ISOM4000D – Generative AI and Large Language Models (2025-26 Spring Term)

General Information

- 3 Credits (Letter Graded)
- Teaching Mode: Face-to-Face
- Lectures/Labs: Wed 03:00pm - 04:20pm LSK-1032
  Fri 03:00pm - 04:20pm LSK-1032
- Instructor: Prof. Jean WANG <jeanwang@ust.hk> (office hour by appointment)
- TA: Mr. Samuel LAI <imsamuel@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 GenAI and LLM-based solutions in business settings. Students will learn core concepts behind modern LLMs, prompt and workflow design, and how GenAI can be integrated into existing data analytics, decision support, and business processes. Through hands-on exercises in Python, students will gain practical experience using GenAI tools and APIs, orchestrating LLMs with other data and software tools, and experimenting with simplified ("toy") model training pipelines. The course will also introduce key issues in evaluation, governance, and risk management, enabling students to critically assess both the potential and the limitations of LLM agentic systems.

Course Intended Learning Outcomes

- Understand the fundamental concepts, evolution, and capabilities of Generative AI and Large Language Models
- Identify and analyze appropriate use cases for GenAI and LLMs in various business and organizational contexts, recognizing both the opportunities and the associated risks
- Evaluate the performance, reliability, and limitations of GenAI and LLM-based applications including issues of accuracy, bias, fairness, safety, and robustness
- Apply effective prompt design, workflow design, and tool integration techniques to develop practical GenAI/LLM-driven solutions that address real-world business needs

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.

Teaching Schedule (tentative)

WK	Lecture Topic	Hands-on Lab *
1	[Feb 4 & Feb 6] Lec01 – Overview of AI and LLMs	*No lab in this week*
2	[Feb 11] Lec02 – LLM Prompt Engineering	[Feb 13] Lab01 – Working with LLM API in Python
3	[Feb 18] No class on public holiday	[Feb 20] Lab02 – Prompt Templates and Engineering
4	[Feb 25] Lec03 – RNN and Word Embeddings	[Feb 27] Lab03 – Using and Training Tokenizers
5	[Mar 4 & Mar 6] Lec04 – LLM Architecture	*No lab in this week*
6	[Mar 11] Lec05 – LLM Training and Finetuning	[Mar 11] Lab04 – Finetuning a Lightweight LLM
7	[Mar 18] Lec06 – Deploying LLMs in Business Applications	[Mar 20] Lab05 – LangChain Agent Fundamentals
8	[Mar 25] Lec07 – LLM Agent Basics	[Mar 27] Lab06 – RAG Agent with Langchain
9	[Apr 1] Lec08 – LLM Agent Components and Tool Use	[Apr 10] Lab07 – Agent with Memory and Tools
10	[Apr 15] Lec09 – Multimodal Models	[Apr 17] Lab08 – No-code Agent Building with n8n
11	[Apr 22 & Apr 24] Lec10 – Open Issues of LLMs	*No lab in this week*
12	[Apr 29] Project Presentation & Discussion	[May 1] No class on public holiday
13	[May 6 & 8] Project Presentation & Discussion	

*** Software Requirement:** A Python coding environment is needed, preferably Anaconda + Jupyter Notebook or Python + Cursor or Google Colab.

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 participant in class. This component of the assessment is based on in-class engagement and the quality of contributions to open discussion. Some pup-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

- 2025 AI Index Report – Stanford University
  https://aiindex.stanford.edu/report/

- State of AI Report 2025 – stateof.ai
  https://www.stateof.ai/

- 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

- Accountable Acceleration: Gen AI Fast-Tracks into the Enterprise – Wharton
  https://knowledge.wharton.upenn.edu/special-report/2025-ai-adoption-report/

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* is a breach of academic conduct. Violators of this policy will be referred to the University for disciplinary