Advanced Artificial Intelligence and Machine Learning: Deep Unsupervised Learning

LMH Summer Programmes are provided by Lady Margaret Hall, a college in the University of Oxford

Course:	Advanced Artificial Intelligence and Machine Learning: Deep Unsupervised Learning
Available:	Programme Session 1: 30th June 2025 to 18th July 2025
Lectures:	18 Hours
Seminars:	12 Hours
Tutorials:	3 Hours
Independent Study:	Approximately 120 Hours
Recommended Credit:	15 CATS / 7.5 ECTS / 4 US Credits

About this Course:	Deep Unsupervised Learning is an exciting emerging area of research in the field of artificial intelligence and machine learning, in which the goal is to develop systems that can learn from unlabelled data. Such systems closely mimic natural human intelligence by finding patterns in data without instructions on what to look for. The course will begin with an introduction to unsupervised learning and clustering algorithms, before exploring generative adversarial networks and deep generative models. You will examine self-supervised learning, anomaly detection, flow-based models, and unsupervised representation learning. The final part of the course focuses on clustering in high-dimensional spaces, semi-supervised learning, energy-based models, and unsupervised learning for reinforcement. This intensive course offers theoretical understanding and practical experience with a focus throughout on real-world applications of deep unsupervised learning across various domains, offering career skills as well as excellent foundations for future research.

Course Overview:	Week 1
	• Introduction to Unsupervised Learning
	    ○ Differences between Supervised and Unsupervised Learning.
	• Clustering Techniques
	    ○ K-Means, Hierarchical, and DBSCAN clustering algorithms.
	    ○ Silhouette scores and Adjusted Rand Index for evaluating clustering results.
	• Dimensionality Reduction in Unsupervised Learning
	• Principal Component Analysis
	    ○ PCA workings and its applications in data compression and visualization.

- Autoencoders and Neural Networks
  - Neural Networks and their application in autoencoder architectures.
- Variational Autoencoders
  - VAEs and their role in generating new data samples.
- Generative Adversarial Networks
  - GANs, their architecture, and their use in generating realistic data.
- Conditional GANs and Deep Generative Models
  - Conditional GANS and image-to-image translation.
  - Advanced generative models, including PixelCNN and WaveGAN.

Week 2

- Self-Supervised Learning and Contrastive Learning
  - Self-supervised learning and its applications in pre-training.
  - Contrastive loss and its role in training models for self-supervised learning.
- Anomaly Detection and One-Class SVM
  - Techniques for anomaly detection, including isolation forests and autoencoder-based methods; pros and cons of the One-Class SVM approach to anomaly detection.
- Flow-Based Models, Real NVP, and Glow
  - Normalizing flow models and how they can be used for density estimation and sampling; architectures of particular flow models, including Real NVP and Glow.
- Unsupervised Representation Learning and Transfer Learning with Unsupervised Data
  - Importance of learned representations and how they benefit downstream tasks.
  - How unsupervised pre-training can enhance transfer learning in various domains.

Week 3

- Clustering in High-Dimensional Spaces
  - Techniques for clustering when dealing with high-dimensional and sparse data.
- Subspace Clustering
  - Spectral clustering and other methods tailored for high-dimensional data.
- Semi-Supervised Learning, Self-Training, and Pseudo-Labelling
  - Methods combining unsupervised and supervised learning for better performance.
  - How self-training and pseudo-labelling work in semi-supervised scenarios.
- Energy-Based Models and Diffusion Models
  - Energy-based models and their use in modelling complex data distributions.
  - Diffusion probabilistic models and score-based diffusion models and they are used for generation.
- Unsupervised Learning for Reinforcement
  - How unsupervised learning can aid in reinforcement learning and policy improvement.
- Applications of Unsupervised Learning
  - Real-world applications of deep unsupervised learning across various domains.

Lady Margaret Hall, University of Oxford
Norham Gardens, Oxford, OX2 6QA, United Kingdom
vacation.programmes@lmh.ox.ac.uk
+44 (0)1865 274295

Key Texts:	Goodfellow, I., Bengio, Y., and Courville, A., *Deep Learning*, Cambridge MA, 2016. Croitoru, F-A., Hondru, V., Tudor Ionescu, R., and Shah, M., "Diffusion Models in Vision: A Survey", in *IEEE Transactions on Pattern Analysis and Machine Intelligence* 14(8): 1-25, 2022. Kingma, D., *An Introduction to Variational Autoencoders*, Ithaca NY, 2019.
Learning Outcomes:	By the end of this course, you will: • Understand the differences between supervised and unsupervised learning and the fundamentals of clustering. • Be able to utilise a range of algorithms and techniques for unsupervised, self-supervised, and semi-supervised learning. • Be able to evaluate the efficacy of real-world applications of deep unsupervised learning across various domains. • be able to demonstrate familiarity with the current state of research into deep unsupervised learning.
Admissions Requirements:	LMH Summer Programmes are designed for students who want to gain and develop knowledge in their chosen subject area. LMH Summer Programmes are intensive courses of study aimed at undergraduates who have completed one, two, or three years of their degree, or entry level postgraduate students. We will consider each applicant's academic ability and expect successful applicants to have a minimum grade point average equivalent to 2:1 level on the British grading scale. For example, this would mean at least a 3.2 GPA on the 4.0 grading scale in the United States, and 80% in China. **This course would suit STEM students with intermediate level experience in artificial intelligence and machine learning concepts and techniques, including those undertaking, or looking ahead to, graduate level study or research.** **Specifically, students on this course must have experience of the following topics:** • **Knowledge of the deep learning libraries.** • **Understanding of deep learning, recurrent neural networks, and convolutional neural networks.** • **Strong background in optimization and probability.** • **Familiarity with the Python programming language.** To participate fully in the programme all students will need to have proficiency in English. English language requirements for students who are not native English speakers: • TOEFL iBT score of 98 • IELTS score of 7.0 (no less than 6.5 in each component) • Duolingo English Test score of 125 (no less than 115 in each section) • Cambridge English Scale score of 185 If the language of instruction in your home institution is English you do not need to provide evidence of your English proficiency.
Teaching Methods:	Core syllabus material will be covered in lectures. Students attend four lectures each week and each lecture lasts 90 minutes. Seminars in smaller groups offer students space to discuss and debate, to dig deeper into difficult concepts, and to explore their own ideas. Student contribution to seminars is vital, and tutors will ensure everyone takes part in discussions. Seminars last 1 hour and students will take part in four seminars each week. Independent study is a crucial part of an LMH Summer Programme and of the Oxford teaching model. Tutors will recommend important reading to do between

	lectures and seminars that will enable students to come to class equipped to understand the information presented and prepared to take part in discussion and debate. Each week students will have an assignment of independent work to complete and submit in advance of the tutorial. There is an appropriate amount of space in the timetable to complete the necessary reading, preparation, and assignments. Students should expect to do around 40 hours of independent study each week. The final class each week is a tutorial, a very small class typically including only 2-4 students and central to the teaching methods used by the University of Oxford and on LMH Summer Programmes. Guided by their tutor, students will receive feedback on their assignments and be challenged to defend, justify, or even rethink their work and ideas. These rigorous academic discussions help develop and facilitate learning in a way that cannot be done with lectures and seminars alone.
Assessment:	On a three-week LMH Summer Programme students produce one piece of assessed work every week, which is submitted to the tutor and then discussed in a tutorial. At the end of each week students will receive a percentage grade for their submitted work. Each week's work counts for a third of the final percentage grade, so the final grade is an average of the mark received for each piece of work. Students who stay for six or nine weeks will receive a separate grade for each 3-week course.
Academic Credit:	Lady Margaret Hall will provide a transcript of students' assessed work, and can send this directly to your home institution if required. LMH Summer Programmes are designed to be eligible for academic credit, and we will communicate with home institution to facilitate this as needed. As a guide, we recommend the award of 15 CATS / 7.5 ECTS / 4 US Credits for each 3-week course.