Module Code	EEEP55C34
Module Name	Advanced AI
ECTS Weighting[1]	5 ECTS
Semester taught	Semester 2
Module Coordinator/s	Dr. Hossein Javidnia
Module Learning Outcomes with reference to the Graduate Attributes and how they are developed in discipline	On successful completion of this module, students should be able to: LO1. Implement and evaluate advanced generative models (GANs, flows, energy-based). LO2. Design and train self-supervised pipelines (VAEs, SimCLR, custom pretext tasks). LO3. Formulate and solve MDPs using MC, TD, Q-learning, SARSA, DQN, REINFORCE, and Actor-Critic. LO4. Build, fine-tune, and distill transformer architectures (tokenizers, LoRA, ViT). LO5. Develop graph neural networks (GCN, GAT, message-passing) for node, edge, and graph tasks. LO6. Apply meta-learning (MAML, Reptile, prototypical, siamese) for few-shot adaptation. LO7. Use LIME, SHAP, and attribution methods to interpret and explain model decisions. LO8. Measure and mitigate bias via fairness metrics (demographic parity, equal opportunity). LO9. Integrate differential privacy, secure MPC, and federated averaging into ML workflows. LO10. Generate adversarial examples and apply robustness defenses with uncertainty estimation. **Graduate Attributes: levels of attainment** To act responsibly - Enhanced To think independently - Attained To develop continuously - Attained To communicate effectively - Enhanced

[1] TEP Glossary

Module Content	This module offers an in-depth examination of contemporary artificial intelligence, combining foundational theory with practical implementation. It begins with a survey of advanced generative modelling, including adversarial networks, flow-based techniques and energy-based frameworks, alongside self-supervised strategies that extract meaningful representations from unlabelled data. The curriculum then advances to reinforcement learning, where intelligent agents are formulated as Markov Decision Processes and trained via Monte Carlo, Temporal-Difference, Q-learning, SARSA, Deep Q-Networks, policy-gradient methods and actor–critic algorithms. Emphasis is placed on stabilisation techniques such as experience replay, target networks and variance reduction, ensuring reliable performance in complex environments. Mid-term content addresses transformer architectures and graph neural networks. Transformer sessions cover subword tokenisation, attention mechanisms, adapter modules (e.g. LoRA) and approaches to model distillation in both natural language and vision contexts. Graph neural network lectures explore spectral and spatial graph convolutions, message-passing schemes, attention-based graph models and applications in social network analysis, knowledge graph reasoning and molecular property prediction. Subsequent topics include meta-learning—with optimisation-based (MAML, Reptile) and metric-based (prototypical, Siamese) algorithms for rapid task adaptation—and explainable and fair AI, utilising LIME, SHAP and demographic-parity techniques to interpret model behaviour and mitigate bias. The final segment examines privacy-preserving and safe AI. Methods for differential privacy, secure multi-party computation, homomorphic encryption and federated learning are studied alongside adversarial-robustness approaches, uncertainty quantification and out-of-distribution detection. Ethical frameworks, AI governance policies and real-world deployment challenges conclude the module. Instructional materials include concise lecture handouts, Jupyter notebooks with end-to-end code examples and curated research readings. Bi-weekly laboratory workshops reinforce each two-week block, enabling hands-on engagement with all major techniques. Upon completion, participants will possess the skills necessary to design, implement and critically evaluate advanced AI systems in research or industry settings.
Teaching and Learning Methods	• **Lectures (2–3 per week):** Structured presentations introduce core concepts, mathematical foundations and algorithmic frameworks. • **Laboratory Workshops (bi-weekly, 2½ hrs):** Hands-on coding sessions in Jupyter notebooks, guided by TAs, consolidating the preceding two weeks' topics through practical exercises.

	• **Tutorials / Seminars (weekly 30 min):** Small-group meetings for paper discussions, problem-solving, code debugging and conceptual Q&A.
Assessment Details[2] **Please include the following:** • **Assessment Component** • **Assessment description** • **Learning Outcome(s) addressed** • **% of total** • **Assessment due date**	see table below

Assessment Component	Assessment Description	LO Addressed	% of total	Week due
Lab Exam 1 (Practical)	Covering advanced generative modelling, self-supervised learning and foundational reinforcement-learning methods.	LO 1, LO 2, LO 3	30%	Week 6
Lab Exam 2 (Practical)	Covering policy-gradient algorithms, transformer fine-tuning and core graph-neural-network techniques.	LO 4, LO 5, LO 6	30%	Week 10
Final Examination	Closed-book computer-based exam with short-answer and problem-solving questions spanning meta-learning, explainability, fairness, privacy-preserving AI and robustness.	LO 7, LO 8, LO 9, LO 10	40%	University Exam Period

Reassessment Requirements

Exam based

Contact Hours and Indicative Student Workload[2]

Contact hours: 43

- **Lectures**: 11 weeks × 2.5 hrs = 27.5 hrs
- **Workshops**: 5 × 2.5 hrs = 12.5 hrs
- **Tutorials**: 6 × 0.5 hrs = 3 hrs

[2] TEP Guidelines on Workload and Assessment

	Independent Study (preparation for course and review of materials): 43 **Independent Study (preparation for assessment, incl. completion of assessment):** 39
Recommended Reading List	• Goodfellow et al. (2014), “Generative Adversarial Networks.” • Sutton & Barto (2018), Reinforcement Learning: An Introduction, 2nd ed. • Vaswani et al. (2017), “Attention Is All You Need.” • Kipf & Welling (2017), “Semi-Supervised Classification with Graph Convolutional Networks.” • Finn et al. (2017), “Model-Agnostic Meta-Learning for Fast Adaptation of Deep Networks.”
Module Pre-requisite	EE5C16 DEEP LEARNING AND ITS APPLICATIONS
Module Co-requisite	
Module Website	
Are other Schools/Departments involved in the delivery of this module? If yes, please provide details.	
Module Approval Date	
Approved by	Prof. Naomi Harte
Academic Start Year	September 2025
Academic Year of Date	2025/2026