

The Hong Kong Polytechnic University

Subject Description Form

Subject Code	CMS6001
Subject Title	Advanced AI Programming
Credit Value	3
Level	6
Pre-requisite/ Co-requisite/ Exclusion	<ul style="list-style-type: none"> - Background in algorithms, data structures, linear algebra, and probability and statistics. - Prior exposure to artificial intelligence or machine learning. - Proficiency in Python programming and development tools.
Objectives	<p>This subject is designed for doctoral students seeking to transition from being consumers of AI solutions to leaders capable of driving enterprise-scale AI transformation. Moving beyond foundational concepts, the subject provides a rigorous, hands-on exploration of advanced AI modelling and programming, algorithm optimization, AI architectural design, and secure deployment of responsible AI systems. Students will engage with the engineering challenges of large-scale AI, including model selection, performance scaling, resilience, and compliance with governance and ethical standards. The curriculum emphasizes practical AI system-building and integration, equipping participants to translate technical mastery into strategic business impact. The subject culminates in a semester project where learners design and implement a sophisticated AI-driven solution aligned with industry context, preparing them to lead transformative AI initiatives in enterprise and regulatory environments. The objectives of the subjects are as follows.</p> <ol style="list-style-type: none"> 1. To cultivate expertise in designing, programming, and deploying robust AI systems using advanced models and frameworks, optimization techniques, secure architectures and pipelines, and following industry standard methodologies. 2. To enable students to critically evaluate AI technologies and models, and develop the skills to formulate, implement, and justify innovative AI solutions for enterprise transformation. 3. To instil a deep understanding and expertise in lifecycle management, resilience, compliance, and governance, ensuring graduates can lead the development of scalable, trustworthy, and industry-ready AI systems.

<p>Intended Learning Outcomes</p>	<p>Upon completion of the subject, students will be able to:</p> <ol style="list-style-type: none"> a. Master advanced AI system engineering: Design, program, and rigorously evaluate complex AI architectures, optimization pipelines, and scalable deployments to address high-impact industry challenges. b. Lead enterprise AI transformation: Formulate and implement strategic AI initiatives by translating business problems into system-level solutions, and effectively communicate proposals to both technical teams and executive stakeholders. c. Ensure resilience, compliance, and governance: Critically assess and manage the lifecycle of AI deployments, applying advanced techniques to guarantee reliability, scalability, and alignment with enterprise standards and international governance frameworks.
<p>Subject Synopsis/ Indicative Syllabus</p>	<p>The syllabus is structured into four core modules, each building upon the last. The approach combines theoretical understanding with intensive, practical implementation.</p> <p><u>Module 1: Foundations of Advanced AI Engineering</u> This module establishes the technical and organizational foundations for building robust, scalable AI systems.</p> <ul style="list-style-type: none"> • Advanced Framework Programming. Deep dive into PyTorch/ TensorFlow APIs beyond the basics. Custom layers, hooks, computation graphs, and modular, reusable code design. • High-Performance and Distributed Training. GPU architecture fundamentals (CUDA basics), parallelism strategies, and distributed training frameworks (e.g., PyTorch DDP, Horovod). Techniques for handling massive datasets and scaling across clusters. • MLOps and Reproducible Research. Machine Learning Operations (MLOps) principles. Experiment tracking (Weights & Biases), data/model versioning (DVC), and containerization (Docker/Kubernetes) for reproducible, production-ready environments. <p><u>Module 2: Advanced Systems Optimization and Deployment</u> This module focuses on engineering challenges in deploying AI at enterprise scale.</p> <ul style="list-style-type: none"> • Optimization Techniques: Model compression, quantization, pruning, and efficient inference pipelines. Balancing accuracy, latency, and cost. • Secure and Resilient Deployment: Adversarial robustness, anomaly detection, and defenses against model inversion or poisoning attacks. • Lifecycle Management: Continuous integration/continuous deployment (CI/CD) for AI, monitoring drift, retraining pipelines, and compliance with industry regulations. <p><u>Module 3: Reliability, Compliance, and Responsible AI</u></p>

	<p>This module equips learners with the tools to ensure AI systems are reliable, auditable, and aligned with governance and ethical standards:</p> <ul style="list-style-type: none"> • System Reliability and Resilience: Designing fault-tolerant AI systems, redundancy strategies, and performance monitoring at scale. • Compliance and Governance: International standards (e.g., ISO/IEC AI guidelines), regulatory frameworks, and enterprise risk management. • Responsible AI and Ethical Use: Principles of fairness, accountability, and transparency. Identifying and mitigating bias in data and algorithms. Practical implementation of explainability techniques (e.g., SHAP, LIME) to support governance. • Strategic Integration: Aligning AI deployments with organizational strategy, ensuring systems meet both technical and societal expectations. <p><u>Module 4: Industry Transformation Project</u> This final module transitions participants toward an applied project that connects technical mastery with business transformation.</p> <ul style="list-style-type: none"> • Strategic Problem Formulation: Identifying high-impact industry challenges and framing them as system-level AI problems. • Solution Design and Implementation: Developing and deploying a sophisticated AI-driven solution tailored to enterprise needs. • Executive Communication and Peer Review: Structuring proposals, presenting to technical and non-technical stakeholders, and defending decisions in peer review. 																							
<p>Teaching/Learning Methodology</p>	<p>Lectures and Tutorials provide foundational knowledge teaching. The instructor presents the core theories, mathematical underpinnings, and architectural concepts of advanced topics.</p> <p>Labs provide hands-on guidance to enable students to translate theory into practice by implementing complex AI systems.</p> <p>Group Discussions will focus on developing critical evaluation and leadership skills. Discussions emphasize strategic alignment, governance, and responsible AI, preparing participants to lead enterprise adoption.</p>																							
<p>Assessment Methods in Alignment with Intended Learning Outcomes</p>	<table border="1"> <thead> <tr> <th rowspan="2">Specific assessment methods/tasks</th> <th rowspan="2">% weighting</th> <th colspan="3">Intended subject learning outcomes to be assessed (Please tick as appropriate)</th> </tr> <tr> <th>a</th> <th>b</th> <th>c</th> </tr> </thead> <tbody> <tr> <td>1. Assignment and Quizzes</td> <td>50%</td> <td>√</td> <td></td> <td>√</td> </tr> <tr> <td>2. Project and Presentation</td> <td>50%</td> <td>√</td> <td>√</td> <td>√</td> </tr> <tr> <td>Total</td> <td>100 %</td> <td></td> <td></td> <td></td> </tr> </tbody> </table>	Specific assessment methods/tasks	% weighting	Intended subject learning outcomes to be assessed (Please tick as appropriate)			a	b	c	1. Assignment and Quizzes	50%	√		√	2. Project and Presentation	50%	√	√	√	Total	100 %			
Specific assessment methods/tasks	% weighting			Intended subject learning outcomes to be assessed (Please tick as appropriate)																				
		a	b	c																				
1. Assignment and Quizzes	50%	√		√																				
2. Project and Presentation	50%	√	√	√																				
Total	100 %																							

	<p>Explanation of the appropriateness of the assessment methods in assessing the intended learning outcomes:</p> <p>The assessment is balanced to test both technical mastery and strategic application.</p> <ul style="list-style-type: none"> • Assignments and Quizzes (50%) cover ILO (a) and (c) to check hands-on skills in programming, optimization, and responsible deployment, ensuring students can design and evaluate advanced AI systems. • Project and Presentation (50%) integrates all three ILOs by requiring students to solve an industry challenge, build a system-level solution, and present it effectively to technical and business audiences. <p>This structure ensures graduates can both execute complex AI work and lead enterprise transformation.</p>	
<p>Student Study Effort Expected</p>	<p>Class contact:</p>	
	<ul style="list-style-type: none"> ▪ Lecture/Tutorials/Labs 	<p>30 Hrs.</p>
	<p>Other student study effort:</p>	
	<ul style="list-style-type: none"> ▪ Self-Study and After-class Reading, Assignments and Projects, Preparations for Quizzes and Presentations. 	<p>90 Hrs.</p>
	<p>Total student study effort</p>	<p>120 Hrs.</p>
<p>Reading List and References</p>	<ol style="list-style-type: none"> 1. Goodfellow, I., Bengio, Y., & Courville, A. (2016). Deep Learning. MIT Press. 2. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L., & Polosukhin, I. (2017). Attention is all you need. Advances in Neural Information Processing Systems, 30. 3. Russell, S. J., & Norvig, P. (2021). Artificial Intelligence: A Modern Approach (4th ed.). Pearson. 4. Rombach, R., Blattmann, A., Lorenz, D., Esser, P., & Ommer, B. (2022). High-resolution image synthesis with latent diffusion models. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (pp. 10684–10695). IEEE. 5. O'Neil, C. (2016). Weapons of Math Destruction: How Big Data Increases Inequality and Threatens Democracy. Crown. 6. Bishop, C. M. (2024). Deep Learning: Foundations and Concepts. Springer. 7. Lundberg, S. M., & Lee, S.-I. (2017). A unified approach to interpreting model predictions. Advances in Neural Information Processing Systems, 30. 8. Barocas, S., Hardt, M., & Narayanan, A. (2023). Fairness and Machine Learning: Limitations and Opportunities. MIT Press. 	