Advanced topic in Computing (LLMs

Code	Credit Hours
CS490	2+1

Course Description

The "Advanced Topics in Computing - Large Language Models" course is designed for students and professionals seeking an in-depth understanding of cutting-edge natural language processing technologies, focusing on large language models. This course explores the theoretical foundations, practical applications, and ethical considerations surrounding the development and utilization of large language models.

Text Book:

1. https://www.cs.princeton.edu/courses/archive/fall22/cos597G/
2. https://stanford-cs324.github.io/winter2022/

Reference Book:

1. https://github.com/mlabonne/llm-course/tree/main

Prerequisites

None

ASSESSMENT SYSTEM FOR THEORY (75 % of the course)

Quizzes	10%
Assignments	10%
Mid Terms	20%
ESE	35%

ASSESSMENT SYSTEM FOR LAB (25 % of the course)

Assessments (Viva and Quizzes)	40%
Lab Work and Report	30%
Lab Project	30%

Teaching Plan

Week No	Topics	Learning Outcomes
1	Introduction	Intro: Revisiting Deep Learning, AI project life cycle, evaluation metrics

2	NLP recap	Text Preprocessing, Feature Extraction Techniques (TF-IDF, n-grams)
3	NLP recap	Word Embeddings (Word2Vec, GloVe, and FastText), RNNs (LSTMs and GRUs) and its limitations
4	Transformers	Architecture, self-attention mechanism, multi-head attention
5	Language model categories	Encoder-only models: Bert, Roberta, Encoder decoder models: T5, BART, Decoder-only models: GPT
6	LLM Modeling (complete life	architecture, tokenization; LLM training: pretraining, supervised finetuning with instruction tuning datasets
7	LLM Modeling (complete life	alignment with human values: Reinforcement learning with human feedback (RLHF)
8	LLM Modeling (complete life cycle)	alignment with human values: recent techniques: DPO, self rewarding LLMs
9	**Mid Semester Break**	
10	Inference, benchmarking, and scaling laws of LLMs	Text generation, sampling techniques; Benchmarking and evaluating LLMs; Scaling laws of LLMs
11	Interacting with LLMs	Prompting, few-shot prompting, and in-context learning
12	Interacting with LLMs	Reasoning of LLMs, Emerging properties of LLMs, LLLMs as knowledge basis
13	LLM adaption for downstream tasks	Parameter efficient fine-tuning techniques (PEFT: LoRA, Prefix tuning; Training pipeline for downstream task adaptation
14	Efficiency, security and safety-alignment of LLMs	Optimization of LLMs during training and inference; Security aspects of LLMs: Prompt hacking, Jailbreak; Safety-alignment of LLMs: Conflict of objectives: Instruction following vs Safety following training

15	New trends, specialized language assistants and LLMs in Computer Vision New trends	MoEs, RAG, Model merging, positional embeddings; Specialized Language assistants; LLMs in Computer Vision
16	Limitations	Limitations and directions to improve LLMs: Limited context size; hallucination; Sleeper Agents; LLMs for mobile devices
17	**Project Presentations**	
18	**End Semester Break**	

Labs :

Lab #	Description
1	Computer vision (Introductory Tasks)
2	Introduction to NLP
3	N-gram Models
4	word2vec part 1
5	word2vec part 2
6	Sentiment Analysis (Open Ended)
7	Encoder, Decoder
8	Transformers
9	BERT
10	Prompting LLMs
11	LoRA
12	Reinforcement Learning
13	Semester Project
14	Semester Project

15	Semester Project
16	Semester Project