

# Advanced Artificial Intelligence and Machine Learning: Computer Vision

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

Course:	Advanced Artificial Intelligence and Machine Learning: Computer Vision
Available:	Programme Session 2: 15 <sup>th</sup> July 2024 to 2 <sup>nd</sup> August 2024
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:	<p>From self-driving cars and augmented reality to intelligent medical imaging helping doctors identify diseases more quickly, computer vision is a rapidly-growing field within artificial intelligence and machine learning. In this course, students who are already familiar with the key theoretical foundations of artificial intelligence and machine learning will dive deeper into the exciting capabilities of this area of research and its applications.</p> <p>You will begin with computer vision algorithms for classification, recognition, detection, and their implementation in deep learning libraries, before exploring autoencoders and variational autoencoders, and gaining insights into the training and application of generative adversarial networks. You will proceed to an in-depth examination of diffusion models, including score-based diffusion models, latent diffusion models, and Stable Diffusion. The final part of the course explores even more advanced topics, including the representation of 3D objects, vision transformers, video classification, and text to image generation.</p> <p>This intensive course offers students theoretical understanding and practical experience in a range of advanced computer vision concepts and techniques, offering career skills as well as excellent foundations for future research.</p>
Course Overview:	<p>Week 1</p> <ul style="list-style-type: none"> <li>• Computer Vision Supervised Machine Learning Applications             <ul style="list-style-type: none"> <li>○ Computer Vision algorithms for classification, recognition, detection, and their implementation in Deep Learning libraries.</li> </ul> </li> <li>• Unsupervised Learning: Autoencoders and Variational Autoencoders             <ul style="list-style-type: none"> <li>○ Unsupervised Learning and Autoencoders, including architecture and training for feature extraction and reconstruction.</li> </ul> </li> </ul>

	<ul style="list-style-type: none"> <li>○ Variational Autoencoders, their probabilistic approach, and how to implement them for latent Representation Learning.</li> <li>• Generative Adversarial Networks <ul style="list-style-type: none"> <li>○ GANs, their architecture, training dynamics, and application in data generation.</li> </ul> </li> <li>• Conditional GANs, CycleGANs, and StyleGANs <ul style="list-style-type: none"> <li>○ Advanced GANs and their distinct contributions to controlled data synthesis and image generation.</li> </ul> </li> </ul> <p>Week 2</p> <ul style="list-style-type: none"> <li>• Diffusion Models <ul style="list-style-type: none"> <li>○ Diffusion Models, their sample generation capabilities, and their applications across domains.</li> </ul> </li> <li>• Score-based Diffusion Models <ul style="list-style-type: none"> <li>○ Score-based Diffusion Models, their training process, the role of score functions, and their advantages in generative modelling.</li> </ul> </li> <li>• Latent Diffusion Models / Stable Diffusion <ul style="list-style-type: none"> <li>○ The concept of Stability, and how latent variables improve Diffusion Models.</li> </ul> </li> <li>• 3D Data <ul style="list-style-type: none"> <li>○ 3D object representation, and Mesh Cloud Generation.</li> </ul> </li> </ul> <p>Week 3</p> <ul style="list-style-type: none"> <li>• Recurrent Neural Networks and Attention Mechanism <ul style="list-style-type: none"> <li>○ RNNs, Attention, and Self-Attention.</li> </ul> </li> <li>• Vision Transformers <ul style="list-style-type: none"> <li>○ Vision Transformers and their applications in Computer Vision.</li> </ul> </li> <li>• Learning from Videos <ul style="list-style-type: none"> <li>○ Video classification, activity recognition, automated driving, and trajectory prediction.</li> </ul> </li> <li>• Vision and Text <ul style="list-style-type: none"> <li>○ Image/Video captioning, CLIP, and Text to Image Generation.</li> </ul> </li> </ul>
Key Texts:	<p>Goodfellow, I., Bengio, Y., and Courville, A., <i>Deep Learning</i>, Cambridge MA, 2016.</p> <p>Croitoru, F-A., Hondru, V., Tudor Ionescu, R., and Shah, M., “Diffusion Models in Vision: A Survey”, in <i>IEEE Transactions on Pattern Analysis and Machine Intelligence</i> 14(8): 1-25, 2022.</p> <p>Kingma, D., <i>An Introduction to Variational Autoencoders</i>, Ithaca NY, 2019.</p> <p>Olah, C., and Carter, S., <i>Attention and Augmented Recurrent Neural Networks</i>, Distill, 2016.</p>
Learning Outcomes:	<p>By the end of this course, you will:</p> <ul style="list-style-type: none"> <li>• Understand computer vision algorithms for classification, recognition, and detection, and their implementation in deep learning libraries.</li> <li>• Know the different types of generative adversarial network and their distinct contributions to controlled data synthesis and image generation.</li> <li>• Be able to identify different diffusion models and assess their advantages in generative modelling.</li> <li>• Be able to demonstrate awareness and understanding of the latest key research areas in computer vision.</li> </ul>
Admissions Requirements:	<p>LMH Summer Programmes are designed for students who want to gain and develop knowledge in their chosen subject area. LMH Summer Programmes are intensive</p>

	<p>courses of study aimed at undergraduates who have completed one, two, or three years of their degree, or entry level postgraduate students.</p> <p>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.</p> <p><b>This course would suit STEM students with intermediate level experience in artificial intelligence, machine learning, and computer vision concepts and techniques, including those undertaking, or looking ahead to, graduate level study or research.</b></p> <p><b>Specifically, students on this course must have experience of the following topics:</b></p> <ul style="list-style-type: none"> <li>• Knowledge of the deep learning libraries.</li> <li>• Understanding of deep learning and convolutional neural networks.</li> <li>• Strong background in optimization and probability.</li> <li>• Familiarity with the Python programming language.</li> </ul> <p>To participate fully in the programme all students will need to have proficiency in English.</p> <p>English language requirements for students who are not native English speakers:</p> <ul style="list-style-type: none"> <li>• Overall TOEFL score of 85;</li> <li>• or IELTS score of 6.5 (no less than 6.0 in each component);</li> <li>• or CET-4 at 550 or CET-6 at 520.</li> </ul> <p>If the language of instruction in your home institution is English you do not need to provide evidence of your English proficiency.</p>
Teaching Methods:	<p>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.</p> <p>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.</p> <p>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.</p>
Assessment:	<p>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</p>

	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.