AI University Public Sector

Online Learning Module for Legal Professionals

A ten-week module that gives you a firm grasp of the technical, legal and ethical foundations of artificial intelligence and their applications in the public sector, including algorithmic decision making.

What is the module about?

AI University Public Sector, led by University of Oxford experts, is a series of three learning units lifting the lid on artificial intelligence (AI) from both computer science and legal viewpoints. The module is aimed at public sector legal professionals, regardless of experience.

This ten-week long exclusive series, delivered by the Oxford LawTech Education Programme (OLTEP), first gives participants a foundational, technical understanding of modern AI from a legal perspective. It then explores, in detail, selected legal issues pertaining to algorithmic decision making, a central application of AI in the public sector. Finally, it gets the participants up to speed with the central positions and trends in digital ethics that underpin and complement AI regulation.

What will you gain?

On completing this module, participants will:

- have developed a foundational understanding of how artificial intelligence works and an awareness of strengths, limitations, and challenges,
- understand the role of public law in ensuring optimal use of algorithmic decision making (ADM) systems,
- have gained applicable insights into the central positions and trends in digital ethics,
- be able to identify the challenges and opportunities of AI and gain the ability to assist clients through knowledge and insight,
- have had the opportunity and gained the confidence to engage with other professionals working in the field of AI,
- have exclusive access to the Oxford LawTech Education Programme (OLTEP) alumni network,
- have earned a personal certificate from the module that highlights their understanding of the topic.

Register your interest in further OLTEP activities

KEEP ME INFORMED

Is this module for you?

AI University Public Sector is designed for public sector legal professionals. It is also suited for policymakers and those who will engage with legal issues relating to artificial intelligence and algorithmic decision making in their profession.

The module is suited for professionals at all levels of experience and seniority who want to develop their skills and knowledge in the field of AI. No prior technical or legal knowledge of artificial intelligence is required.

If you are working in the private sector, you may want to participate in the AI University Foundations module.

How does the module work?

AI University Public Sector is delivered over ten weeks to a cohort of enrolled participants, with limited numbers to ensure a high-quality learning experience for everyone. This enables us to host live tutorial discussion sessions, conducted personally by our expert lecturers, and provides an opportunity for professional networking. Participants can enrol individually or join through a group booking that has been made by their organization.

The module is delivered entirely online, with elements of synchronous and asynchronous learning. Each learning unit includes on-demand video lectures, a downloadable copy of the lecture materials designed for convenient notetaking, core and further readings, key takeaways, and self-assessments. All these can be studied at any convenient time, so that you can progress your learning on your own schedule and alongside a busy professional life. All core reading materials are open access, so that you can study them at no extra cost and retain them for future reference.

In addition, you can put questions directly to the lecturers—and engage in discussion with other members of the cohort—at any time, on our online discussion forum. These interactions are visible to all participants in the cohort, ensuring that everyone learns from the discussion and the questions that are asked and answered.

You will also have the opportunity to join cohort-wide, live online discussions with our expert lecturers. These are hosted by live videoconferencing in four 60-minute sessions, scheduled regularly during the module. Each session is run twice on each occasion, at different times of the day, giving the best opportunity for our participants across the globe to attend.

The module concludes with an optional, hands-on capstone exercise, the successful submission of which leads to a certificate of completion.

The minimum expected time commitment is around 15 hours spread over the duration of the whole module or, on average, 1.5 hours weekly. This includes 6 hours of video lectures and 4.5 hours of live sessions. That said, many participants choose to dedicate up to 20 hours in total to their learning (or around 2 hours weekly).

Our module handbook guides you through the module and provides advice on scheduling your studies, to make your learning as successful as possible. You will receive regular email signposts on the recommended progression journey.

What does the capstone exercise involve?

The capstone exercise is a short, hands-on written assignment that participants can complete after they have studied all the learning units in the module and submit to the OLTEP Team for review and feedback. Its pedagogical function is to reinforce your learning across all units of the AI University Public Sector module. We find, in our professional teaching at Oxford, that when our students have the opportunity to produce something written and test how to apply the newly gained knowledge, this greatly consolidates and deepens their learning.

Participants who do the exercise will spend no more than 1.5 to 2 hours completing it during the last three weeks of the module. They will submit a write-up of three pages to us by the end of the module. We review these and provide feedback, though not a numerical grade. Successful completion of the exercise, along with successfully passing all the self-assessment components in the module, allows us to issue a certificate of completion.

OLTEP certificates are issued as both a PDF file that you can download and an online certificate that anyone can access via a unique URL. We automatically verify all online certificates for 365 days from the date they are issued, thereby attesting to your up-to-date knowledge in the field. You can, of course, download the PDF certificate and use it as permanent evidence of your professional development.

Participants whose exercise submission falls substantially short of the required standard are given one opportunity to revise and resubmit the exercise, within three weeks. The feedback provided for the first submission will give clear guidance on what needs to be improved.

How much will it cost and when is the next cohort?

We run the AI University Public Sector module regularly throughout the year. For the latest information on dates and prices, see the OLTEP website. We also send email announcements to all those who have registered their interest by signing up to our mailing list on the OLTEP website.

We offer a 10% discount for group bookings of at least 10 places. We also offer a 15% discount for participants working in public-sector or non-commercial organizations, a 10% discount for OLTEP alumni, and a 5% discount for the University of Oxford alumni. Only one of these discounts can apply for any booking.

Applicants with personal financial barriers to participation can seek support from one of our OLTEP Bursaries.

If you wish to enquire about an exclusive run of the module for your group or company, please contact us at info@oltep.ox.ac.uk.

Testimonials – AI University Public Sector

“I really enjoyed the OLTEP AI module. …. [T]he module made me think about AI in a different way than I do on a day-to-day basis in my work. I would highly recommend the course!” – Maartje Nugteren, International Policy Lead, Centre for Data Ethics and Innovation.

It was great to have the instructors "[b]reaking things into categories clearly, so we have the blocks of understanding". – participant feedback, Government Legal Department.

“This was a very practical course - both in terms of content but also how it is designed to fit in with working lives.” – participant feedback, Government Legal Department.

“A fully holistic and comprehensive approach to AI, which was interesting and refreshingly different to other more basic and generic AI training I have been on.” – participant feedback, Government Legal Department.

“Fantastic leaders, great delivery of lectures and talks, …” – participant feedback, Government Legal Department.

“The course was excellent and I think all government lawyers need to do it!” – participant feedback, Government Legal Department.

Register your interest in further OLTEP activities
KEEP ME INFORMED

THE LAWYER'S AI PRIMER

This unit introduces law professionals to the computer science foundations of modern 'artificial intelligence' (AI). Using non-technical language and accessible examples, it explains Artificial General Intelligence (AGI) and Applied AI—the two principal aims of AI research, past and present. It introduces the two primary technical approaches to AI and the main kinds of machine learning, and it explains the technology of artificial neural networks. The unit then discusses how we evaluate the quality of AI machine learning models. It builds on these ideas to introduce the technology that underpins today's generative AI based on large language models, and their chatbot applications. It concludes by highlighting the main risks we have with using AI in the real world, elucidating their technical origins and the limits of technical solutions available today.

Lecturer: Tom Melham
Contact: tom.melham@cs.ox.ac.uk.

Lecturer bio: Tom Melham is a Professor of Computer Science at the University of Oxford and a Fellow of Balliol College, where he is Tutor in Computation. During 2012–15, he was Associate Head of Oxford's Mathematical, Physical and Life Sciences Division. He was elected a Fellow of the Royal Society of Edinburgh in 2002 and a Fellow of the British Computer Society in 2015 and is a Chartered Engineer. His research interests include AI and technology for legal services and the justice system, and testing and evaluation of AI-based systems.

Learning outcomes:

- A foundational understanding of how AI works and an awareness of strengths, limitations, and challenges.
- Awareness of the differing aims and approaches of modern artificial intelligence.
- Acquaintance with the three main kinds of modern machine learning, a prominent branch of AI today.
- A non-technical grasp of machine learning models and their key characteristics.
- Knowledge of how we evaluate the quality of AI models.
- Understanding of the basic technology underpinning generative AI, such as GPT models, and what this technology can and cannot do.
- A technically informed and balanced understanding of key AI concepts such as prompting, reasoning models, and agentic AI.
- Awareness of the main pitfalls and challenges in deploying modern AI in practice, including problems of bias and explainability.

Register your interest in further OLTEP activities
KEEP ME INFORMED

Algorithmic Decision Making and Public Law

The attractions of algorithmic decision making (ADM) are obvious; its efficiency makes it very appealing as an adjunct to, or replacement for, human decision-making, particularly in a time of austerity. ADM systems can also offer other advantages, such as improved accuracy. But the deployment of ADM can also cause problems on a huge scale.

Public law therefore has a significant role to ensure we harnesses the benefits of ADM while minimising risks and mitigating them when they do arise. This unit examines how public law might respond and develop existing grounds of judicial review and how we might make use of legislation such as the GDPR and Public Sector Equality Duty to enable the law to perform this crucial task. The public law framework is also critical for the governance and operation of algorithmic decision making in the private sector, and so this unit will be relevant to anyone interested in ADM more generally.

Lecturer: Rebecca Williams
Contact: rebecca.williams@law.ox.ac.uk

Lecturer bio: Rebecca Williams is a Professor of Public Law and Criminal Law at the University of Oxford, in association with Pembroke College. Rebecca's principal teaching interests are criminal law and public law. Her work includes examining optimum methods of decision-making and the use of criminal law as a form of regulation. Increasingly her work also focuses on the relationship of law and technology and the ways in which the law will need to develop to keep pace with technological developments. Her work has been cited by courts in the UK, EU and Australia.

Learning outcomes:

- Understanding of what public law is and what it is supposed to do.
- A solid grasp of the basic principles of public law.
- Understanding of how and why Algorithmic Decision Making (ADM) raises new challenges for public law.
- Knowledge of how relevant legislation (GDPR, Public Sector Equality Duty) applies to ADM.
- Understanding of the grounds of judicial review and how these will apply to ADM, including relevant case law.
- Appreciation of the kinds of issues that public authorities will need to consider when adopting or deploying ADM.
- Knowledge of the potential for the use of public law reasoning in contexts where ADM is used.

DIGITAL ETHICS

The development and deployment of AI systems require ethical judgment, especially when law does not answer relevant questions, or when it answers them unsatisfactorily. What should your client do? And why? Non-action and non-advice are not an option. The chief aim of this unit is to equip lawyers with the skills and conceptual toolbox to face difficult question around the practical applications of AI systems.

This unit introduces digital ethics as an academically rigorous yet profoundly practical discipline that gives you confidence to spot and solve ethical puzzles concerning simple as well as complex AI cases or scenarios. Participants gain an understanding of the 'methods of ethics' and familiarise themselves with major trends in digital ethics. They also learn three practical lessons that will help them deploy digital ethics in their daily work.

Lecturer: Václav Janeček
Contact : vaclav.janecek@cs.ox.ac.uk

Lecturer bio: Václav Janeček is a Senior Lecturer in Law at the University of Bristol and an Associate Member of the Oxford University's Department of Computer Science. Having previously researched how lawyer and computer scientists can work together more efficiently, Václav is currently leading a new master's programme in Law, Innovation and Technology. In the past, Václav held principal research appointments at the Oxford Law Faculty, the Oxford Internet Institute's Digital Ethics Lab, and the Masaryk University's Institute for Law and Technology.

Learning outcomes:

- Understanding of the distinction between law and ethics and their relevance for the development, deployment, and maintenance of AI systems.
- A firm grasp of digital ethics as a practical and rigorous discipline.
- Familiarity with the basic trends and major debates in digital ethics.
- Understanding of the 'method of ethics', which includes how to spot an ethical issue and how to go about solving it.
- Knowledge, gained through three practical lessons, of how to deploy digital ethics in daily work as a legal advisor, policymaker and innovation champion.
- Ability to track developments in digital ethics and AI regulation.
- Awareness of strengths and weaknesses in AI regulatory and AI alignment approaches.

About OLTEP

The Oxford LawTech Education Programme (OLTEP) was founded as a joint initiative between the Oxford Law Faculty the Oxford Department of Computer Science. Our mission is simple. We educate future leaders in the legal market: adept providers of tech-enhanced legal services and confident advisers who can spot, analyse and utilise trends in digital technology.

Read an article about us on the University of Oxford Social Sciences Division website.

Over 1,300 participants from private, public and third sector organisations (including the Government Legal Department, Slaughter and May, BUPA, Burges Salmon, the Law Commission, and the Office for AI) have successfully completed our education programme.

Our approach is research-driven and builds on our world-leading expertise in digital technologies and relevant issues regarding both substantive law and the practice of law. To see examples of our work, visit the Unlocking the Potential of Artificial Intelligence for English Law project webpage.

OLTEP is a non-degree programme of professional education operated by Oxford University Innovation Limited (Company No. 2199542), a wholly owned subsidiary of the University of Oxford, and directed by experienced Oxford academics. If you are looking for a formal qualification from Oxford University, please see the University's website.

Register your interest in further OLTEP activities
KEEP ME INFORMED