



# Artificial Intelligence, Education and Assessment at UCL Laws

## Current Thinking and Next Steps for the UK Legal Education Sector

We are academics at the Faculty of Laws, University College London involved in the development of educational policy and practice concerning artificial intelligence. In this discussion paper, we lay out our thinking on how the legal education sector should respond to generative AI. We are very supportive of thoughtful use of technology to change our world and our work for the better, but developing excellent legal expertise does and will remain crucial for lawyers, whether working with AI tools or in the many situations where such tools will be less helpful, absent or limited. There is a critical balance between using AI as a true study aide, facilitating meaningful learning, and using it to cognitively offload tasks in a way that hinders learning. We believe this balance to be delicate, and that academic judgement is crucial to achieving it. By highlighting relevant abilities and skills we regard as essential outcomes of learning the law, we identify the barriers AI poses to their acquisition, and the implications the technology has for assessment. We explain the thinking behind the policies we have enacted as a Faculty, including the requirement that all modules have 50–100% secured assessment—forms of assessment that reliably safeguard against the use of generative AI. Our position is that, in assessing our degrees, integrity takes priority and should not be put at risk, even as we recognise the value of diverse forms of assessment. More broadly, we see a compelling case for our sector to come together to respond to the future, double down on its values, cease passively responding to the ongoing evolution of business models of AI and cloud companies, and instead actively work together to determine the sector's direction and to assert a positive and ambitious vision of learning well in a rapidly changing technological environment.

**Cite as:** Michael Veale, Isra Black, Mark Dsouza, Matthew Fisher, Mujtaba Ghaus, Thea Gibbs, Orla Lynskey, Colm O'Conneide, Eloise Scotford, Olga Thomas, Alina Trapova, Kimberley Trapp (2025) 'Artificial Intelligence, Education and Assessment at UCL Laws: Current Thinking and Next Steps for the UK Legal Education Sector', *UCL Legal Studies Research Paper Series 04/2025*.

### The challenge

#### Old and fundamental skills for new AI

- Wading through cheap text
- Standing out from the AI slop
- Knowing what to ask
- Connecting with each other
- Working with the lights off

#### Challenges of AI to skill acquisition

- Learning to learn
- Taking intellectual risks

### Assessment

- Integrity must be the priority
- Diversity and security
- Designing security into assessment questions fails us
- Policing does not ensure security
- Students should not be assessed as content creators

### Concluding thoughts

- Next steps for the sector

---

## The challenge

Our task as a law faculty is to ensure that our degrees are, and continue to be, both **transformative educational journeys** and **powerful, internationally recognised and durable signals of our students' achievements**. We design and maintain our programmes so that students obtain rigorous legal knowledge and skills for both predictable and novel situations; forge an ability to question received wisdom and formulate effective arguments; develop their capacity to independently, critically and analytically engage with complicated and often lengthy material; engage in contextual assessment and problem solving; communicate powerfully and effectively; and are exposed to spaces and experiences that allow them to develop their personal and professional values and ethics.

The emergence of new artificial intelligence (AI) systems does not change the core of that mission. AI systems, particularly generative systems which produce text and media, are the latest in a string of technologies to influence and apply pressure to legal education and practice. The way we respond to these pressures defines us as an institution of knowledge-building and dissemination.

The present situation is far from the first time that law schools have experienced pressure to shape legal education around particular tools. Legal research tools have long been big business, and external commercial dynamics have impacted education for some time. Expensive legal databases such as LEXIS were made available to law schools in the 1980s at discounted academic rates on condition that those schools provided specific instruction in them to ensure a pipeline of future customers.<sup>1</sup> In the face of this pressure to let industry tools determine the structure of education, legal academia chose to craft its own fate, pushing back against attempts to shape and enclose access to legal information and to provide free and universal access instead. Academics founded legal informatics institutes and similar initiatives throughout the 1990s and early 2000s, such as [LII](#) (Cornell Law School), [AustLII](#) (UTS and UNSW), [HKLII](#) (HKU), and in the UK and Ireland, [BAILII](#) (UC Cork and IALS, University of London). UCL Laws operated an extensive free online database of translated foreign statutes and judgments, including what was then the most [extensive English-language database of German case law](#).

Responding to AI similarly requires us to tread carefully and with purpose. Law schools must approach AI critically to avoid being instrumentalised as part of the ongoing hype-cycle. It is easy to be wowed by tales of economic growth, scientific discovery through AI, or the widespread transformation of legal work, including the prospect of reduced workloads. From what we have seen, talk of this transformation far outweighs current evidence that it is happening as advertised. Where AI tools have so far proven useful, they have been applied carefully, with as much attention paid to when not to use AI as to when to apply it. There is unambiguous promise in the use of AI in supporting legal information retrieval and work on mountains of documents, such as in redaction or disclosure review. The role of AI in more complex yet essential tasks, such as legal reasoning or complex drafting, is considerably less clear, particularly where the stakes of failure are high.

Our response as legal educators is informed by the unreliable nature of AI as it stands, but importantly, our thinking does not hinge on the current capabilities of these systems. We note that AI transcription systems hallucinate things that were not said;<sup>2</sup> that reviews of legal language model benchmarks reveal even the most advanced models perform

<sup>1</sup> Abdul Paliwala, 'Creating an Academic Environment: The Development of Technology in Legal Education in the United Kingdom' (1991) 5 International Review of Law, Computers & Technology 136, 146.

<sup>2</sup> Allison Koenecke and others, 'Careless Whisper: Speech-to-Text Hallucination Harms', *Proceedings of the 2024 ACM Conference on Fairness, Accountability, and Transparency* (Association for Computing Machinery 2024).

unsatisfactorily on basic legal tasks;<sup>3</sup> and that, when pitted against each other, AI summaries from frontier models have been found to be considerably worse than human-authored summaries on all criteria.<sup>4</sup> However, AI tools and their capabilities are in constant development, so we have been careful to ensure that none of our analysis below hinges on AI not being ‘good enough’ yet.

Even though AI might not always be a reliable or authoritative tool, it does not mean it will not change the sector and bring productivity gains in some respects. At the same time, there are risks for the development of the law, and for the provision of high-quality legal reasoning and services. The production of huge volumes of automated content presents challenges for a discipline based on carefully constructed text and closely scrutinised evidence. The access to justice crisis presents an opportunity for technology, but also a risk that those in need will flock to such solutions or services first, however inadequate they may be, with consequences that will resonate through the court system.<sup>5</sup> Where some parties to a dispute have access to tools and the expertise to use them with due precision, and others do not, there is further concern that this inequality of arms might exacerbate existing structural imbalances in the law.

Change to a sector can present opportunities for improvement. We are very supportive of thoughtful use of technology to change our world and our work for the better. Firms and commentators regularly talk about how transformative AI is and will be. Many of the forms of transformation they envisage would be progressive, and some extremely commercially advantageous for select AI businesses. These visions are rarely grounded in the context of the legal sector and discipline that we and our extended community know so well, or the values central to our educational mission as a law school. Our role is to cut through the hype around technological progress to understand and assess the opportunities and challenges any new technology brings, before making active decisions to help bring about the world we want to see. This paper is about those decisions.

In the sections below, we will first outline the **abilities and skills** we believe AI encourages us to heighten our focus on imparting as educators, particularly as teachers of law. However, these competences have to be imparted in the context of significant **acquisition challenges** that AI exacerbates. We do not just teach the substance of our disciplines. Higher education is about learning how to learn as much as it is about learning specific content and skills. We should not be complacent about AI’s effect on attitudes to, and capacities for, knowledge acquisition, and on the willingness to take intellectual risks. Finally, we consider another key societal role of universities – **robust signalling and assessment** of knowledge and skill acquisition. Dealing with these challenges together is not without tension, and some require careful consideration and trade-offs. We hope our thinking – and our signalling – will be useful to other institutions (both higher education institutions and beyond) as they develop their own responses.

---

## Old and fundamental skills for new AI

Graduates of the future will be working in a world where text and AI-quality responses are ‘cheap’. Our degrees are already designed to impart a wide variety of abilities and skills, and AI encourages us to further sharpen our efforts to highlight certain aspects of these skills to our graduates. In setting out below, and in the sections that follow, some aspects of these skills to which we believe particular attention should be paid in the face of AI, we

<sup>3</sup> Zhiyu Chen and others, ‘[A Survey on Large Language Models for Critical Societal Domains: Finance, Healthcare, and Law](#)’ (2024) 11 Transactions on Machine Learning Research.

<sup>4</sup> Amazon Web Services for Australia Securities and Investment Commission, [Generative Artificial Intelligence Document Summarisation Proof of Concept](#) (ASIC 2024).

<sup>5</sup> Natalie Byrom, [Where Has My Justice Gone?](#) (Nuffield Foundation 2024).

are (re)committing to the pedagogical vision and values outlined above. That is, legal education requires (at least) gaining basic legal knowledge, including a reflective understanding of the purpose(s) of law in society, as well as fundamental abilities and skills of critical thinking, principled argumentation, ethical assessment, written and interpersonal communication, and creative problem solving. This means that our graduates will be—as they need to be—equipped to master AI tools in their future as lawyers, having developed expertise to harness the power of this technology, to spot risks in its use, and to continue being excellent lawyers as social and technological forces continue to wax and wane. . In this context, which is presenting real challenges for how we organise, govern and regulate ourselves in society, we believe that well educated lawyers are needed more than ever. It should also be noted that the skills highlighted below are non-exhaustive. Our point is that thoughtful understanding of AI and its limits is critical in designing and delivering high quality legal education.

### **Wading through cheap text**

Knowing how to navigate a world of easy (and often questionable) information generation will be a highly valued skill (even more so than it is already). AI tools may play a support role here, but they will require creativity and critical thinking to use masterfully. Such creativity and critical analysis is a human trait that cannot be simply or practically offloaded to AI tools. While AI systems may help navigate the weight of documentation produced in legal proceedings, approaches to bury relevant information are likely to develop alongside an increased ability to navigate and search using computers.<sup>6</sup> We will guide students in building the skills to find useful information and critically appraise it in a world of cheap, abundant text, and prime them to develop and improve these skills in a changing world.

### **Standing out from the AI slop**

The online world is under siege, as banal ‘AI slop’ drowns out thoughtful and reliable content, and changes the way entire systems function.<sup>7</sup> It seems unlikely that courts will allow AI to swell filings to unreadable proportions in a similar way. In a world of generic and cheap text, we think that style, incisiveness, and parsimony will matter more than ever. Our degrees will train students to find and refine their voice, and to write effectively, including with new analysis and argument.

### **Knowing what to ask**

AI systems are tools, and tools need a purpose. Rigorous and responsible use requires the user to stay in control of the questions they ask for support with. This is all the more important as AI systems readily attempt poorly defined tasks, and are often sycophantic, in the sense that they confirm user framings and the biases inherent in the questions asked.<sup>8</sup> Users of these technologies, just like analysts in any complex task, need themselves to appraise the situation, prioritising and contextualising next steps. Our degrees will give graduates these problem-framing skills.

### **Connecting with each other**

When information becomes unreliable, overwhelming, or both, interpersonal skills will be crucial. Consider issues of access to justice. AI will not change the fact that people often have the most acute legal needs when life has hit them hardest, when they fall through the cracks in systems or face what seem like insurmountable challenges.<sup>9</sup> Digital divides

<sup>6</sup> See generally the field of adversarial machine learning: Apostol Vassilev and others, [Adversarial Machine Learning: A Taxonomy and Terminology of Attacks and Mitigations](#) (NIST 2024).

<sup>7</sup> Jason Koebler, [‘AI Slop Is a Brute Force Attack on the Algorithms That Control Reality’](#) (404 Media, 17 March 2025).

<sup>8</sup> Mrinank Sharma and others, [‘Towards Understanding Sycophancy in Language Models’](#) in the *Twelfth International Conference on Learning Representations* (ICLR 2024).

<sup>9</sup> Hazel Genn, [Paths to Justice: What people do and think about going to law](#) (Hart 1999).

relating to technology might even exacerbate issues of who law works for, and how.<sup>10</sup> A similar situation arises in the high-pressure environments of corporate boardrooms. The people skills of a solution-oriented lawyer can make all the difference between a high value corporate deal coming together or falling apart.

Whatever area our graduates choose to work in—inside or outside law—our degrees will prepare them with crucial interpersonal skills for these complex, multifaceted situations where human connection, rather than technological solutionism, fundamentally matters.

### **Working with the lights off**

Our graduates' clients or jurisdictions of practice may not be at the forefront of technology or data availability. Many jurisdictions are still at the earliest stages of the digitisation of law, let alone the use of AI. Many languages are deprioritised in the training of the most advanced models, and some linguistic communities are so small in terms of relative number of speakers that it seems unlikely that AI tools in these languages will ever match up in terms of data volumes or investment.<sup>11</sup> Many legal environments also generate specific challenges that may impact on AI usage. Lawyers working with sensitive evidence, or in contexts such as national security, are unlikely to have the standard array of AI tools at their disposal. In court as an advocate or in meetings, thinking on their feet in response to a question or in cross-examination, lawyers are not going to have the same resources they have at a computer. Our degrees will train graduates who can operate with advanced tools, on their feet, and in a power-cut.

---

## **Challenges of AI to skill acquisition**

AI does not just require us to emphasise the abilities and skills that are constitutive of legal education. It also presents challenges to acquiring both these and wider competences.

### **Learning to learn**

Law schools do not just teach substantive law. We intend our law degrees to impart an ability to acquire knowledge in an ongoing manner, including at the intersection with other fields and disciplines. Learning to learn is an important part of a UCL Laws degree, and we believe we must adapt our degrees to heighten our focus on this crucial life skill in response to AI.

Using AI during study may tempt many students. We do not doubt that it may facilitate the production of some good work if it is consulted responsibly and handled critically (which includes, importantly, that a student's work is always authored by the student). AI may, when used thoughtfully, improve studying and learning through explaining and linking concepts, or through locating information for further study (with the usual caveats around the veracity problems of modern AI systems).

Thoughtful use of AI, however, goes beyond checking for sources or correctness. Thoughtful use is not some horizontal, generic skill that can be acquired in a module on *Using AI at University* and stashed away. Thoughtful use, as the name indicates, requires thinking. Thinking is not domain agnostic; it requires foundational knowledge. Students who learn primarily to act as passive conduits for AI-produced information—regardless of whether they check its veracity—are not going to reach the potential they have to use the tool well, let alone the much broader potential they have as independent thinkers.

<sup>10</sup> Natalie Byrom, [Digital Justice: HMCTS data strategy and delivering access to justice](#) (The Legal Education Foundation 2019).

<sup>11</sup> Hannah Klaus, [Now You Are Speaking My Language: Why Minoritised LLMs Matter](#) (Ada Lovelace Institute 28 November 2024).



Learning at university also goes beyond learning to know—we teach students to interact, and to improve understanding and problem-solving in the process. We hope to give them the foundations to develop these fundamental communicative abilities throughout their careers. Study of the law is by its nature discursive and participatory. In our classrooms, there is growing interest in and adoption of technologies that can transcribe, live-translate, summarise, and even live-prompt during a presentation. Careful and targeted uses of these tools may transform the ways students with disabilities can learn with us, and the possibilities along this frontier are exciting. Other uses might give us reason for pause. How does live translation software change in-the-moment engagement, reflexive and self-critical knowledge development, community building, and the acquisition of the lifelong skill of language through immersion, even where it seems daunting?

More simply, AI tools may undermine basic learning processes. Many studies have indicated the proven educational benefits of handwriting over typing.<sup>12</sup> Might we find parallel issues from students that do not summarise material themselves, but rely too heavily on language model support? AI tools assume that all content and learning processes are digitally mediated and delivered. This is not a robust assumption when we are aiming to educate lawyers to be the thinkers of the future armed with deep knowledge of the law and legal reasoning skills. Furthermore, learning the value of independent and critical thought, knowledge and skill acquisition requires trust. Our academic colleagues are excellent at thinking and learning. It is their job and, more often than not, a large part of their identity. But technology-exacerbated shortcuts to independent learning are tempting for students in a world of distraction, oversimplified bite-sized content, and the many stresses (and the excitement) of studying at university, particularly in London. It can be just as hard to look ahead to a time when you have certain knowledge and skills as it can be for academics teaching you to remember how difficult it was to acquire them. Unlike technology giants, academics do not use algorithmic tools to access students' attention in order to achieve our goals. While we share with technology firms an ambition to attract student attention, our means to do this are engagement, trust and the development of community. Technology firms, by contrast, often push the 'gamification' of education, a highly unproven proposition that risks undermining our core processes of deep and multi-sensory learning, and of building a discipline through community engagement, peer-reviewed standards and mutual accountability.<sup>13</sup> In the varied community moments at a physical university—the classroom, the lecture theatre, our offices, or a snatched second after a public event—we hope to share with students a lesson we believe in: *do not take that shortcut, the long way round is so much richer—and more socially useful—in the end.*

Building trust may be, and may become, all the more challenging because we do not inherit students as a blank slate. Students' previous education, whether at school, in earlier degrees or in the workplace, may not have been great at showing them that what looks like a shortcut, actually is not one. Schools, often in perpetual financial need or pressured to prepare students for an underspecified 'digital future', are particularly vulnerable to

<sup>12</sup> Frederikus Roelof Ruud van der Weel and Audrey Lucia Hendrika van der Meer, '[Handwriting but Not Typewriting Leads to Widespread Brain Connectivity: A High-Density EEG Study with Implications for the Classroom](#)' [2024] 1-9; Giuseppe Marano and others, '[The Neuroscience Behind Writing: Handwriting vs. Typing—Who Wins the Battle?](#)' (2025) 15 Life 345.

<sup>13</sup> Research on gamification in education remains extremely inconclusive on even its core assumptions, such as that gamification is motivating, gamification is engaging, gamification can improve attendance and participation. See generally Amina Khaldi, Rokia Bouzidi and Fahima Nader, '[Gamification of E-Learning in Higher Education: A Systematic Literature Review](#)' (2023) 10 Smart Learning Environments 10; Christo Dichev and Darina Dicheva, '[Gamifying Education: What Is Known, What Is Believed and What Remains Uncertain: A Critical Review](#)' (2017) 14 International Journal of Educational Technology in Higher Education 9.

technology providers appearing with free tools that in practice train students to rely on particular vendors' platforms.<sup>14</sup>

We similarly must be wary with our own provision of technologies in the university. Universities do not have much control over the platforms they use; UK universities in particular have migrated heavily to reliance on Microsoft, and independent technology expertise has been fast altered to expertise in one specific platform's toolset.<sup>15</sup> "The first dose is free" said Google's now chief economist, Hal Varian, of software in 2006. "Once you start using a product, you keep using it".<sup>16</sup> Universities cannot fully escape the political economy of software platforms, and must be pragmatic. Yet, we also need to be mindful, and wary not to become an instrumental piece of a larger business model designed less for education than for ensuring a steady pipeline of future customers. There is a serious risk that if we go passively down this road, private firms, rather than our teachers and scholarly communities, will lead in determining pedagogy and practice.<sup>17</sup>

The balance here is ultimately between using AI as a true study aid in a way that facilitates meaningful learning, and using it to cognitively offload tasks in a way that hinders learning. We believe the balance here is more delicate than much of the current discourse in higher education leads us to believe.

We should not rush to integrate educational AI into modules, learning management systems, reading lists or office software. It is true that some students are going to use AI technologies anyway. That does not mean we need to rush into mitigation mode, shifting our educational focus immediately to 'using AI well'. We should have the confidence to stay rooted in our existing practices for imparting knowledge where these are foundational for learning well (and for using AI well in their future legal lives), and to focus on connecting with students to develop the trust that maybe—just maybe—we know a little bit about learning and knowing, and the value of scholarly community. From those roots and that confidence, which we should focus on inculcating in our students, the rest—including careful uses and integration of AI where it does not inhibit learning to learn—can follow. Our eyes are fixed squarely on our academic mission, not on technologies or the noise around them.

## **Taking intellectual risks**

In our experience, almost all students have moments of fear or apprehension in taking intellectual risks in a higher education environment. In law schools, this very predictably happens early in undergraduate degrees when students are first confronted with writing an analysis and argument in relation to a question for which there is no set answer. Similar moments of trepidation happen at postgraduate level, at a higher level of difficulty. Students might be used to having model answers, extensive scaffolding, and multiple cycles of feedback before finalising work. In higher education, however, we expect students to become confident independent and critical thinkers through the course of their studies. Even with support from academics teaching them, there will be multiple moments where students need to challenge themselves and move well beyond their intellectual comfort zones to develop their cognitive abilities. This may be even more of a challenge for

<sup>14</sup> Ben Williamson and Anna Hogan, *Pandemic Privatisation in Higher Education: Edtech and University Reform* (Education International 2021); Michael Veale, 'Schools must resist big EdTech – but it won't be easy' in Sonia Livingstone and Kruke Pothong (eds), *Education Data Futures: Critical, Regulatory and Practical Reflections* (5Rights 2022).

<sup>15</sup> Tobias Fiebig and others, 'Heads in the Clouds: Measuring the Implications of Universities Migrating to Public Clouds' (2023) 2 Proceedings on Privacy Enhancing Technologies; Agathe Balayn and Seda Gürses, 'Beyond Debiasing: Regulating AI and Its Inequalities' (*European Digital Rights (EDRI)*, 2021) 110.

<sup>16</sup> Charles Piller, 'How Piracy Opens Doors for Windows' (*Los Angeles Times*, 9 April 2006).

<sup>17</sup> Elana Zeide, 'Robot Teaching, Pedagogy, and Policy' in Markus D Dubber, Frank Pasquale and Sunit Das (eds) *The Oxford Handbook of Ethics of AI* (Oxford University Press 2020).

students who have first studied in other jurisdictions. Based on our extensive experience as law teachers, we have every confidence that our very able students, from all educational backgrounds, can make these intellectual leaps and flourish as lawyers. Unfortunately, AI tools can offer tempting alternatives for students when confronted with discomfiting intellectual challenge, even more tempting when the apparent short-term risks—particularly of employability—seem too high to take a leap into the intellectual unknown.

Having read material from the academics teaching them, or at the least heard their angles and interpretations during lectures and tutorials, we know that very able students can be tempted simply to study and replicate the former's stances. If done competently, students will not do too badly in assessments, but meaningful academic success typically requires independent thinking and taking intellectual risks. Most simply, this requires a student to make, independently, the analytical connections on which the material being taught is based, in order to foster deep understanding of material. Beyond this, independent thinking might involve reading further; piecing together the materials in an unguided and seemingly unconventional way; or trying to identify unmade connections, draw them out and express them with one's own voice. This is all going beyond what one is explicitly taught, into the unknown. What can seem like the biggest risk of all is to challenge the lecturer's work itself with an ambitious and thoughtful critique. This apparent risk often pays off with a smile on the lecturer's face as they award a high mark and see their job of stimulating thinking is done.

It is not the case that AI cannot support the taking of intellectual risks when used as a careful tool. However, it cannot do the thinking for students. It cannot play around with legal ideas and analysis in a student's own mind to the point that independent legal thinking is acquired as a skill. Furthermore, by default, AI steers students down a safe path, producing more of the same. AI generates statistically median content, median structure, median style and median substance. Even when explicitly prompted to be *wild and out there*, it will slip immediately into the median of what *wild and out there* represents in its learned patterns.

Will median directions be interesting or novel ones? Will reading recommendations point people in unusual and untrodden directions? AI researchers have long considered how to build in concepts like serendipity and surprise, but fall into the trap that optimising for these and building them in undermines them conceptually forecloses the possibility of serendipity, surprise, or ultimately really much challenge at all.<sup>18</sup> It is possible that, without careful use, generative AI might become equivalent to stabilisers on a bike. Training wheels may offer comfort and safety, but there is joy to be found riding on two wheels alone. The process of building knowledge and understanding is one of uncertainty and risk, challenge and reward. This is even more so with interdisciplinary study, or study at the frontier of a research field.

When thinking about these 'risky' but foundational aspects of the learning process, AI tools are unhelpfully advertised as being there in a moment of need. This advertised benefit plays well into the history of business models in the digital economy trying to engineer reliance. Yet moments of intellectual need at university are moments to break through barriers and to progress. When we talk to our alumni around the world, what they remember most fondly and vividly are often these moments of need—of acute intellectual challenge when they were confronted with real difficulty in taking the next steps of legal reasoning, argumentation or understanding—and their surprise and satisfaction in those moments at

<sup>18</sup> Sylvie Delacroix and Michael Veale, 'Smart Technologies and Our Sense of Self: Going Beyond Epistemic Counter-Profiling' in Mireille Hildebrandt and Kieron O'Hara (eds), *Life and the Law in the Era of Data-Driven Agency* (Edward Elgar 2020).



just how far the skills they had acquired (and been challenged to acquire) carried them through.

Similarly, we hear from students how transformative it was to go beyond their comfort zones in clinical education, such as in our legal advice clinic, or by participating in the many mooted competitions we support. Would these experiences be the same with a statistical language model whispering in a student's ear? It is not clear they would be. They may look the same from the outside. They may even look slicker. But would these experiences have the same effect on people's lives? Education (and acquiring a good legal mind as a result) is not about trying to achieve plausible text, or a veneer of passable professionalism. It is the process and the experiences that matter most, developing the strongest capacity for critical and independent legal thinking. The unwarranted use of AI bypasses process and experience. Risk-taking can be tricky to preserve. Law firms appear earlier and earlier in students' degrees, checking grades for sought-after placements, which can leave less and less time for students to feel free to take intellectual risks. Students from less advantaged backgrounds, without much of a material safety net, may inherently feel less able to take risks. None of these observations concern only AI, but they will interact with a technology that inherently sells the path most trodden. We are committed to retaining the space and opportunity for intellectual risk-taking as an essential part of an excellent legal education.

---

## Assessment

Many of the most publicised concerns around AI in education relate to assessment. At UCL Laws, we have a policy that between 50-100% of every taught module, both undergraduate and postgraduate, is assessed securely, that is, in a way that can guarantee that AI does not substitute for the skills or knowledge acquisition being evaluated. Secure assessments include in person examinations (which also have advantages of fairness and reducing the mental stress of alternative assessment formats over more extended, overlapping time periods) and other in person assessments, including oral assessments. In addition, for the remaining non-secure types of assessment, such as coursework, the use of AI that can create or alter content is prohibited by default unless explicitly authorised by the module convenor for a valid pedagogical reason. In this section, we would like to explain and justify how and why we arrived at this approach.

### Integrity must be the priority

Integrity of our students' educational achievements is as important to us as the integrity of the research and knowledge produced by our community of scholars.

UCL Laws degrees qualify individuals to continue training both as a barrister in England and Wales and as a lawyer in many other common law jurisdictions. It functions as a trusted and authenticated marker of achievement. These jurisdictions trust our degrees as having imparted foundational legal knowledge. Maintaining the integrity of this marker is crucial, both inherently for the value of our education provision and additionally for the role we play in the legal sector. Unlike some other regulated workplaces, legal services regulators also typically require disclosure, truthfully from the candidate or through us via the reference process, of findings of plagiarism, academic misconduct or cheating in educational assessments. This is in order to assess the character and suitability requirements of a legal professional practice.<sup>19</sup> Under UCL's regulations, unauthorised AI usage is counted as plagiarism, and would require disclosure of this type.<sup>20</sup>

<sup>19</sup> See eg Solicitors Regulation Authority, '[What is character and suitability?](#)' (SRA 19 December 2024).

<sup>20</sup> UCL, [Academic Manual \(2024-25\)](#), ch 6 s 9.2(5)(a).

## Diversity and security

Assessment at UCL Laws is designed to assess acquired abilities, skills and knowledge. The diversity of competences we seek to assess—those described earlier in this document are a sample—necessitates a diversity of assessment methods. We also monitor how our equality, diversity and inclusion goals interact with assessment methods.

Our approach as a faculty is to facilitate assessment diversity where appropriate whilst not compromising *complete* assurance of academic integrity throughout our programme. We recognised that diverse assessments provide important learning experiences and allow our diverse community of students to showcase their skills and understanding. However, certain assessment types have always been structurally more vulnerable to risks to integrity, such as contract cheating and essay mill services. The practical impossibility of perfect detection and policing, combined with the societal importance in assessing skills such as research, is why so many jurisdictions have stepped upstream to make running or advertising such services a criminal offence.<sup>21</sup> We still see real value in such assessments, but must recognise that AI transforms a serious step of a student enlisting an illegal third party into a few easy clicks, even integrated into software we provide them. We have felt able to robustly socialise students into seeing contract cheating and essay mill usage as unacceptable. However, socially deterring misconduct through AI, in the context of political and economic discourse encouraging its usage, is much more challenging. We cannot, even as we continue to enforce academic integrity in all assessment, assume that tasks such as coursework alone robustly safeguard the integrity of our degrees as they once did.

As types of assessment that we would like to include throughout our programme for reasons of pedagogically-informed diversity (such as coursework) are not able to be fully secured against academic integrity issues arising from AI, they need to be avoided or otherwise complemented throughout the programme of study by *an equal or higher proportion of secured assessment* in order to provide assurance of the security of the degree qualification. This does not mean that AI usage or reliance in insecure assessment is permitted; using AI in any form of coursework (unless permitted for pedagogical reasons) will be a disciplinary violation. We accept a lower level of security in these specified, counter-balanced, parts of the qualification precisely because of the benefits of diversity and the reliance we can place on the majority of the assessment base. This is a carefully designed situation for assessing our degrees, where integrity takes priority and diverse assessment methods have their place in educating good lawyers. We are keeping this position under review to ensure that the integrity of each student's assessment performance can be guaranteed.

## Designing security into assessment questions fails us

Contrary to a significant amount of commentary in the area, we cannot achieve the necessary level of assured integrity of assessment through maintaining insecure forms of assessment combined with either AI-resistant design principles or *ex post* checking. Some of our observations come from generic analysis of the current situation in the AI business, while others are law-specific.

Some modules can make successful design interventions that add to AI resistance—highly bespoke assessments; niche, contemporary content; empirical research projects—but this largely depends on the type of module, the skills it is trying to impart, and the possibility to assess in that way. Some tasks are simply conceptually harder for AI to currently complete or assist with. Module size and teaching format matters too. Many of our modules are taught by teams, where each might cover a few weeks on highly specialised topics—not

<sup>21</sup> See eg Education and Training Act 2020 (New Zealand) s 393(1)(a); Qualifications and Quality Assurance (Education and Training) Act 2012 (Ireland) s 43A; Tertiary Education Quality and Standards Agency Act 2011 (Australia) s 5; Skills and Post-16 Education Act 2022 (England and Wales) s 26(2).

enough time to become so familiar with individual students so that an unusual assessment can be reliably spotted.

Foundational topics are, in some ways, more vulnerable to the use of AI to achieve passing grades. Law degrees, by intent and by regulation, impart foundational legal knowledge. There are few topics taught at universities around the world with more structured text existing online about them than common law systems. Language models are soaked in core legal knowledge because law is an open, text-driven field. Judgments, statutes, commentaries and more have been in highly structured legal databases for decades. Such texts are so perfect for computational modelling that the European Union even bundles up its legislative texts just for natural language process training in its [Acquis Communautaire](#). Companies like Westlaw and Nexis are fast training legal AI systems for sale that will be available to students either through university systems or through the workplaces of the practitioner parents of students. AI systems can answer basic questions about the law well, and they can analyse problem questions at a basic level of competence. They might not do so *excellently*, but this is not the core issue. Where this is the difference between a pass and fail grade, it is important we do not let this undermine the integrity of our programmes. Awarding a law degree to someone without the appropriate legal knowledge and skills both threatens our values and creates a real danger to people who may rely on this individual for life-changing advice in the future.

Common tropes, such as having students critique an essay they create using an AI prompt, do not seem to work educationally either. Experience with this practice appears to indicate that students who go into this without acquiring foundational knowledge only manage cosmetic critiques.<sup>22</sup> Slotting this kind of exercise once during a programme might be fair, but within multiple modules, it will clearly get rote, tiring and stale. Consider that students may, before long, have had similar tasks at school or previous degrees, and the exercise looks like a pretty painful—and pedagogically ineffectual—one to inflict. The unfortunate truth is that there is no easy way to address the multiple directions of challenge that AI presents to examination and learning with such simplistic assessment fixes.

Furthermore, applying the constraint that non-secured assessment must be AI-resistant by design consumes the freedom staff need to design pedagogically good tasks that assess skills in the right way *for that skill*. Staff would be required to understand the current (and continually updating) capabilities of AI systems which, frankly, is not their job. These change regularly—for example, colleagues until recently were looking for fake citations to identify AI-generated work, but more recent AI systems ground their findings in the Web.<sup>23</sup> Systems can already ingest significant information into either their prompt windows or a connected database—for example syllabi, lecture transcripts or subtitles, briefing tasks, previous drafts of work, textbooks and readings, all of which are conveniently very machine-readable as accessibility law requires them (rightly!) to be for screen readers.<sup>24</sup> In short, it is not enough to design questions to connect to class topics or approaches—these systems can take that context into account in their answers.

### **Policing does not ensure security**

Ex post detection as a method to *ensure* the security of an assessment does not work. AI detection tools are unreliable and do not function as meaningful safeguards. They exhibit

<sup>22</sup> Sarah O'Connor, '[Students Must Learn to Be More than Mindless "Machine-Minders"](#)' (*Financial Times*, 4 March 2025).

<sup>23</sup> Reiichiro Nakano and others, '[WebGPT: Browser-Assisted Question-Answering with Human Feedback](#)' (arXiv 17 December 2021).

<sup>24</sup> The Public Sector Bodies (Websites and Mobile Applications) (No. 2) Accessibility Regulations 2018; This is typically either done using large context windows or retrieval-augmented generation (RAGs), see Patrick Lewis and others, '[Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks](#)', *Advances in Neural Information Processing Systems* (Curran Associates, Inc 2020). All market leaders in AI at the time of writing offer both.

bias. They are trivially circumvented in many ways. Their false positive rates of detection are too high given the regulatory consequences of findings of academic malpractice for potential lawyers. Students can cognitively offload the demonstration of skills onto AI systems even where they do not copy exactly, for example delegating the generation of structure, points and counterpoints. This ‘structural’ support offered by AI is a point of particular concern to us as legal educators, as our students need to develop these analytical skills independently—we have identified this as a more invisible point of AI use where explicit educational policy is required.

Students are constantly being illegally marketed AI detection circumvention tools on social media, in an arms race that colleagues do not have the time or energy to keep up with.<sup>25</sup> Even if AI systems did choose to add watermarks, companies like OpenAI refuse to do so anyway as it would put off their userbase.<sup>26</sup> Openly available language models that can run on custom devices or cloud services without watermarking features are already widely available, licensed permissively and irrevocably, and can and will be rapidly deployed by the large contract cheating/essay mill industry.<sup>27</sup> While companies in the EU are facing an obligation to watermark language model output, no such requirement exists in English law.<sup>28</sup> Students can already download powerful models to their devices and use them locally with zero surveillance or potential watermarking using tools like *Ollama* or *GPT4All*. Moreover, the process of detection throws suspicion onto students who have done nothing wrong. Staff become anxious, looking for signs of AI when it may well not be there at all, increasing workload during marking periods when they are already stretched and deadlines are tight. Organising investigatory vivas or panels is logistically extremely difficult in the short time frames and at the scales we face, and within staff workload constraints, although we are trialling these approaches. Precariously balancing the academic integrity of a degree on staff reliably suspecting and detecting AI-facilitated misconduct may hold back the staff–student trust building we have argued above is crucial for learning to learn, and learning to take intellectual risks.

This does not mean integrity in non-secured assessment does not matter, and that we do not and will not seek to enforce violations of academic integrity when they occur. It means that the guaranteed integrity of our overall assessment matters even more, and that we apply proportionate, regularly reviewed measures to any (minority) non-secured assessment that we retain for pedagogical reasons. We do not expect these measures to do the heavy lifting of providing assurances of integrity to our entire degree programme, and that potential policy violations (which we will continue to ward off culturally and pedagogically) are at a level with which, because of our programme and module level assurances and guarantees and the benefits of assessment diversity, we are comfortable.

### **Students should not be assessed as content creators**

We are actively monitoring whether other approaches, such as requiring students to write in software that maintains a document history, might prove useful in the assessment process. Such software still allows significant cognitive offloading—students relying on AI to think and produce points and structure, even if they type manually.

<sup>25</sup> Noëlle Gaumann and Michael Veale, ‘[AI Providers as Criminal Essay Mills? Large Language Models Meet Contract Cheating Law](#)’ (2024) 33 Information & Communications Technology Law 276.

<sup>26</sup> Deepa Seetharaman and Matt Barnum, ‘[There’s a Tool to Catch Students Cheating With ChatGPT. OpenAI Hasn’t Released It](#)’ (*Wall Street Journal*, 4 August 2024).

<sup>27</sup> Gaumann and Veale (n 25).

<sup>28</sup> [Regulation \(EU\) 2024/1689](#) of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence and amending Regulations (EC) No 300/2008, (EU) No 167/2013, (EU) No 168/2013, (EU) 2018/858, (EU) 2018/1139 and (EU) 2019/2144 and Directives 2014/90/EU, (EU) 2016/797 and (EU) 2020/1828 (Artificial Intelligence Act) OJ L, art 50(2).

There are voices arguing that heavy use of AI in, for example, coursework, is acceptable—these tools will be available in the workplace, so why should students not have access to them in a university? Perhaps they can be used in assessment, as long as they are acknowledged? We disagree. Students at UCL Laws are assessed on underlying abilities and skills rather than a capacity for ‘content creation’. This is part of what distinguishes tasks in higher education from tasks in workplaces. Materials produced in a workplace may look similar to those produced in a law school (e.g. a brief, a presentation, or a piece of legal analysis) but, in an educational context, the skills demonstrated in the process are more crucial than the output. If we are just training students to be “content creators”, functionally-sound language model-generated text is an appropriate response to an assessment.<sup>29</sup> Lawyers do indeed need to be able to produce content — sometimes vast amounts, at very short notice. But this content creation is just one skill amongst many. It is a grave mistake to think that content creation proxies for the many foundational understanding and skills that legal education should be trying to impart.

Some abilities and skills can be assessed regularly and in concert with others. For example, nearly all assessments in a law school will assess, to some degree, style, knowledge, structuring of ideas, and critical thinking. Sometimes, we might want to assess the ability to create convincing and professional-looking content. However, assessing the ability to create content with all the tools at hand that one might expect in a workplace (including Generative AI) today undermines the assessment of many of these other skills in a way that previous assistive tools for rapid text production (like spelling and grammar checkers or translators) did not. This means that we need, to some degree, to isolate the assessment of this skill from other skills in order to ensure that we are getting a true and holistic picture of the students we are educating.

---

## Concluding thoughts

AI does not change everything about legal education. While UCL Laws will, as it always has done in the face of new technologies, update and adapt, we will also stay true to our values and use every opportunity to strengthen our commitments to them. AI is an opportunity, but it is not an opportunity to abandon foundational legal education because Microsoft has decided to embed a chatbot into their ubiquitous office software suite. It is an opportunity to heighten the imparting of abilities and skills that will change students’ lives and empower them to be valuable citizens and effective lawyers (or professionals in other contexts), and to refocus on those competences that will remain transformative, given the future we might face in and around the law. As a law school, we have shifted significantly to secured forms of assessment, but in many respects, this is also a return to recent practice following the huge shift to coursework necessitated by the coronavirus pandemic. The commitment to assessment integrity is a long-standing commitment to ensuring that our students graduate with independent, and verified, abilities to think well as a lawyer.

## Next steps for the sector

We can alter our own educational policies at UCL Laws, but these issues affect the entire legal education sector (and the legal and education sectors more broadly). In our view, the sector should make explicit its own thinking, ideally collectively. In addition, each institution should consider how AI connects to their educational mission and values, and be bold in relation to the policies they adopt. We have laid out our current thinking here, as well as the policies we have decided upon as a faculty, and recognise not all institutions may share it, but we wish to facilitate the conversation. Thinking and policymaking also takes time, energy and effort, and overloaded colleagues can be excused for not wanting another issue

<sup>29</sup> For a related argument, see Sonja Drimmer and Christopher J Nygren, ‘[How We Are Not Using AI in the Classroom](#)’ (2025) 1 *ICMA News* 25.



on their plate. We hope this discussion paper can bootstrap and support such work, providing a starting point.

We will be in touch with colleagues across the sector to work on facilitating and convening discussion on AI and legal education. Please reach out to us if you would like to be involved in this.

All considered, it is clear to us that conscious decisions must sit at the heart of universities' approaches to AI and education. Universities must not be passive rule-takers. We must not simply 'adjust to' speculative educational and professional visions of the future marketed by technology firms in order to sell more cloud computing and increase reliance on tools with questionable and uncertain utility. Universities must steer, and if necessary, themselves create, the technology they need for their missions. Working together makes this viable. We educate the next generation, and provide them with a framework for thinking about the world that is just as, if not more, important than the technological infrastructures they will find themselves in. A response to AI in education is a response to the future.

These are the first steps in our thinking; we do not plan to stop thinking any time soon.