The Hong Kong Polytechnic University

Subject Description Form

Please read the notes at the end of the table carefully before completing the form.

Subject Code	COMP6713
Subject Title	Advanced Large Language Models and Beyond
Credit Value	3
Level	6
Pre-requisite/ Co-requisite/ Exclusion	Nil (but knowledge in Artificial Intelligent/ Machine Learning is preferable)
Objectives	• Provide students with foundational aspects of large language models (LLM), including modeling, theory, ethics, and systems. • Enhance students' expertise in multimodal LLMs (MLLM) and text-to-video generation technologies such as stable diffusion. • Provide students with hands-on experience using these models in applications like chatbots, search, advertising, and emerging robotics technologies.
Intended Learning Outcomes *(Note 1)*	Upon completion of the subject, students will be able to: a. Gain a thorough understanding of LLMs' fundamentals, theory, ethics, and systems, supplemented by hands-on experience. b. Demonstrate the ability of critical reading and analysis through paper reviews and panel discussions, enhancing abilities to write reviews and analyze papers from multiple perspectives. c. Gain practical experience with language models such as GPT-3 and Mistral, evaluating their capabilities and risks, tailored to future projects and interests. d. Build LLMs using a designated compute budget to train and refine, enhancing to systematically evaluate and improve models based on chosen characteristics.
Subject Synopsis/ Indicative Syllabus *(Note 2)*	1. **Fundamental:** Feed forward networks, backpropagation, sequence-to-sequence models, adaptation framework, perplexity, data processing and ecosystems, tokenization, embeddings, layer normalization, positional encoding, transformers, attention mechanisms, transfer learning, multimodal inputs, fusion 2. **Advanced:** Data parallelism, model parallelism, pipeline parallelism, pretraining, scaling laws, mixture of experts, memory-augmented (retrieval) models (RAG), probing, prompting, zero-shot/few-shot/many shots learning, fine-tuning and lightweight fine-tuning, cross-modal learning,

3. **Applications:** natural language understanding (NLU) and generation (NLG), sentiment analysis, information extraction, question answering systems, personalised recommendations, language translation, image and video captions, visual question answering (VQA), cross-modal information retrieval, sentiment analysis with visual context, robotics.

Teaching/Learning Methodology

(Note 3)

Lecture and tutorial sessions equip students with a solid foundation in natural language processing, using detailed examples and interactive Q&A segments to enhance understanding.

Lab sessions enable students to master practical techniques and the necessary tools to replicate and refine advanced models for a range of natural language applications.

During paper panel review and discussion sessions, students develop their critical reading and analytical skills. These sessions are designed to deepen understanding of what makes an academic review both thorough and insightful, and to improve the ability to write comprehensive reviews across a variety of research papers. Panel discussions complement these objectives by encouraging analysis from diverse perspectives, enriching the learning experience.

Assessment Methods in Alignment with Intended Learning Outcomes

(Note 4)

Specific assessment methods/tasks	% weighting	Intended subject learning outcomes to be assessed (Please tick as appropriate)			
		a	b	c	d
1. Assignment and Quiz	30%	✓		✓	✓
2. Project and Presentation	40%			✓	✓
3. Paper Panel Review and Discussion	30%		✓		
Total	100 %				

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

Quizzes assess students' comprehension of the fundamental concepts in natural language processing. Assignments and presentations gauge students' understanding of the latest developments in the field, their ability to learn independently, and their critical thinking skills, as well as their proficiency in both written and oral communication. Projects evaluate students' critical and creative thinking capabilities, problem-solving skills, and their ability to apply the techniques and tools they have learned to real-world applications. The purpose of the panel reviews is to critically read and analyze research papers, while the panel discussions aim to explore these papers from multiple perspectives, enriching the analytical process.

Student Study Effort Expected

Class contact:	
▪ Lecture/Tutorial/Lab/Paper panel	39 Hrs.
Other student study effort:	

	▪ Completing Assignment and Projects, Reviewing and preparing for quiz, reading papers, and presentation	85 Hrs.
	Total student study effort	124 Hrs.
Reading List and References	(1) T. Brown, et al., Language Models are Few-Shot Learners, arxiv: 2005.14165 (2) L. Ouyang, et al., Training language models to follow instructions with human feedback, arxiv:2203.02155 (3) R. Bommasani, et al., On the Opportunities and Risks of Foundation Models, arxiv:2108.07258 (4) S. Xie, et al., DoReMi: Optimizing Data Mixtures Speeds Up Language Model Pretraining, arxiv:2305.10429 (5) S. Gehman, et al., RealToxicityPrompts: Evaluating Neural Toxic Degeneration in Language Models, arxiv:2009.11462 (6) L. Gao, et al., The Pile: An 800GB Dataset of Diverse Text for Language Modeling, arxiv:2101.00027 (7) N. Carlini, et al., Extracting Training Data from Large Language Models, USENIX 2021 (8) Z. Dai, et al., Transformer-XL: Attentive Language Models beyond a Fixed-Length Context, ACL 2019 (9) K. Clark, et al., ELECTRA: Pre-training Text Encoders as Discriminators Rather Than Generators, ICLR 2020 (10) DeepSpeed Team, DeepSpeed: Extreme-scale model training for everyone, microsoft (11) J. Kaplan, et al., Scaling Laws for Neural Language Models, arxiv:2001.08361 (12) P. Lewis, et al., Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks, NeurIPS 2020 (13) B. Lester, et al., The Power of Scale for Parameter-Efficient Prompt Tuning, EMNLP 2021 (14) Y. Wang, et al., Exploring the reasoning abilities of multimodal large language models (mllms): A comprehensive survey on emerging trends in multimodal reasoning, arxiv:2401.06805 (15) S. Chan, Tutorial on Diffusion Models for Imaging and Vision, arxiv:2403.18103	

Note 1: Intended Learning Outcomes
Intended learning outcomes should state what students should be able to do or attain upon subject completion. Subject outcomes are expected to contribute to the attainment of the overall programme outcomes.

Note 2: Subject Synopsis/Indicative Syllabus

The syllabus should adequately address the intended learning outcomes. At the same time, overcrowding of the syllabus should be avoided.

Note 3: Teaching/Learning Methodology
This section should include a brief description of the teaching and learning methods to be employed to facilitate learning, and a justification of how the methods are aligned with the intended learning outcomes of the subject.

Note 4: Assessment Method
This section should include the assessment method(s) to be used and its relative weighting, and indicate which of the subject intended learning outcomes that each method is intended to assess. It should also provide a brief explanation of the appropriateness of the assessment methods in assessing the intended learning outcomes.