COURSE SPECIFICATION DOCUMENT

Academic School / Department: Science, Innovation & Technology

Programme: MSc Artificial Intelligence

***FHEQ* Level:** 7

Course Title: Machine and Deep Learning Systems

Course Code: COMP 7101

Total Hours: 200 (Lev 7) (4 US Credits)
Timetabled Hours: 39
Learning Hours 21
Independent Learning Hours: 140

Credit 20 UK CATS credits
10 ECTS credits
4 US credits

Course Description:
Machine and Deep Learning offers an in-depth and practice-focused exploration of state-of-the-art neural network methods and their deployment in real-world scenarios. The course develops advanced expertise in modelling complex, high-dimensional data using modern architectures such as convolutional networks, recurrent networks, transformers, and generative models. Students learn to design, train, evaluate, and deploy deep learning systems using industry-standard frameworks (e.g., PyTorch or TensorFlow). Strong emphasis is placed on rigorous experimental methodology, optimisation strategies, interpretability, and responsible use of large-scale models.

Prerequisites:
None

Aims and Objectives:
Aim: Develop advanced understanding of deep learning theory, architectures, and optimisation strategies.
Objectives:

- Enable students to design, implement, and train neural networks using modern software frameworks.
- Equip students with the ability to evaluate architectures for vision, language, and multimodal domains.
- Provide hands-on experience in configuring training pipelines, hyperparameter tuning, and debugging.

- Foster expertise in experimental design, interpretability, and critical assessment of model behaviour.
- Prepare students for research, engineering, and industrial applications of deep learning technologies.

Programme Outcomes:
A2, A4, A5, B1, C1, C2, D1,D5

A detailed list of the programme outcomes are found in the Programme Specification. This is located at the archive maintained by Registry and found at:
https://www.richmond.ac.uk/programme-and-course-specifications/

Learning Outcomes:

1. Critically analyse deep learning architectures including CNNs, transformers, and generative models and evaluate their suitability for specific AI applications. (A2, B1)
2. Design, implement, and optimise complex neural networks using PyTorch/TensorFlow, following professional engineering practices. (A4, C1, C2)
3. Apply advanced training and regularisation strategies such as learning-rate scheduling, batch-normalisation, augmentation, and transfer learning. (B1, C2)
4. Conduct rigorous experimental evaluations using appropriate metrics, validation frameworks, and interpretability tools. (A5, C2)
5. Develop applied deep learning solutions for image, text, and multimodal tasks while demonstrating professional judgement. (C1, D1, D5)

Indicative Content:

- Neural network foundations and modern training dynamics
- Optimisation strategies: SGD variants, AdamW, schedulers, regularisation
- Convolutional neural networks and architecture evolution (LeNet,ResNet, ConvNeXt)
- Sequence modelling: RNNs, LSTMs, attention, transformers
- Generative modelling: autoencoders, VAEs, GANs, diffusion models
- Explainability: saliency maps, Grad-CAM, SHAP
- Metrics, experimental design, and error analysis
- Model compression, quantisation, and deployment considerations

Assessment:
This course conforms to the University Assessment Norms approved at Academic Board and located at: https://www.richmond.ac.uk/university-policies/

Teaching Methodology:
Delivered through lectures, programming labs, problem-based learning, weekly exercises, research-paper discussions, guided independent study, and guest speakers. Students engage in hands-on experimentation and iterative model development supported by structured feedback.

Indicative Text(s):

- Chollet, F. (2023). *Deep learning with Python*. 2nd edn. New York: Manning Publications.
- Géron, A. (2023). *Hands-on machine learning with Scikit-Learn, Keras & TensorFlow.* 3rd edn. Boston: O'Reilly Media.
- Goodfellow, I., Bengio, Y. and Courville, A. (2016). *Deep learning*. Cambridge, MA: MIT Press.

Journals

- *Journal of Machine Learning Research (JMLR).*
- *Applied Artificial Intelligence.*

See syllabus for complete reading list.

Change Log for this CSD:

Nature of Change	Date Approved & Approval Body (School or AB)	Change Actioned by Registry Services
Guided Learning Hours menu updated	October 2025	
Total Hours Updated	October 2025	