

ISOM4110 – Large Language Models and Agentic AI (2026-27 Fall Term)

General Information

- 3 Credits (Letter Graded), IS major elective / IS AI option required
- Pre-requisite: ISOM3320 OR ISOM3400
- Lectures/Labs:

Mon	12:00pm – 01:20pm	LSK-1005
Wed	12:00pm – 01:20pm	LSK-1005
- Instructor: Prof. Jean WANG <jeanwang@ust.hk> (office hour by appointment)
- TA: Mr. Tommy NG <imtng@ust.hk> (office hour by appointment)

Course Description

Generative AI (GenAI) and Large Language Models (LLMs) are transforming how organizations create content, manage knowledge base, automate workflow, and make decisions. At the same time, these emerging technologies introduce new risks and challenges related to reliability, bias, security, ethics, and alignment with organizational goals.

This course trains students to design, implement, and deploy LLM-based agentic solutions in business settings. Students will learn core concepts and common practices behind modern LLMs, including training, finetuning, prompting, retrieval augmentation and safety guardrails. The course will also introduce how LLM-based agentic AI can be deployed and integrated into existing data analytics, decision support, and business processes. Through hands-on exercises in Python, students will gain practical experience using LLM APIs and prompt templates, finetuning lightweight LLMs, orchestrating LLMs agents, and experimenting with simplified RAG (Retrieval Augmented Generation) pipelines. The course will also introduce key issues in evaluation, governance, and risk management, enabling students to critically assess both the potential and limitations of LLM agentic systems.

Course Intended Learning Outcomes

- Understand the fundamental concepts, capabilities and evolution of Large Language Models and Agentic AI systems
- Explain the core architectures and training techniques underlying LLMs, and analyze how design choices affect model behavior and performance
- Evaluate the performance, reliability, and limitations of LLM-based agentic systems, including issues of accuracy, bias, fairness, safety, and token cost
- Apply effective prompt design, pre-training, parameter efficient fine tuning, workflow design, context engineering and tool integration techniques to develop practical LLM driven agentic solutions

Required Background and Prerequisite Knowledge

This course assumes foundational proficiency in Python, as it is essential for the hands-on labs and exercises. Additionally, a basic understanding of Machine Learning concepts, particularly in training and finetuning Neural Networks, is recommended to fully grasp how Large Language Models are developed, deployed, and applied in real-world scenarios.

Software Needed

A Python coding environment is needed, preferably [Cursor](#) or Google [Colab](#).

Teaching Schedule (tentative)

WK	Lecture Topic	Hands-on Lab *
1	[Sep 2] Lec01 – Overview of AI and LLMs	
2	[Sep 7] Lec01 (cont'd)	[Sep 9] Lab01 – Working with LLM API in Python
3	[Sep 14] Lec02 – LLM Prompt Engineering Lec03 – Tokenizer and Word Embeddings	[Sep 16] Lab02 – Prompt Templates and Engineering
4	[Sep 21] Lec04 – Model Architecture: RNN, Seq2Seq & Transformer	[Sep 23] Lab03 – Using and Training Tokenizers
5	[Sep 28, 30] Lec04 – Model Architecture: LLM	
6	[Oct 5] Lec05 – LLM Training	[Oct 7] Lab04 – Training a Lightweight LLM
7	[Oct 12, 14] Lec06 – LLM Deployment: Efficient Fine-Tuning and Quantization	
8	[Oct 19] <i>No class on public holiday</i>	[Oct 21] Lab05 – Finetuning a Lightweight LLM
9	[Oct 26] Lec07 – LLM Deployment: RAG	[Oct 28] Lab06 – RAG with LangChain
10	[Nov 2] Lec08 – LLM Deployment: Agent	[Nov 4] Lab07 – LangChain Agent with Tools
11	[Nov 9] Lec09 – Multimodal Models	[Nov 11] Lab08 – Receipt-to-Email Agent with n8n
12	[Nov 16] Lec10 – LLM Evaluation, Guardrails and Governance	[Nov 18] Lab09 – New Scrapping Agent with n8n
13	[Nov 23, 25, 30] Project Presentation & Discussion	

Course Assessments

Components	Weighting
Class Participation	15%
Lab Submissions	25%
Group Project	20%
Final Examination	40%
Total:	100%

- **Class Participation (15%):** week 1 to week 11
Students are required to attend all lectures and actively participate in class. This component of the assessment is based on in-class engagement and the quality of contributions to open discussion. Some pop-up quizzes will be conducted in class to evaluate students' understanding of the knowledge and concepts discussed in lectures.
- **Lab Submissions (25%):** week 2 to week 11
These are individual assessments. In the hands-on labs, students are required to use Python and/or AI agent framework/platform to accomplish a specific LLM deployment or prototyping task. Upon completion, students must submit their implementation code and any other related files.
- **Group Project (20%):** week 10 to week 13
This is a group assessment. 2-3 Students will form a group to conduct research and present their findings on a specific LLM-related topic. Groups may also propose a topic of their own interest, subject to the instructor's approval.
- **Final Exam (40%):** end of semester
The logistic arrangement of the final exam will be announced in Canvas later.

Recommended Readings

- 2026 AI Index Report – Stanford University
<https://aiindex.stanford.edu/report/>
- State of AI trust in 2026: Shifting to the Agentic Era - McKinsey
<https://www.mckinsey.com/capabilities/tech-and-ai/our-insights/tech-forward/state-of-ai-trust-in-2026-shifting-to-the-agentic-era>
- The State of AI in 2025: How Organizations are Rewiring to Capture Value – McKinsey
<https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai>
- The Wharton Blueprint for AI Agent Adoption – Wharton
<https://knowledge.wharton.upenn.edu/special-report/wharton-blueprint-ai-agent-adoption/>

Course Content Appropriate Access Policy

Appropriate access to course materials is given for personal academic study and review purposes only. Unless otherwise stated in writing, the content may not be shared, distributed, modified, transmitted, reused, sold, or otherwise disseminated. Unethical sharing of course materials on commercial websites, such as *CourseHero* and *Studocu*, is a breach of academic conduct. Violators of this policy will be referred to the University for disciplinary purposes.