Course Name		Applied Machine Learning with LLMs			
Duration		30 hours (6 hours/day)	Break time		1hour/day
Module	Topics	Things to be covered	Must have / good to have	Rough time (hours)	Labs
1	Basic Statistics	Probability, Distributions, Bayes, Central Limit Theorem, p-value	Must	3	no
2	Basics of ML	Regression / OLS	Must		No
		What is Machine Learning	Must		No
		Supervised, Unsupervised Learning	Must		No
		Terminologies: Training, Testing, Inference	Must		No
		DS Projects flow, CRISP-DM	Must		No
		EDA, Pre-processing, feature engineering, NLP techniques	Lab good to have		EDA Implementation: Predicting Ad Clicks Dataset
3	Supervised ML	Classification vs Regression	Must	8	No
		Regression: Linear	Must		Yes
		Classification: Logistic Reg. k-nn, Trees, Bagging & boosting (high level)	Must		Implementation with scikit-learn (ML) and Tensorflow (DL): Predicting Ad Clicks Dataset
		Networks: Neural Networks, CNN, RNN (very high level)	Must		No
4	Unsupervised ML	Clustering, PCA	Good to have		Implementation with scikit-learn (ML): Predicting Ad Clicks Dataset
5	Miscellaneous Topics	Forecasting, Encoders, Vector Embeddings	Must	2	Demo
6	Large Language Models	Introduction to LLMs: Transformer architecture, and how LLMs work.	Must	1	No
		Prompt engineering: Crafting effective prompts to elicit desired outputs from LLMs.	Must	2	Yes
		RAG: How to bring your own data to LLMs. Introduction to Langchain and its components Types of RAG Basic and generic RAG Multi-Modality RAG	Must	5	Embedding & Retrieval Mechanisms: (advertisement industry pdf) RAG Implementation with Langchain: (advertisement industry pdf)
		Agentic Architecture: Building complex flows via LLM to power business use case. Open Source Agent framework- Autogen , CrewAI , Langgraph	Must	2	Creating agents using LangChain: Basic demonstration Autogen: Coding and Financial Agent, Blogpost writing agent
		Limitations and challenges of LLMs: Bias, factual accuracy, and responsible use.	Must	1	No