

10. Classify movie reviews as positive or negative from the IMDB movie dataset of 50K movie reviews. (Link for dataset: <https://www.kaggle.com/datasets/lakshmi25npathi/imdb-dataset-of-50k-movie-reviews>)
11. Build and train a text classifier for the given data (using textblob or simpletransformers library)
12. Generate text using a character-based RNN using an appropriate dataset. Given a sequence of characters from a given data (eg "Shakespear"), train a model to predict the next character in the sequence ("e").

DISCIPLINE SPECIFIC CORE COURSE – 20: AdvancedAI Systems Design

CREDIT DISTRIBUTION, ELIGIBILITY AND PRE-REQUISITES OF THE COURSE

Course title & Code	Credits	Credit distribution of the course			Eligibility criteria	Pre-requisite of the course (if any)
		Lecture	Tutorial	Practical/Practice		
AdvancedAI Systems Design	4	3	0	1	Class XII pass with Mathematics	ML , DL , NLP

Learning Objectives:

1. To build AI systems by applying machine learning, deep learning, and NLP concepts learned previously.
2. To understand Artificial General Intelligence (AGI) and its relationship with current AI technologies.
3. To develop practical AI applications using Python, APIs, and integration frameworks.

Learning Outcome:

After completing the course, students will be able to:

1. Apply ML/DL/NLP knowledge to build intelligent AI systems.
2. Explain AGI concepts and current AI limitations.
3. Develop AI applications using APIs and multimodal capabilities.
4. Understand ethical, social, and safety implications of AI systems.

UNIT-I

(10 hours)

Introduction to AI and Intelligent Architectures: Definition of AI and AGI, narrow AI vs general AI, applying ML/DL models in AI systems, symbolic AI and expert systems, search algorithms (BFS, DFS, A*), knowledge representation, cognitive architectures (SOAR, ACT-R), intelligent agents and multi-agent systems, combining supervised/unsupervised learning with

reasoning systems.

UNIT-II

(15 hours)

Building AI Applications using APIs and Frameworks: Architecture of AI systems, integrating pre-trained ML/DL models, using LLM APIs (OpenAI GPT, Claude, Gemini), computer vision APIs (YOLO, MediaPipe), applying NLP models in chatbots and virtual assistants, multimodal AI (text + image + speech), prompt engineering and API chaining, building intelligent workflows using LangChain, creating AI applications with Streamlit/Gradio.

UNIT-III

(10 hours)

AI Ethics and Responsible AI: Bias in ML models and mitigation strategies, fairness in AI systems, explainable AI and model interpretability, privacy concerns in AI applications, ethical use of generative AI, AI safety and alignment, legal implications of automated decision-making, responsible AI development practices.

UNIT-IV

(10 hours)

Current AI Trends and Applications: Overview of major AI companies and their products, generative AI applications (GPT, DALL-E, Stable Diffusion), AI in computer vision (object detection, face recognition), conversational AI and chatbots, AI in robotics and autonomous systems, reinforcement learning applications (gaming, robotics), future of AI in real-world applications and education.

References

1. Russell & Norvig – Artificial Intelligence: A Modern Approach
2. Melanie Mitchell – Artificial Intelligence: A Guide for Thinking Humans
3. Marcus & Davis – Rebooting AI
4. OpenAI, DeepMind, and SingularityNET research papers and blogs
5. Documentation from OpenCog, LangChain, GPT, and Hugging Face

List of Practicals:

1. Build a rule-based expert system using Python
2. Create an intelligent chatbot using LLM APIs
3. Develop a computer vision application using pre-trained models
4. Implement multi-agent system simulation
5. Build a multimodal AI application (text + image)
6. Create an AI web application using Streamlit
7. Test bias detection and mitigation in ML models
8. Develop a recommendation system using collaborative filtering