

CSE473: LARGE LANGUAGE MODELS AND AGENTIC AI

L:2 T:0 P:2 Credits:3

Course Outcomes: Through this course students should be able to

CO1 :: understand the foundations, architecture, and functioning of transformer-based LLMs and intelligent agents.

CO2 :: apply prompting and reasoning strategies, including chain-of-thought and structured prompting, to develop LLM-driven applications.

CO3 :: apply supervised fine-tuning, parameter-efficient techniques, and reinforcement learning strategies to adapt LLM systems.

CO4 :: build retrieval pipelines for knowledge-grounded reasoning and coordination mechanisms for multi-agent systems.

CO5 :: demonstrate the integration of LLMs into agentic frameworks, utilizing tool-use and generative workflows like LangChain and ReAct.

CO6 :: understand metrics, safety guardrails, ethical considerations, and principles for deploying reliable LLM and agentic systems.

Unit I

Foundations of LLMs and Agentic Systems : Introduction to intelligent agents and their types, Emergence of Agentic AI, Transformer architecture basics, Attention mechanism, Tokenization approaches, Pre-training objectives in LLMs, Agent perception-action cycle, Agent-environment interaction, Overview of popular LLMs

Unit II

Prompt Engineering, Reasoning and Agent Architectures : Zero-shot, few-shot prompting and chain-of-thought prompting, Structured prompting including function calling, Prompt strategies for task planning, Goal-oriented action planning, Agent memory and state management, Basic reasoning strategies

Unit III

Learning, Adaptation and Fine-Tuning for LLMs and Agents : Supervised fine-tuning and instruction tuning, Parameter-efficient techniques (LoRA), Quantization for deployment, Reinforcement learning foundations (Q-learning basics), Reward shaping, and learning from human feedback for agent behavior

Unit IV

Retrieval-Augmented Generation and Multi-Agent Collaboration : Embeddings and semantic search, Chunking and indexing strategies, Vector databases, Retrieval pipelines for knowledge-grounded reasoning, Foundations of multi-agent systems, Multi-agents communication and coordination mechanisms, Multi-agent basic task allocation strategies

Unit V

LLM-Driven Agents and Generative Workflows : Integration of LLMs into agentic frameworks, Chain-of-thought reasoning for planning, Memory augmentation, Tool-use and API orchestration by agents, ReAct framework, Introduction to LangChain Agents and AutoGPT, Autonomous task delegation, and applications in automation and decision support

Unit VI

Evaluation, Safety and Deployment of LLM and Agentic Systems : Metrics for evaluating LLM-driven agents, Safety guardrails and Prompt filtering, Explainability of autonomous decisions, Ethical considerations and Responsible AI guidelines, Testing approaches, Scalable serving frameworks, Monitoring systems, Principles for deploying reliable LLM and agentic systems

List of Practicals / Experiments:

List of Practicals

- Implement tokenization and visualize attention maps of a small transformer model for sample text.
- Design and compare zero-shot, few-shot, and chain-of-thought prompts for reasoning tasks using an open-source LLM.
- Build a simple goal-based agent that selects actions based on objectives in a toy environment.

- Perform LoRA-based fine-tuning of a small LLM on an instruction dataset and evaluate improved responses.
- Implement Q-learning for a grid-world agent and analyze how reward shaping influences policy learning.
- Generate embeddings and build a semantic search index using FAISS for document retrieval.
- Simulate interaction between two collaborative agents performing task allocation in a shared environment.
- Develop a basic RAG pipeline combining a vector database and an LLM to answer knowledge-grounded queries.
- Create an LLM-driven ReAct-style agent that uses at least one external tool to complete a multistep task.
- Evaluate LLM outputs for hallucination and safety risks and apply guardrails to improve output reliability.

Text Books:

1. BUILDING AGENTIC AI SYSTEMS: CREATE INTELLIGENT, AUTONOMOUS AI AGENTS THAT CAN REASON, PLAN, AND ADAPT by ANJANAVA BISWAS, WRICK TALUKDAR, PACKT PUBLISHING

References:

1. NATURAL LANGUAGE PROCESSING WITH TRANSFORMERS: BUILDING LANGUAGE APPLICATIONS WITH HUGGING FACE by LEWIS TUNSTALL, LEANDRO VON WERRA, THOMAS WOLF, O'REILLY