The Hong Kong Polytechnic University

Subject Description Form

Subject Code	COMP5584
Subject Title	Large Language Models for Agents
Credit Value	3
Level	5
Pre-requisite/ Co-requisite/ Exclusion	Nil (but knowledge in Machine Learning, Artificial Intelligence, Natural Language Processing, and Data Analytics is preferable).
Objectives	This subject aims to equip students with an advanced understanding of large language models (LLMs) for their deployment as intelligent agents. It contains advanced topics such as reasoning, planning, multi-agent collaboration, and the use of tools for task execution, alongside the technical infrastructure required for agent-based systems. Students will critically explore applications in robotics, code generation, web automation, and scientific discovery, while addressing ethical, privacy, and safety challenges inherent to LLM agents. With rigorous analysis, hands-on projects, and research-driven inquiry, students will acquire the expertise to evaluate, design, and innovate in the field of LLM agents, preparing them for cutting-edge roles in academia or industry.
Intended Learning Outcomes	Upon completion of the subject, students will be able to: a. Demonstrate an advanced understanding of the foundational principles behind large language models (LLMs) in reasoning, planning, and tool use, and their role in building intelligent agents. b. Design and implement LLM-driven agents to solve complex tasks in domains such as robotics, web automation, scientific discovery, and code generation. c. Analyse and understand the limitations, risks, and ethical considerations of LLM agents, including issues related to privacy, safety, and human-agent interaction. d. Develop and assess novel LLM agent infrastructures, incorporating techniques like retrieval-augmented generation, multimodal integration, and multi-agent collaboration.
Subject Synopsis/ Indicative Syllabus	1. Foundations of LLMs for Intelligent Agents - Reasoning, planning, and tool use in agent systems. - Understanding LLM agent infrastructures. 2. Advanced Techniques and Architectures of LLM Agents - Retrieval-augmented generation (RAG). - Tool and API integration for task execution. - Multimodal agents. - Applications in robotics and beyond. 3. Applications of LLM Agents

	- Code generation and software engineering. - Data science and knowledge discovery. - Case studies: real-world deployment in diverse industries. 4. Evaluation and Ethics - Metrics and benchmarking for LLM-driven agent systems. - Addressing privacy, safety, and ethical concerns. 5. Human-Agent Interaction and Collaboration - Personalization in agent design and deployment. - Human-agent interaction. - Multi-agent collaboration.
Teaching/Learning Methodology	1. Lectures and Seminars Core theoretical concepts, foundational principles, and advanced topics such as reasoning, planning, and LLM agent infrastructures will be introduced through lectures. Seminars will focus on deeper discussions of cutting-edge advancements, including retrieval-augmented generation, multimodal agents, and ethical considerations. To ensure active learning, in-class activities such as guided discussions, group exercises, and case study reviews will be incorporated. These activities will encourage critical thinking and allow students to explore complex ideas collaboratively. Topics like privacy, safety, and alignment in LLM agents will also be debated in seminar sessions to develop critical perspectives. 2. Labs and Tutorials Hands-on experience will be provided via labs and tutorials to reinforce theoretical concepts and develop practical skills. Students will work on guided mini-projects and experiments, e.g., - Building simple LLM agents for specific tasks. - Exploring retrieval-augmented generation techniques. - Integrating APIs and tools to enhance agent functionality. Labs will also include practical exercises in deploying multimodal agents, benchmarking agent performance, and addressing challenges like hallucinations. Tutorials will complement lab sessions by providing step-by-step guidance on technical implementations and troubleshooting. 39 hours of class activities including - lectures, tutorials, lab, and seminars where applicable.

Assessment Methods in Alignment with Intended Learning Outcomes

Students' performance in this subject will be assessed using a letter-grading system in accordance with the University's convention from grade F (failure) to A+. The relative weighting of the different assessment components are as follows:

Specific assessment methods/tasks	% weighting	Intended subject learning outcomes to be assessed (Please tick as appropriate)			
		a	b	c	d
1. Projects	20%		✓		✓
2. Quizzes	10%	✓		✓	
3. Exam	70%	✓	✓	✓	✓
Total	100 %				

Explanation of the appropriateness of the assessment methods in assessing the intended learning outcomes:

Hands-on projects will enable students to design and implement LLM agents to solve complex, real-world tasks (ILO b). These projects also encourage students to develop and assess novel agent infrastructures, incorporating advanced techniques like retrieval-augmented generation and multimodal integration (ILO d).

Quizzes are designed to deepen students' understanding of the foundational principles of LLMs and intelligent agents introduced in lectures (ILO a). Additionally, they include questions that prompt students to critically engage with ethical considerations and analyse the risks and limitations of LLM agents (ILO c).

The comprehensive final exam evaluates students' mastery of theoretical concepts, practical applications, and critical analysis. It assesses their advanced understanding of LLM foundations (ILO a), ability to design solutions using LLM-driven agents (ILO b), and capacity to analyse ethical implications, limitations, and risks (ILO c). The exam also challenges students to propose and evaluate novel agent infrastructures (ILO d).

Student Study Effort Expected		
	Class contact:	
	▪ Class Activities (lectures, seminars, labs, tutorials, and in-class quizzes).	39 Hrs.
	Other student study effort:	
	▪ Self-study, Projects, and Exams.	45 Hrs.
	▪ After-class Reading.	20 Hrs.
	Total student study effort	104 Hrs.

Reading List and References	1. Alto, V. (2024). Building LLM Powered Applications: Create intelligent apps and agents with large language models. Packt Publishing Ltd. 2. Alammar, J., & Grootendorst, M. (2024). Hands-On Large Language Models: Language Understanding and Generation. O'Reilly Media, Inc. 3. Dan, J., & Martin, J. H. (2009). Speech and Language Processing: An Introduction to Natural Language Processing, Computational Linguistics, and Speech Recognition. Prentice Hall Series in Artificial Intelligence. 4. Bouchard, L. F., & Peters, L. (2024). Building LLMs for Production: Enhancing LLM Abilities and Reliability with Prompting, Fine-Tuning, and RAG. Towards AI, Inc. 5. Gupta, M. (2024). LangChain in Your Pocket: Beginner's Guide to Building Generative AI Applications Using LLMs. Mehul Gupta. 6. Papers and articles selected from: - Conference on Neural Information Processing Systems (NeurIPS) - International Conference on Machine Learning (ICML) - International Conference on Learning Representations (ICLR) - Annual Meeting of the Association for Computational Linguistics (ACL) - Conference on Empirical Methods in Natural Language Processing (EMNLP) - Other top-tier conferences and journals in NLP, ML, and AI.