS6701: FOUNDATIONS OF ARTIFICIAL INTELLIGENCE

L	T	P	C
3	0	0	3

Course Overview:

This course aims to introduce the fundamental concepts of artificial intelligence (AI). Students will develop a broad understanding of AI technologies, their implications, and their potential applications in various fields. The course will emphasize practical examples and real-world case studies to facilitate comprehension and inspire innovative thinking.

Course Objectives: The students will try to learn

- The fundamental concepts and subfields of AI.
- Real-world applications of AI across various industries.
- The knowledge base and application of reasoning
- First-order logic to solve real world problems.
- Recognize the potential of AI to drive innovation and transformation in different domains.

Course Outcomes: After successful completion of the course, students should be able to

- Explain the foundations, history, and core principles of Artificial Intelligence including intelligent agents, search, logic, and planning.
- Apply propositional and first-order logic for knowledge representation, reasoning, and inference mechanisms.
- Evaluate intelligent agents and problem-solving strategies such as uninformed search, informed search, and game-playing algorithms.
- Analyze knowledge bases using logical reasoning, ontologies, and knowledge engineering principles.
- Develop AI-based solutions for planning and decision-making problems using appropriate algorithms and representations.

Module-I **[10]**

Artificial Intelligence: What is AI, Foundations and History of AI. Propositional and first order logic.

Intelligent Agents: Introduction, how Agents Should Act, Structure of Intelligent Agents, Agent programs, Simple reflex agents, Goal based agents, Utility based agents, Environments and Environment programs.

Problem Solving by Search: Problem-Solving Agents, Formulating Problems, Example Problems, Searching for Solutions, Search Strategies (Breadth-first search, Depth-First Search, Bidirectional search).

Module –II **[8]**

Game Playing: Introduction, Games as Search Problems, Perfect Decisions in Two-Person Games, Imperfect Decisions, Alpha-Beta Pruning, Games That Include an Element of Chance, State-of-the- Art Game Programs.

Module –III **[9]**

First-Order Logic: Syntax and Semantics, Extensions and Notational Variations, Using First-Order Logic, Logical Agents for the Wumpus World, A Simple Reflex Agent, Representing Change in the World Building a Knowledge Base: Properties of Good and Bad Knowledge Bases, Knowledge Engineering, The Electronic Circuits Domain, General Ontology, Application: The Grocery Shopping World.

Module –IV **[9]**

Inference in First-Order Logic: Inference Rules Involving Quantifiers, An Example Proof, Generalized Modus Ponens, Forward and Backward Chaining, Resolution: A Complete Inference Procedure, Completeness of resolution.

Module –V **[9]**

Planning: A Simple Planning Agent, From Problem Solving to Planning, Planning in Situation Calculus, Basic Representations for Planning, A Partial-Order Planning Example, A Partial-Order Planning Algorithm, Knowledge Engineering for Planning.

TEXTBOOKS:

1. Artificial Intelligence A Modern Approach, Stuart Russell and Peter Norvig, 3rd Edition, Pearson Education

REFERENCES:

1. Artificial Intelligence, E.Rich and K.Knight, , 3rd Edition, TMH
2. Artificial Intelligence, Patrick Henny Winston, 3rd Edition, Pearson Education.
3. Artificial Intelligence, ShivaniGoel, Pearson Education.
4. Artificial Intelligence and Expert systems – Patterson, Pearson Education

245MDS6702: PRINCIPLES OF MACHINE LEARNING

L T P C
2 0 0 2

Prerequisite: Probability and Statistics, Data Structures

Course Objectives:

- To introduce the fundamental concepts, types, and issues of machine learning, along with the theoretical foundations of supervised and unsupervised learning.
- To develop the ability to design and implement learning systems using models such as perceptrons, multilayer neural networks, decision trees, and support vector machines.
- To explore advanced machine learning techniques including ensemble methods, radial basis functions, and dimensionality reduction approaches.
- To provide insights into evolutionary and probabilistic learning models such as genetic algorithms, Bayesian networks, and hidden Markov models.
- To enable students to apply reinforcement learning and modern optimization strategies for solving real-world machine learning problems effectively.

Course Outcomes: After completion of the course, student should be able to

- Distinguish between, supervised, unsupervised and semi-supervised learning
- Describe the relationship between the brain, neuron models, and learning systems, highlighting their computational perspectives.
- Formulate concept learning tasks and apply hypothesis search techniques such as maximally specific hypothesis and version spaces.
- Implement linear models including linear discriminants, perceptron learning, and linear regression for classification and prediction.
- Evaluate issues and perspectives in machine learning system design by considering challenges such as linear separability and algorithm limitations.

Module – I [10]

Learning: Types of Machine Learning, Supervised Learning, The Brain and the Neuron, design a Learning System – Perspectives and Issues in Machine Learning, Concept Learning Task, Concept Learning as Search – Finding a Maximally Specific Hypothesis, Version Spaces and the Candidate Elimination Algorithm. Linear Regression and Logistic Regression

Module – II [10]

Multi-layer Perceptron, Going Forwards, Going Backwards: Back Propagation Error, Multi-layer Perceptron in Practice Examples of using the MLP, Overview, Deriving Back-Propagation, Support Vector Machines

Module – III [8]

Learning with Trees, Decision Trees, Constructing Decision Trees, Classification and Regression Trees, Nearest Neighbor Methods, Unsupervised Learning: K means Algorithms

Module – IV [10]

Dimensionality Reduction: Linear Discriminant Analysis, Principal Component Analysis, Factor Analysis, Independent Component Analysis, Locally Linear Embedding, Least Squares Optimization

Module – V [10]

Reinforcement Learning: Overview – Getting Lost Example, Markov Chain Monte Carlo Methods, Sampling, Proposal Distribution, Markov Chain Monte Carlo, Graphical Models, Markov Random Fields, Hidden Markov Models, Tracking Methods

TEXT BOOKS:

1. Stephen Marsland, ―Machine Learning — An Algorithmic Perspective, Second Edition, Chapman and Hall/CRC Machine Learning and Pattern Recognition Series.
2. Tom M Mitchell, ―Machine Learning, First Edition, McGraw Hill Education, 2013.

REFERENCE BOOKS:

1. Peter Flach, ―Machine Learning: The Art and Science of Algorithms that Make Sense of Data‖, First Edition, Cambridge University Press, 2012.
2. Jason Bell, ―Machine learning – Hands on for Developers and Technical Professionals‖, First Edition, Wiley, 2014

245MDS6771: Principles of Machine Learning Laboratory

L	T	P	C
0	0	2	1

Prerequisites: Essentials of Problem solving using python Laboratory

Course Objectives:

- To impart knowledge of statistical concepts and their implementation in Python for data analysis.
- To familiarize students with Python libraries such as Statistics, Math, NumPy, SciPy, Pandas, and Matplotlib for data processing and visualization.
- To develop skills in applying supervised learning techniques like Linear Regression, Logistic Regression, Decision Trees, and KNN using sklearn.
- To introduce unsupervised learning techniques such as K-Means Clustering and their applications.
- To enable students to analyze, evaluate, and compare the performance of different machine learning algorithms through mini projects.

Course Outcomes: After completion of the course, students should be able to

- **I**mplement statistical measures and apply Python libraries for data analysis.
- Utilize data processing and visualization tools such as Pandas and Matplotlib effectively.
- Apply supervised machine learning algorithms including Regression, Decision Trees, and KNN using sklearn.
- Demonstrate the use of unsupervised learning techniques such as K-Means clustering.
- Analyze the performance of machine learning models on real-world datasets through mini projects.

List of Experiments

1. Write a python program to compute Central Tendency Measures: Mean, Median, Mode Measure of Dispersion: Variance, Standard Deviation
2. Study of Python Basic Libraries such as Statistics, Math, Numpy and Scipy

3. Study of Python Libraries for ML application such as Pandas and Matplotlib
4. Write a Python program to implement Simple Linear Regression
5. Implementation of Multiple Linear Regression for House Price Prediction using sklearn
6. Implementation of Logistic regression
7. Implementation of Decision tree using sklearn and its parameter tuning
8. Implementation of KNN using sklearn
9. Implementation of support vector Machine.
10. Implementation of K-Means Clustering

TEXT BOOK:

1. Machine Learning – Tom M. Mitchell, - MGH.

REFERENCE BOOK:

1. Machine Learning: An Algorithmic Perspective, Stephen Marshland, Taylor & Francis.

246MDS6703: PRINCIPLES OF DEEP LEARNING

L T P C
3 0 0 3

Course Objectives:

- To introduce the fundamental principles of machine learning, including learning algorithms, bias–variance trade-off, maximum likelihood estimation, and gradient-based optimization.
- To develop an understanding of deep feedforward networks, backpropagation, and the challenges that motivate the use of deep learning architectures.
- To explore advanced regularization and optimization techniques that enhance the performance, robustness, and generalization of deep learning models.
- To provide knowledge of specialized architectures such as convolutional networks, recurrent and recursive neural networks, and their applications in handling structured, sequential, and large-scale data.
- To enable students to apply deep learning methodologies for solving real-world problems in domains such as computer vision, speech recognition, and natural language processing, with appropriate evaluation metrics and practical strategies.

Course Outcomes: After completion of the course, the student should be able to

- Understand fundamental machine learning concepts, including bias–variance trade-off, supervised/unsupervised algorithms, stochastic gradient descent, and backpropagation, for building predictive models.
- Implement regularization and optimization strategies such as dropout, dataset augmentation, adaptive learning rates, and ensemble methods for improving the generalization of deep learning models.
- Analyze convolutional neural networks (CNNs) for processing structured data, including images, and evaluate their performance across different architectures.
- Develop recurrent and recursive neural network architectures for addressing sequence learning tasks in domains such as text, speech, and time-series data.
- Apply deep learning methodologies to real-world applications, including computer vision, natural language processing, and speech recognition.

Module - I [10]

Machine Learning Basics Learning Algorithms, Overfitting and Underfitting, Hyperparameters and Validation Sets, Estimators, Bias and Variance, Maximum Likelihood Estimation, Stochastic Gradient Descent, Building a Machine Learning Algorithm, Motivating Deep Learning Deep Feedforward Networks Learning XOR, Gradient-Based Learning, Hidden Units, Architecture Design, Back-Propagation.

Module - II [10]

Regularization for Deep Learning Parameter Norm Penalties, Regularization and Under Constrained Problems, Dataset Augmentation, Noise Robustness, Early Stopping, Parameter Tying and Parameter Sharing, Sparse Representations, Bagging and Other Ensemble Methods, Dropout, Adversarial Training, Tangent Distance, Tangent Prop, and Manifold Tangent Classifier, Challenges in Neural Network Optimization.

Module- III [8]

Convolutional Networks The Convolution Operation, Motivation, Pooling, Convolution and Pooling as an Infinitely Strong Prior, Variants of the Basic Convolution Function, Structured Outputs, Data Types, Efficient Convolution Algorithms, Random or Unsupervised Features

Module - IV [10]

Recurrent and Recursive Nets Unfolding Computational Graphs, Recurrent Neural Networks, Bidirectional RNNs, Encoder-Decoder Sequence-to-Sequence Architectures, Deep Recurrent Networks, Recursive Neural Networks.

Module- V [10]

Practical Methodology: Performance Metrics, Default Baseline Models, Determining Whether to Gather More Data, Selecting Hyperparameters, Debugging Strategies, Example: Multi-Digit Number Recognition

TEXT BOOK:

1. Deep Learning by Ian Goodfellow, Yoshua Bengio and Aaron Courville, MIT Press.

REFERENCE BOOKS:

1. The Elements of Statistical Learning. Hastie, R. Tibshirani, and J. Friedman, Springer.

2. Probabilistic Graphical Models. Koller, and N. Friedman, MIT Press.

3. Bishop. C.M., Pattern Recognition and Machine Learning, Springer,2006.

4. Yegnanarayana, B., Artificial Neural Networks PHI Learning Pvt. Ltd, 2009.

246MDS6772: PRINCIPLES OF DEEP LEARNING LABORATORY

L	T	P	C
0	0	2	1

COURSE OBJECTIVES:

- To build the foundation of Deep Learning
- To understand how to build the Neural Network
- To enable students to develop successful machine learning concepts
- To discuss fundamentals concepts of deep learning techniques in the area of Natural language Processing
- To build the natural language processing applications using deep learning concepts

COURSE OUTCOMES: After completion of the course, the student should be able to

- Learn the fundamental principles of Deep Learning
- Identify the Deep Learning algorithms for various types of learning tasks in various domains
- Implement Deep Learning algorithms and solve real-world problems
- Apply Deep Learning methods to natural language processing applications
- Train and evaluate sentiment analysis models using RNN layers with LSTM/GRU to classify text data effectively

LIST OF EXPERIMENTS:

1. Setting up the Spyder IDE Environment and Executing a Python Program
2. Installing Keras libraries and making use of them
3. Installing Tensorflow and Pytorch libraries and making use of them
4. Implement Backpropogation algorithm.
5. Applying the Convolution Neural Network on computer vision problems
6. Image classification on MNIST dataset (CNN model with Fully connected layer)
7. Applying the Deep Learning Models in the field of Natural Language Processing
8. Train a sentiment analysis model on IMDB dataset, use RNN layers with LSTM/GRU notes
9. Applying the Autoencoder algorithms for encoding the real-world data
10. Applying Generative Adversial Networks for image generation and unsupervised tasks.

TEXT BOOKS:

1. Deep Learning, Ian Goodfellow, Yoshua Bengio and Aaron Courville, MIT Press
2. The Elements of Statistical Learning, T. Hastie, R. Tibshirani, and J. Friedman,

Springer
3. Probabilistic Graphical Models, Koller and N. Friedman, MIT Press

REFERENCES:

1. Pattern Recognition and Machine Learning, Bishop C. M., Springer, 2006
2. Artificial Neural Networks, Yegnanarayana B., PHI Learning, 2009
3. Matrix Computations, Golub G. H. and Van Loan C. F., JHU Press, 2013
4. Neural Networks: A Classroom Approach, Satish Kumar, Tata McGraw-Hill Education, 2004

247MDS6704: NATURAL LANGUAGE PROCESSING

L T P C
3 0 0 3

COURSE OBJECTIVES:

- To explain text normalization techniques and n-gram language model
- To discuss part of speech methods and naïve bayes classification techniques
- To understand word sense disambiguation techniques and process of building question answering system
- To introduce the concepts of chatbots, dialogue systems, speech recognition systems and text to speech recognition methods
- To develop chatbots and dialogue systems, including speech recognition and synthesis integration

COURSE OUTCOMES: After the completion of the course, the student should be able to

- Apply normalization techniques on a document and evaluate a language model
- Implement parts of speech tagging and classification techniques on the words
- Establish relationships among words of a sentence using word net and also build the question answering system
- Analyze chatbots, dialogue systems, and automatic speech recognition systems
- Design chatbots, dialogue systems, and integrate automatic speech recognition and text-to-speech technologies

Module-I: [10]

Introduction, Regular Expressions, Text Normalization, Edit Distance: Words, Corpora, Text Normalization, Word Normalization, Lemmatization and Stemming, Sentence Segmentation, The Minimum Edit Distance Algorithm.

Module-II: [10]

N-gram Language Models: N-Grams, Evaluating Language Model, Sampling sentences from a language model

Sequence Labeling for Parts of Speech and Named Entities: Part-of-Speech Tagging, Named Entities and Named Entity Tagging

Module-III: [10]

Naive Bayes and Sentiment Classification: Naive Bayes Classifiers, Training the Naive Bayes Classifier, Optimizing for Sentiment Analysis, Naive Bayes as a Language Model, Evaluation: Precision, Recall, F-measure, Test sets and Cross-validation

Word Senses and WordNet: Word Senses, Relations between Senses, WordNet: A Database of Lexical Relations, Word Sense Disambiguation, WSD Algorithm: Contextual Embeddings

Module-IV: [8]

Information Retrieval, IR-based Factoid Question Answering, IR- based QA: Datasets, Entity Linking, Knowledge-based Question Answering, Using Language Models to do QA, Classic QA Models

Module-V: [10]

Chatbots & Dialogue Systems: Properties of Human Conversation, Chatbots, GUS: Simple Frame-based Dialogue Systems, The Dialogue-State Architecture, Evaluating Dialogue Systems, Dialogue System Design,

Automatic Speech Recognition and Text-to-Speech: The Automatic Speech Recognition Task, Feature Extraction for ASR: Log Mel Spectrum, Speech Recognition Architecture

TEXT BOOKS:

1. Speech and Language Processing, Dan Jurafsky and James H. Martin, 3rd Edition, Pearson Publications
2. Natural Language Processing with Python, Steven Bird, Ewan Klein, and Edward Loper, O'Reilly, 2007

REFERENCES:

1. Practical Natural Language Processing: A Comprehensive Guide to Building Real- World NLP Systems, Sowmya Vajjala, Bodhisattwa Majumder, Anuj Gupta, Harshit Surana
2. Foundations of Statistical Natural Language Processing, Christopher Manning and Hinrich Schütze
3. Natural Language Processing in Action- Understanding, Analysing, and Generating Text with Python, Hobson Lane, Cole Howard, Hannes Max Hapke
4. The Handbook of Computational Linguistics and Natural Language Processing (Blackwell Handbooks in Linguistics), 1st Edition

247MDS6705: COMPUTER VISION AND ROBOTICS

L T P C
3 0 0 3

COURSE OBJECTIVES:

- To understand the fundamental concepts related to computer vision
- To get knowledge on boundary tracking techniques
- To learn Hough transform for lines, circle, and ellipse detections
- To understand the geometry of multiple views
- To Implement geometric camera models and calibration techniques for real-world computer vision applications

COURSE OUTCOMES: After completion of the course, the student should be able to

- Understand fundamental image processing techniques required for computer vision
- Implement boundary tracking techniques
- Apply chain codes and other region descriptors, Hough Transform for line, circle, and ellipse detections
- 3D vision techniques and Implement motion related techniques. And develop applications using computer vision techniques
- To delve into geometric camera models, calibration, and model-based vision for practical applications like medical imaging and robot localization

Module-I: [10]

Radiometry – Measuring Light: Light in Space, Light Surfaces, Important Special Cases. Sources, Shadows, And Shading: Qualitative Radiometry, Sources and Their Effects, Local Shading Models, Application: Photometric Stereo, Interreflections: Global Shading Models.
Color: The Physics of Color, Human Color Perception, Representing Color, A Model for Image Color, Surface Color from Image Color.

Module-II: [10]

Linear Filters: Linear Filters and Convolution, Shift Invariant Linear Systems, Spatial Frequency and Fourier Transforms, Sampling and Aliasing, Filters as Templates.
Edge Detection: Noise, Estimating Derivatives, Detecting Edges.

Texture: Representing Texture, Analysis (and Synthesis) Using Oriented Pyramids,
Application: Synthesis by Sampling Local Models, Shape from Texture.

Module-III: [10]

The Geometry of Multiple Views: Two Views

Stereopsis: Reconstruction, Human Stereopsis, Binocular Fusion, Using More Cameras. Segmentation by Clustering: What Is Segmentation? Human Vision: Grouping and Getstalt,
Applications: Shot Boundary Detection and Background Subtraction, Image Segmentation by Clustering Pixels, Segmentation by Graph-Theoretic Clustering,

Module-IV: [10]

Segmentation by Fitting a Model: The Hough Transform, Fitting Lines, Fitting Curves, Fitting as a Probabilistic Inference Problem, Robustness.

Segmentation and Fitting Using Probabilistic Methods: Missing Data Problems, Fitting, and Segmentation, The EM Algorithm in Practice.

Tracking With Linear Dynamic Models: Tracking as an Abstract Inference Problem, Linear Dynamic Models, Kalman Filtering, Data Association, Applications and Examples.

Module-V: [8]

Geometric Camera Models: Elements of Analytical Euclidean Geometry, Camera Parameters and the Perspective Projection, Affine Cameras and Affine Projection Equations
Geometric Camera Calibration: Least-Squares Parameter Estimation, A Linear Approach to Camera Calibration, Taking Radial Distortion into Account, Analytical **Photogrammetry, An Application:** Mobile Robot Localization

Model-Based Vision: Initial Assumptions, Obtaining Hypotheses by Pose Consistency, Obtaining Hypotheses by pose Clustering, Obtaining Hypotheses Using Invariants, Verification, and Application: Registration in Medical Imaging Systems, Curved Surfaces and Alignment.

TEXT BOOK:

1. Computer Vision – A Modern Approach, David A. Forsyth and Jean Ponce, Pearson Education, 2015

REFERENCES:

1. Computer and Machine Vision – Theory, Algorithms and Practicalities, E. R. Davies, 4th Edition, Elsevier (Academic Press), 2013
2. Digital Image Processing, R. C. Gonzalez and R. E. Woods, Addison Wesley, 2008

247MDS6706: EXPERT SYSTEMS

L T P C
3 0 0 3

Course Outcomes: After completion of the course, the students should be able to

- Apply the principles of LISP and PROLOG programming to solve problems involving symbolic manipulation, logical reasoning, and knowledge-based representation.
- Implement various AI search strategies such as breadth-first, depth-first, heuristic-based, and game-playing algorithms with alpha–beta pruning.
- Analyze knowledge using predicate logic, semantic networks, frames, inheritance, and rule-based deduction systems.
- Design expert systems by understanding their architecture, characteristics, and processes for knowledge acquisition and representation.
- Utilize expert system tools and techniques for developing simple knowledge-based applications, applying concepts of knowledge engineering and system-building aids.

Module-I [10]

Lisp and Prolog Languages: **PROLOG**: The declarative semantics, procedural semantics and the interpreter. Reading in programs, Input and Arithmetical predicates.

LISP: Fundamental principles of LISP, LISP expression, The form, Procedural abstraction in LISP, Variables and their scopes ,Symbol manipulation , Control structures, The structure, Input and output

Module-II [10]

Artificial Intelligence: Blind search strategies, Breadth-first – Depth-first – Heuristic search techniques Hill Climbing – Best first – A Algorithms AO* algorithm – game tress, Min-max algorithms, game playing – Alpha-beta pruning.

Module–III [10]

Knowledge representation: issues predicate logic – logic programming Semantic nets-frames and inheritance, constraint propagation; Representing Knowledge using rules, Rules-based deduction systems.

Module–IV [10]

Expert Systems: Architecture of expert system, Representation and organization of knowledge, Basics characteristics, and types of problems handled by expert systems.

Building an Expert System: Expert system development, Selection of the tool, Acquiring Knowledge, Building process.

Module-V [8]

Expert System Tools: Techniques of knowledge representations in expert systems, knowledge engineering, system-building aids, support facilities, stages in the development of expert systems.

TEXT BOOKS:

1. Principles of Expert Systems Peter J.F. Lucas & Linda C. van der Gaag.

2. Elain Rich and Kevin Knight, "Artificial Intelligence", Tata McGraw-Hill, New Delhi,

REFERENCES:

1. Stuart Russel and other Peter Norvig, "Artificial Intelligence – A Modern Approach", Prentice-Hall,

2. Patrick Henry Winston, "Artificial Intelligence", Addison Wesley,

3. Patterson, Artificial Intelligence & Expert System, Prentice Hall India,1999.

247MDS6707: GENERATIVE AI

L T P C
3 0 0 3

COURSE OBJECTIVES:

- To familiarize with generative models and their role in AI
- To develop expertise in text generation techniques and frameworks
- To master image generation with advanced techniques
- To apply generative models to create Art, Music, and Interactive Systems
- To gain Proficiency with Open-Source Generative Tools and Deployment

COURSE OUTCOMES: After completion of the course, the student should be able to

- Understand foundational concepts of Generative AI
- Apply generative models for natural language processing (NLP)
- Analyze image generation techniques using GANs, VAEs, stable diffusion, and transformer-based models (CLIP, ViT, DALL-E, GPT-4V).
- Experiment with generative models in creative domains
- Develop generative AI solutions

Module-I: [10]

Introduction to GEN AI: Historical Overview of Generative modeling - Difference between Gen AI and Discriminative Modeling – Importance of generative models in AI and Machine Learning – Types of Generative models – GANs, VAEs, autoregressive models and Vector quantized Diffusion models - Understanding if probabilistic modeling and generative process - Challenges of Generative Modeling – Future of Gen AI – Ethical Aspects of AI – Responsible AI – Use Cases.

Module-II: [10]

Generative Models for Text: Language Models Basics – Building blocks of Language models - Transformer Architecture – Encoder and Decoder – Attention mechanisms - Generation of Text – Models like BERT and GPT models – Generation of Text - Autoencoding – Regression Models – Exploring ChatGPT – Prompt Engineering – Designing Prompts– Revising Prompts using Reinforcement Learning from Human Feedback (RLHF) - Retrieval Augmented Generation – Multimodal LLM – Issues of LLM like hallucination.

Module-III: [10]

Generation of Images: Introduction to Generative Adversarial Networks – Adversarial Training Process – Nash Equilibrium – Variational Autoencoders – Encoder-Decoder Architectures - Stable Diffusion Models – Introduction to Transformer-based Image Generation – CLIP – Visual Transformers ViT- Dall-E2 and Dall-E3, GPT-4V – Issues of Image Generation models like Mode Collapse and Stability.

Module-IV: [8]

Generation of Painting Music and Play: Variants of GAN – Types of GAN - Cyclic GAN – Using Cyclic GAN to Generate Paintings – Neural Style Transfer – Style Transfer - Music Generating RNN – MuseGAN – Autonomous agents – Deep Q Algorithm – Actor-critic Network.

Module-V: [10]

OpenSource Models and Programming Frame Works: Training and Fine tuning of Generative models – GPT4All - Transfer learning and Pretrained models - Training vision models – Google Copilot - Programming LLM – LangChain – Open Source Models – Llama - Programming for TimeSformer – Deployment – Hugging Face

TEXT BOOK:

1. Transformers for Natural Language Processing and Computer Vision, Denis Rothman, 3rd Edition, Packt Books, 2024

REFERENCES:

1. Generative Deep Learning, David Foster, O'Reily, 2024 2. Generative AI for Everyone, Altaf Rehmani, BlueRose One, 2024

247MDS6708: SPEECH AND VIDEO PROCESSING

L T P C
3 0 0 3

COURSE OBJECTIVES:

- To understand the mechanisms of human speech production system
- To learn the concepts related to speech analysis and speech recognition algorithms
- To get knowledge on digital video processing
- To familiarize with object tracking and video sequence concepts
- To focus on object tracking, video segmentation methods, and compression techniques for efficient video processing

COURSE OUTCOMES: After completion of the course, the student should be able to

- Describe the mechanisms of human speech production systems and methods for speech feature extraction
- Understand basic algorithms of speech analysis and speech recognition
- Explain basic techniques in digital video processing, including imaging characteristics and sensors
- Apply motion estimation and object tracking algorithms on video sequence
- Develop skills in object tracking, video segmentation, and interframe compression for real-time video applications

Module-I: [10]

Speech processing concepts: The speech production mechanism, Discrete time speech signals, Pole-Zero modeling of speech, relevant properties of the fast Fourier transform for speech recognition, convolution, linear and nonlinear filter banks, spectral estimation of speech using DFT. Linear Prediction analysis of speech.

Module-II: [10]

Speech recognition: Real and Complex Cepstrum, application of cepstral analysis to speech signal, feature extraction for speech, static and dynamic feature for speech recognition, robustness issues, discrimination in the feature space, feature selection, MFCC, LPCC, Distance measures, vector quantization models. Gaussian Mixture model, HMM.

Module-III: [10]

Basics of Video Processing: Video formation, perception and representation: Principle of color video, video cameras, video display, pinhole model, CAHV model, Camera motion, Shape model, motion model, Scene model, two-dimensional motion models. Three-Dimensional Rigid Motion, Approximation of projective mapping.

Module-IV: [10]

Motion Estimation Techniques: Optical flow, motion representation, motion estimation criteria, optimization methods, pixel-based motion estimation.Block matching algorithm, gradient Based, Intensity matching, feature matching, frequency domain motion estimation, Depth from motion.

Motion analysis applications: Video Summarization, video surveillance.

Module-V: [8]

Object Tracking and Segmentation: 2D and 3D video tracking, blob tracking, kernel based counter tracking, feature matching, filtering Mosaicing, video segmentation, mean shift based, active shape model, video short boundary detection. Interframe compression, Motion compensation.

TEXT BOOKS:

1. Fundamentals of Speech Recognition, L. Rabiner and B. Juang, Prentice Hall
2. Digital Video Processing, A. Murat Tekalp, Prentice Hall
3. Discrete-Time Speech Signal Processing: Principles and Practice, Thomas F. Quatieri, Coth