Libera Università Internazionale
degli Studi Sociali Guido Carli

LUISS

Syllabus

Academic Year	2027/2028
Program	Data Science and Management
Course	Foundation Models
Term	I semester
Year	2
SSD	ING-INF/05
Credits	6

INSTRUCTIONAL GOALS	The course introduces students to the foundation models that power the latest AI applications, particularly large language models (LLMs). Upon successful completion of this course, students will: 1. Master the Transformer architecture and the mathematical foundations of self-attention mechanisms. 2. Analyze pre-training methodologies and the scaling laws that govern model performance and compute efficiency. 3. Implement adaptation techniques, including Parameter-Efficient Fine-Tuning (PEFT) and Retrieval-Augmented Generation (RAG), for specialized domains. 4. Critically evaluate model alignment, safety, and ethical implications to deploy robust AI systems in real-world contexts.
INTENDED LEARNING OUTCOMES They describe what a learner is expected to know, understand and be able to demonstrate after completion of a learning path.	**Knowledge and Understanding.** Students will demonstrate a deep theoretical understanding of: • The mathematical foundations of the Transformer architecture, specifically multi-head self-attention and positional encoding. • The taxonomy of Foundation Models (Encoder, Decoder, and Hybrid) and the empirical scaling laws that govern their performance. • The principles of self-supervised pre-training and the alignment mechanisms (RLHF, DPO) used to guide model behavior. **Applying Knowledge and Understanding.** By the end of the course, students will be able to: • Implement and optimize Parameter-Efficient Fine-Tuning (PEFT) techniques, such as LoRA, for domain-specific tasks. • Architect and deploy Retrieval-Augmented Generation (RAG) systems to integrate external knowledge bases with LLMs. • Apply data-curation and tokenization strategies to prepare large-scale datasets for model training or adaptation. **Making Judgements.** Students will develop the ability to:

	• Critically evaluate the trade-offs between model size, computational cost (FLOPs), and inference latency. • Assess the ethical risks of foundation models, including bias, hallucinations, and privacy leaks, and propose mitigation strategies. • Select the most appropriate model architecture and adaptation strategy based on specific application constraints and performance benchmarks. **Communication Skills.** Students will be able to: • Articulate complex architectural concepts and experimental results through technical reports and structured presentations. • Communicate the societal and business implications of AI deployment to non-technical stakeholders, balancing technical limitations with potential value. • Defend technical design choices during peer reviews or collaborative project defense. **Learning Skills.** Students will acquire the skills to: • Autonomously navigate and synthesize the rapidly evolving research literature (e.g., ArXiv, top-tier AI conferences) to keep pace with the state-of-the-art. • Independently master new deep learning frameworks and open-source tools as they emerge in the AI ecosystem. • Develop a mindset for continuous experimentation, allowing them to troubleshoot and refine models in a fast-changing technological landscape.
Pre-requisites	Knowledge of machine learning ad deep learning frameworks.
Course content	• Evolution of Foundation Models • The Transformer Architecture • Tokenization and Data Curation • Pre-training Objectives • Scaling Laws and Efficiency • Parameter-Efficient Fine-Tuning (PEFT) • Prompt Engineering and In-Context Learning • Retrieval-Augmented Generation (RAG) • Human Alignment (RLHF & DPO) • Evaluation, Ethics, and Safety
Reference Books	Textbook: ***Foundations of Large Language Models*** by Tong Xiao & Jingbo Zhu (2024). Slides set of the course and other teaching materials (available on the class web site). Suggested additional readings: • ***Build a Large Language Model (From Scratch)*** by Sebastian Raschka (2025). • ***The LLM Engineering Handbook*** by Paul Iusztin & Maxime Labonne (2025/2026). • ***The Hundred-Page Language Models Book*** by Andriy Burkov (2024).

Teaching Methods	• Classroom-based lectures • Case studies • Projects
Assessment	• Mid-term written exam (1/3) • Project (2/3).