COURSE CONTENT

LARGE LANGUAGE MODELS								
II Semester: CSE								
Course Code	**Category**	**Hours / Week**			**Credits**	**Maximum Marks**		
		L	**T**	**P**	**C**	**CIA**	**SEE**	**Total**
2525815	**Foundation**	3	0	0	3	40	60	100
Contact Classes: 45	**Tutorial Classes: Nil**	**Practical Classes: Nil**				**Total Classes: 45**		
Prerequisites: Mathematical foundations, programming skills, and machine learning knowledge.								

Course Overview:

This course is an advanced course that explores the theory, architecture, training, and application of modern AI systems capable of understanding and generating human language. The course bridges deep learning and natural language processing, focusing on transformer-based architectures, large-scale training methods, and practical deployment strategies.

Course Objectives:

1. To understand the fundamentals of Natural Language Processing and Large Language Models.
2. To develop practical skills in tokenization, tagging, and language model implementation using Python and NLP frameworks.
3. To implement and fine-tune deep learning and LLM models using tools such as PyTorch and TensorFlow.
4. To explore advanced LLM techniques including prompt engineering, RAG, PEFT, and model evaluation methods.
5. To design and develop real-world AI applications using LLM-based technologies while addressing ethical and safety considerations.

Course Outcomes: After Completion of the Course, Students should be able to

1. Apply transformer architectures and pretraining objectives to build baseline natural language and generation systems.
2. Analyze fine-tuning and domain adaptation methods to customize LLMs for specialized industry tasks.
3. Evaluate prompt engineering strategies to optimize performance in real-world applications as chatbots and code assistants.
4. Distinguish evaluation metrics and deployment strategies for efficient approaches in scaling LLM based enterprise solutions.
5. Create responsible LLM-driven applications by integrating fairness, safety, and retrieval augmented generation techniques for domains such as healthcare, education, and legal services.

UNIT - I: Foundations of Large Language Models

Introduction to LLMs: Definition, scope, and historical evolution from statistical NLP to transformers. The Transformer architecture: Attention mechanisms, self-attention, multi-head attention. Pretraining objectives: Masked language modeling (MLM), Causal language modeling (CLM). Evolution of LLMs: BERT, GPT series, T5, LLaMA, Mistral.

UNIT - II: Training and Fine-Tuning LLMs

Pretraining datasets and tokenization: BPE, Sentence Piece, Word Piece. Fine-tuning approaches: Full fine- tuning, LoRA, adapters, instruction tuning. Domain adaptation and few-shot/zero-shot learning. Data augmentation for LLMs and prompt-based tuning.

UNIT - III: Prompt Engineering and Applications

Principles of prompt design: Zero-shot, few-shot, and chain-of-thought prompting. System prompts, role prompting, and context length optimization.

Use cases: Text generation, summarization, code generation, question answering, chatbots. Tools & frameworks: Lang Chain, Llama Index, Hugging Face Transformers.

UNIT - IV: **Evaluation and Deployment of LLMs**

Evaluation metrics: Perplexity, BLEU, ROUGE, METEOR, human evaluation.

Benchmark datasets: GLUE, SuperGLUE, HELM, BIG-bench.

Deployment strategies: API-based deployment, on-prem deployment, inference optimization. Scaling and latency considerations; quantization and pruning for LLMs.

UNIT - V: Ethics, Safety, and Future Directions

Bias, fairness, and toxicity in LLMs. Hallucination problem and mitigation techniques. Legal and regulatory issues: Copyright, data privacy, AI Act. Trends in LLM research: Multimodal LLMs, retrieval-augmented generation (RAG), open-source LLM ecosystems.

TEXT BOOKS:

1. Vaswani, A. et al. (2017) Attention Is All You Need – NIPS Conference Paper.
2. Lewis, P. et al. (2021) Language Models are Few-Shot Learners – OpenAI Research Paper.
3. Tunstall, L., von Werra, L., & Wolf, T. (2022) Natural Language Processing with Transformers – O'Reilly Media.

REFERENCE BOOKS:

1. Bommasani, R. et al. (2021) On the Opportunities and Risks of Foundation Models – Stanford CRFM.
2. Jurafsky, D., & Martin, J. H. (2023) Speech and Language Processing (3rd Edition draft) – Pearson.
3. Mollick, E., & Mollick, L. (2024) Co-Intelligence: Living and Working with AI – Little, Brown Spark.

ELECTRONIC RESOURCES:

1. **https://www.coursera.org/learn/natural-language-processing**
2. **https://www.coursera.org/learn/transformers-nlp**
3. **https://www.deeplearning.ai/learning-paths/deep-learning-specialization**
4. **https://huggingface.co/learn/nlp-course**

MATERIALS ONLINE:

1. Course template
2. Tutorial question bank
3. Tech talk and Concept Video topics
4. Open-ended experiments
5. Definitions and terminology
6. Assignments
7. Model question paper – I
8. Model question paper – II
9. Lecture notes
10. E-Learning Readiness Videos (ELRV)