

Testing the Frontier: Generative AI in Legal Education and Beyond

Cari Hyde-Vaamonde and Anat Keller

King's College London, United Kingdom

Abstract

Unlike previous AI applications in law, which focused on search and prediction, generative AI (GenAI) has the capacity to produce, on some level, coherent legal writing. It is this capacity that is driving legal educators to fundamentally reconsider approaches to academic integrity and pedagogical practice. This study investigated how legal education can productively integrate GenAI into higher education settings while maintaining academic integrity, specifically examining: (1) how students critically evaluate AI-generated legal content; (2) what limitations they identify; and (3) how collaborative approaches can develop effective guidelines for responsible GenAI use in legal curricula. We employed a novel three-stage intervention, using metacognitive modelling and collaborative co-creation, involving over 125 law students from King's College London. Data were collected through workshop observations, student evaluations of AI outputs, collaborative guideline development and follow-up interviews. Students consistently demonstrated sophisticated critical evaluation of AI-generated legal content, identifying significant limitations including superficial analysis, a lack of argumentative coherence, citation inadequacies and absence of nuanced understanding. Most notably, students strongly preferred content that demonstrated originality and critical thinking – precisely where AI systems under-performed. Exposure to AI limitations fostered responsible usage attitudes and enhanced students' confidence in their own analytical capabilities. Our findings demonstrate that critical engagement with AI tools enhances rather than diminishes academic standards. The co-created guidelines offer a transferable model centred on fostering a 'culture of trust' rather than prohibition. This transferable approach prepares future legal professionals for an AI-augmented workplace while preserving core values of legal education: critical thinking, ethical reasoning and intellectual rigour.

Keywords: Generative AI; legal education; academic integrity; critical thinking; technology in education.

1. Introduction

The project 'Testing the Frontier: Generative AI in Legal Education and Beyond' aimed to enhance students' critical thinking and AI literacy within legal education, with broader implications for scholarship and research. The primary focus was to ensure that students critically evaluated AI-generated content, recognising its limitations, inaccuracies and biases. By doing so, students could learn to apply generative AI (GenAI) responsibly and effectively in their academic work, thereby maintaining academic integrity and rigour.

The project emerged from our recognition that legal education needed a proactive rather than reactive approach to the advent of ubiquitous GenAI tools.¹ As authors teaching in fintech and AI in law – fields where AI has been a feature for some time and where AI adoption has been accelerating – we were well positioned to anticipate the technology's potential impact on legal

¹ GenAI tools as described are typified by ChatGPT (Chat Generative Pre-trained Transformer), released by OpenAI. This is a tool with an interface that allows a user without any technical training to interact using (mostly text-based) messages that are responded to with human-like responses. The earliest version in this format was released to the public in November 2022.



Except where otherwise noted, content in this journal is licensed under a [Creative Commons Attribution 4.0 International Licence](https://creativecommons.org/licenses/by/4.0/). As an open access journal, articles are free to use with proper attribution. ISSN: 2652-4074 (Online)

practice and education. In preparing the project, we felt it was essential that, rather than waiting for problems to emerge or pursuing policies based on assumptions about how these tools might be employed, we should seek to understand how students would actually engage with them and explore whether collaborative approaches could yield more effective solutions. This proactive stance was supported by King's College London's Teaching Fund (CTF) Awards, with this being one of the 19 projects selected by two members of staff from the College and two students, with oversight from the CTF fund chair.

The objectives of the project included:

- *Enhancing AI literacy*: equipping students to view GenAI outputs through a critical lens, promoting a deeper understanding of the potential limitations of these outputs.
- *Promoting academic integrity*: fostering a culture of trust rather than policing students, and incentivising them to use GenAI only where there is added value and they can do so in a manner that does not jeopardise academic integrity and achievement.
- *Preparation for future workplaces*: ensuring that students are ready for AI-integrated workplaces by providing them with practical experience in scrutinising GenAI outputs and using them responsibly in their writing process to maintain academic standards.

The study builds upon and contributes to the emerging scholarly literature on the potential of AI in higher education settings.² Studies indicate that AI has benefits in enhancing students' learning experiences and supporting the development of critical thinking in various fields and disciplines. These studies often also highlight the potential detriments of such use to learning, including the risk of over-dependence on AI, which could potentially curtail learners' capacities for critical thinking and intricate problem-solving, and diminish human interaction and engagement.³

The findings of this study can serve as a valuable resource for teaching staff at higher education institutions in adapting teaching methods to foster critical thinking and support creativity and problem-solving in law in an AI-driven environment. This study aims to assist educators in bridging the gap between the benefits and limitations of using AI in education by introducing a reflective method that leverages the limitations of GenAI to augment the benefits it can bring to students and their learning experience. Further, when one notes the different cultural dimensions to how GenAI use is perceived, as highlighted by Yusuf and others, engagement with the diverse student body becomes crucial, making the collaborative approach central to formulating effective policies and strategies in relation to these tools.⁴

Karim Lakhani of Harvard Business School remarked that "AI will not replace humans but humans with AI will replace humans without AI"; this statement, while deliberately provocative, indicates the need to face the issue of GenAI usage front on.⁵ What we found in this project was nuanced. Rather than the division being between those who simply use GenAI and those who do not, the crucial element of critical thought is likely to be a distinguishing factor in legal education and beyond.

2. Theoretical Framework

The 3P model proposed by Biggs highlights that three key elements can influence learning outcomes: student-dependent factors, teaching-dependent factors and interactive impacts from the whole system.⁶ The framework is helpful in the context of GenAI tools in education, as these tools have the capacity to impact all three elements significantly. Understanding how a system responds to change, how interconnected elements impact each other, requires a systems thinking approach, such as that set out by Meadows.⁷ Meadows' analytical framework goes beyond the simple interaction of rules and describes how, in complex systems, there are particular core paradigms that constitute leverage points that can transform the entire system. In the context of legal education, critical thinking can be seen as such a paradigm. Rather than being one part among many, it is foundational to not only legal education but the legal profession itself. If a technology or a practice threatens the integrity of this paradigm, the system itself may be critically undermined.

² Ipek, "Educational Applications."

³ Gonsalves, "Generative AI's Impact," 1.

⁴ Yusuf, "Generative AI," 26.

⁵ Lakhani, "AI Won't Replace Humans."

⁶ Chan "Students' Voices," 14 exploring Biggs "Teaching for Quality Learning."

⁷ Meadows, Thinking in Systems.

The paradigmatic importance of critical thinking, and its vulnerability to GenAI has been illustrated in the literature, although perhaps has not been stated in these terms before. Naqvi and colleagues⁸ emphasise the ‘risk of cognitive complacency’ in health sciences education, arguing that embedding critical thinking regarding GenAI outputs ensures that ‘researchers and clinicians remain architects of their intellectual landscapes’ rather than becoming dependent on technological outputs. At the same time, Gonsalves – who identified how students ‘moved fluidly between cognitive stages’ – highlights how this challenges traditional approaches such as the hierarchical structure of Bloom’s Taxonomy when using GenAI.⁹ The challenge is further complicated when one considers the cultural variations in GenAI perceptions¹⁰ and the risk that negative student attitudes may lead to surface learning approaches.¹¹ Lee and Low also emphasise the benefit of involving GenAI to challenge students’ ability to ‘think critically and enhance their human interactions’.¹²

This is a turning point, where protecting critical thinking as the foundational paradigm of legal education requires intentional pedagogical interventions that engage students as active participants in evaluating AI-generated content.

3. Research Objectives

AI has the potential to enhance students’ learning experience and support creativity and problem-solving in law. However, there is a gap in AI literacy and the essential skills needed to harness this potential effectively and ethically while maintaining academic rigour and integrity. There is also an increasing risk of inconsistency and uncertainty among academic staff regarding how to advise students accordingly.¹³

Our project aimed to encourage students to develop critical thinking when using GenAI tools and view AI outputs through a critical lens. This reflective process allowed students to acknowledge the limitations of GenAI tools, identify possible inaccuracies and biases, and find ways to respond to these challenges and refine GenAI outputs.

Given that the Russell Group represents 24 universities that have collectively agreed on principles for GenAI integration, their framework offered relevant considerations for maintaining academic standards while navigating technological change. Our methodology involved student voices and collaborative development, ensuring the approach was grounded in student experience rather than institutional hierarchy. Moreover, the principles’ focus on academic integrity transcends the type of institution.

The impact of this process is twofold. First, while in desperation some students may act contrary to best practices, understanding and engaging with the limitations will disincentivise the dishonest uses of GenAI in learning. It will foster a culture of trust rather than detection or policing of students and, more broadly, promote the adoption of Russell Group principles in the use of AI in education. Second, it is a first step in preparing students for their future workplaces, which will undoubtedly integrate AI tools. In addition, the project will equip staff to support students in using AI effectively and ethically, to become leaders in an increasingly AI-enabled world in accordance with the Russell Group principles for AI in education: ‘Universities will ensure academic rigour and integrity is upheld.’¹⁴

In ‘Navigating the Jagged Technological Frontier’, the authors described the point at which GenAI technologies, such as ChatGPT, assisted in office-based tasks as a non-uniform boundary, making it difficult to reliably predict in different contexts.¹⁵ Given the rapid development of technology at present, this boundary will also flex and change over time. In this project, we encouraged students to probe and test this frontier in an academic context, equipping them with the tools to do so. It is an effort in keeping with Prakash and Nair’s paper, which asserts that ‘Law school policymakers must establish adaptable, transparent policies for AI use that are regularly reviewed to keep pace with rapid advancements’.¹⁶ The project described here engaged in the actual experience of students, and involved staff in the process.

⁸ Naqvi, “Critical Thinking,” 2, 4.

⁹ Gonsalves, “Generative AI’s Impact,” 2, 11.

¹⁰ Yusef, “Generative AI,” 25.

¹¹ Chan, “Students’ Voices,” 4.

¹² Lee, “Using GenAI in Education.”

¹³ For example, Lee “Using GenAI in Education”; Coffey, “Students Outrunning Faculty.”

¹⁴ Russell Group, “New Principles,” 3. King’s College London contributed to and subscribes to the Russell Group’s five principles on the use of GenAI tools in education.

¹⁵ Dell’Acqua, “Navigating the Jagged Technological Frontier.”

¹⁶ Prakash, “Integrating Generative AI,” 67

Students start with varying levels of ability. The introduction of GenAI into the legal landscape can cause students to doubt their abilities, so it is important to build confidence, and the best way to achieve this is not necessarily didactic training.¹⁷

When encountering something new, if learners are provided with preapproved forms of use, this can limit spontaneous exploration and discovery, termed the ‘double-edged sword of pedagogy’.¹⁸ When analysing texts and discovering the impact of GenAI on critical thinking, formalised modes of teaching that didactically present approved modes of GenAI use actually jeopardise the overall goal of the intervention itself: to improve critical thought.

This study addresses the gap in legal education practices by providing empirical evidence of how law students evaluate AI-generated legal content and what this reveals about maintaining academic integrity, while also preparing students for AI-integrated legal practice. Unlike previous work focused on institutional policies, we examine the pedagogical process itself: how students learn to critically evaluate AI outputs when given structured opportunities to do so.

4. Methodology

We combined workshops with quantitative surveys and semi-structured group interviews and received ethical approval¹⁹ for students and staff participation in a low-risk setting. Students taking part in the workshops signed consent forms via the online KEATS platform.

The workshop methodology employed elements of metacognitive modelling, a pedagogical approach where teachers make their analytical processes visible by talking through their reasoning and demonstrating their thought processes as they work through problems.²⁰ This technique, which involves active demonstration of analysis rather than simple explanation, was adapted here from its previous application in legal text analysis to guide students through the critical evaluation of AI-generated content. By demonstrating the evaluation process and then observing students’ analytical approaches, we were able to both model critical thinking and gather insights into how students naturally approach GenAI outputs.

The modelling approach was particularly prominent in the second-stage workshop, where students worked in groups to analyse GenAI-created essay content based on the marking rubric of postgraduate studies at King’s College London. Teaching staff present offered their assessments along with those of the students in a non-didactic way.²¹

4.1 Stage 1: Interactive Workshop

Facilitated by staff from the Law School and external consultants from AND Digital, over 100 students from the financial and fintech law pathways attended the workshop on campus and an additional 25 students and staff attended online.²²

The workshop was recorded and became available on the dedicated KEATS page that was created to disseminate the project. Over 200 students enrolled on the KEATS page and had access to the resources.

The aim of the workshop was to enable staff and students to build foundational knowledge and AI literacy regarding the effective, appropriate and ethical use of AI within an education setting before the second stage commenced. Adam Hall from AND Digital discussed the importance of tech skills and digital upscaling, the development of GenAI and the way it works (‘you can think of it as an autocomplete on steroids’). This included some practical advice – for instance, rather than asking it to summarise the key points, ask where you would find a certain argument in an academic paper.

According to a survey, which was conducted during the session, students’ prior experience with GenAI varied from ‘not having even opened ChatGPT before’ to frequent use. Most students had used GenAI in the past, and many confirmed using it in some way for their coursework.

During the last 10 minutes, students worked in groups and considered how they would use a model such as ChatGPT as if it were their legal assistant and what questions they would ask.

¹⁷ Bandura, “Self-Efficacy,” 197.

¹⁸ Bonawitz, “The Double-Edged Sword.”

¹⁹ King’s College London Ethical Clearance Reference Number: MRSU-23/24-40780.

²⁰ Hyde-Vaamonde, “Teaching in a Transitional Space.”

²¹ Ahn, “Do as I Do.”

²² These pathways were chosen due to the accessibility of the authors to students as well as their knowledge of the subject matter that was required for them to critically assess GenAI outputs.

The demonstrations used AND Digital's proprietary model, which performed similarly to ChatGPT, although when students tested prompts independently, they used various GenAI models of their own choice. The study adopted a model-agnostic approach as the focus was on discovery and reflection rather than didactic instruction.

4.2 Stage 2: One-day Focus Group Workshop

Twelve students participated in an exercise in which they were provided with an essay-type question and an AI-generated output related to financial law. Over a period of four hours, students were then asked to critically analyse the text, arguments and sources in order to identify inaccuracies, biases or other limitations in the GenAI output text and find ways to respond to these limitations and challenges, refine GenAI outputs and reflect on the process they had carried out.

Stage 2 focused on the students scrutinising the quality of the output so they could 'use GenAI tools effectively and appropriately in their learning experience', ensuring that 'academic rigour and integrity are upheld'.²³

4.3 Stage 3: Interactive Staff–Student Workshop

Four students who participated in the Stage 2 workshop presented their learning experience, including the challenges they faced in the process of scrutinising GenAI output and their recommendations for rectifying the identified limitations in these outputs. The presentations were followed by an open discussion and brainstorming of students and members of staff to draft guidelines for the effective and ethical use of GenAI in academic writing tailored to the law discipline. This co-creative experience brought to the surface a more concrete meaning and real-world implications of the Russell Group AI Principles in a law higher education setting.

4.4 Final Group Interview Stage

Three students who were involved in all stages of the project met online and discussed their experience and the lessons they had learnt from the project using a series of open-ended questions. The open-ended questions were drafted by the authors, but one of the students took the role of the interviewer and asked the questions, responding and adapting them accordingly.

5. Findings

5.1 Students Naturally Identify AI Limitations Across Multiple Dimensions

Through the workshop process, students frequently pointed out concerns with the GenAI output, often without prompting. Critiques included that the output was too 'generic' or 'superficial', not showing in-depth knowledge.

Students identified concerns with the development of argument, observing that the content resembled a list in which there was no hierarchy to the ideas or connections between them: 'argument doesn't flow ... they have proposition a, b, c, but no threading of sentences together'. They also noted where it gave equal weight to different arguments, lacking in nuance. They noted that main issues are addressed, but that there was no detailed discussion, saying 'it just stated facts and that's it'.

They identified the output as 'not human-like' and lacking the patterns of human writing, expanding and focusing on themes rather than going through them 'in a mechanical manner'. Although they anthropomorphised the author of the essay (as discussed later), they felt the content did not exhibit personality.

As might have been expected, the students picked up on an absence of citations for particular assertions, as well as the use of non-scholarly citations. Beyond these easily verifiable errors, more subtle errors were identified by the students. For example, they noted that the arguments 'mix objectives vs rationales' and the subject matter of the essay 'mixes regulation vs oversight and policymaking'.

Noting the structure of the essay itself, they observed that the piece used 'outdated illustrations and case studies', but on other occasions gave no practical examples. They also had concerns about the content: 'The introduction doesn't introduce sections of the essay or their stance.'

Despite recognising these significant limitations, students did not ignore certain strengths in the AI output. They appreciated that it was clearly structured, with coherent English expression ('great-sounding English'), and noted that it successfully

²³ Russell Group, "New Principles," 3.

identified many key terms and concepts relevant to the topic. They also appreciated the ability of AI to simplify complex legal concepts, which had the potential to make them more accessible. However, this recognition of AI's capabilities did not diminish their critical assessment of its fundamental shortcomings.

Indeed, the above concerns resulted in observations from the students themselves that illustrated their caution about the use of GenAI: 'It is important to acknowledge that GenAI cannot substitute critical thinking but possesses the ability to summarise a variety of information provided to it.'

5.2 Exposure to AI Limitations Builds Confidence and Responsible Usage

The exercise showed students how important it is to be 'suspicious' and inquisitive, and to have confidence in their knowledge of the subject matter and their ability to produce creative and original ideas. They provided suggestions on how to address the challenges presented by this new technology, commenting: 'It is better not to prevent students from using GenAI – it's there, so it is better to educate rather than forbid.'

They showed a nuanced approach to its use: 'to use GenAI successfully, you must know the subject matter, and do the readings. GenAI is trained to be convincing and deceptive.' And 'It is useful with careful usage.'

Recognising the limits of the technology, they stated that, 'if given a broad spectrum essay, GenAI struggles with it', and 'GenAI is an add-on so you need to be vigilant when using it beyond brainstorming', but it is 'better to come up with your own ideas first'. When asked what advice they would give to fellow students, they encouraged their peers to 'trust in yourself, at the moment GenAI cannot replace or mimic human critical thinking'.

It was interesting to see how students would anthropomorphise the output of AI, making comments such as, 'I really like the way this person structured their essay', 'I don't want to fail him [laughter]' and 'At least he tried [laughter]'. This anthropomorphising demonstrated students' application of peer assessment skills to AI outputs. It revealed the pedagogical value of treating AI-generated content as material for collaborative critique.²⁴

At first instance, it seemed that the essay's mark would fall within the bracket for a merit, but once critically analysed in more detail, students thought the essay would, at best, only pass.

We felt the exercise itself was very useful to students in terms of developing critical thinking skills.²⁵ By exposing the limitations of GenAI, we felt that students would be able to use it responsibly. It showed students that the use of GenAI can potentially limit them. For instance, it may not provide recent relevant examples to demonstrate their arguments or it may not include the full breadth and scope of the topic. Our most interesting discovery was that students wanted outputs to be personal, original and nuanced, just as teaching staff would in standard marking.

When asked to consider how an academic setting might address the advent of GenAI, students attached particular importance to the language used. For instance, instead of saying 'you should not use GenAI', they recommended saying 'you would be better off using GenAI in a particular manner'. They felt that certain core elements of learning might be impacted by GenAI, and that staff would need to be clear which learning experiences should be performed without the involvement of GenAI, protecting fundamental learning experiences where necessary.

Students also emphasised that faculty members should be trained on the fallibility of GenAI-detection tools, which are unable to identify all aspects of GenAI usage and are therefore inherently unreliable.

Notably, students felt strongly that in the early years of undergraduate study, where foundational legal writing skills are taught, GenAI should not significantly influence their thinking. As one student articulated, 'In those areas ... I would definitely would not want generative AI to influence my way of thinking ... there I would be, at least in the early years, I would be vigilant regarding use of generative AI beyond the brainstorming process.' This reflects a concern that relying on AI at the writing stage could undermine the development of essential skills such as structuring arguments, threading sentences together and demonstrating in-depth knowledge. Allowing students to bypass entirely the process of developing their own coherent arguments by relying on AI could impede their ability to master these fundamental aspects of legal writing. Another student

²⁴ For instance, on the value of peer review to the development of critical thinking skills, see Ardill, "Peer Feedback," 1. On providing opportunities for discussion and dialogue as a successful tool in supporting the development of critical thinking skills, see Abrami, "Strategies for Teaching Students," 275.

²⁵ As described in the previous note.

echoed this sentiment, stating that ‘there are more fundamental learning experiences in an undergraduate course where you would not want GenAI being involved. You would want students to learn ... [to] have those learning experiences without the involvement of GenAI.’ They contrasted this with master’s courses, which are advanced, implying postgraduate students might be expected to use GenAI differently due to their existing foundation of knowledge and critical skills. The observation was even made that the GenAI-generated essay analysed in the workshop ‘looked like an undergraduate essay and was not a postgraduate essay’, suggesting that relying on AI might produce work that lacked the depth and sophistication expected at a higher level.²⁶

5.3 Student-Generated Guidelines (Output)

During the final workshop, students worked with staff to develop guidelines for the responsible and effective use of GenAI in legal essay writing. Students produced these guidelines for their peers, to translate what they had learned through the project into a guide for writing essays in the age of GenAI while maintaining academic rigour. As these were current students, they chose to adapt an essay-writing guide produced by Dr Keller to incorporate a reference to GenAI usage. The flowchart they produced is described below.

According to the guidelines, the essay-writing process begins with broad readings and research conducted prior to the release of the essay question. Importantly, when the essay question is first released, the student should avoid using GenAI. The next phase involves focused reading based on the established reading list, followed by a brainstorming session where the student determines their central argument. Again, GenAI should be avoided at this stage. The brainstorming will naturally lead to ideas for additional reading tailored to establish whether the tentative argument is robust, allowing the student to identify any areas that would benefit from further investigation.

Following these stages, the student may optionally query GenAI about reading areas or general viewpoints on their topic. Here, it is important to exercise extreme caution, as the content may be misleading. The student must treat any AI output with the same scepticism they would apply to comments on internet forums or social media. The content may contain serious errors or fail to address important topics with appropriate gravity. Near the end of the research phase, the student should structure the essay using insights from brainstorming and additional reading. The student should understand that the initial structure is likely to continue to evolve as work progresses.

When planning, the student should be aware that the essay should incorporate both descriptive and analytical elements, and critical analysis is important for achieving distinction-level work. Depth should be prioritised over breadth. This means not merely covering as many different points as possible but going into detail about a more limited number of key areas. More space and a word count should be given to key points. The student should consider using case studies.

The student may use GenAI to explore alternative structural approaches once they have established a solid foundation. Crucially, the student should have confidence in their own assessment of the materials and their initial structure, and only make amendments where genuinely necessary.

Before submission, the student reviews the essay. While unlikely, it may be necessary to do some additional reading. Citations should have been gathered throughout the process and consistently added according to the relevant citation style.

5.4 Staff Training and Awareness Regarding GenAI

A key finding from the project concerns the urgent need to prepare staff to support students in using AI effectively and ethically. Students actively contributed suggestions for addressing this issue, recognising that staff training is a crucial first step. One student proposed a multi-step approach: ‘Step one: Educate the professors, educate the students. Step two: Let the professors take a call as per their own individual subjects. How do they want to deal with it? Step three: Keep having ... feedback session[s] and make sure this recommendation or whatever the professor decided is renewed every year or every time the module starts so that they can always evolve with it.’ This student-led recommendation emphasises not only initial training but also ongoing learning and flexibility due to the rapid evolution of the technology.

Students felt faculty members should be advised on how GenAI works and they should be given the training to understand specific ways in which it can be used, such as structuring, explaining difficult concepts or finding literature.

²⁶ One of the key concerns related to appropriate citation of sources. This is in line with research conducted in other disciplines – for instance, Cohen, “Generative Artificial Intelligence.” 2, highlighting issues with producing inaccurate information particularly with referencing external sources.

Students identified the fallibility of AI detection tools as a crucial area for staff awareness. As one student stated, ‘faculty needs to be made aware ... it is not accurate in its detection, it can only hypothesise it was written by it’. They referenced ChatGPT’s own FAQs, and gave a notable example: ‘If you run the American Declaration of Independence through GPTZero, it will turn out as AI-generated.’ This underscores the students’ concern that faculty ‘should be made aware that its accuracy is not 100 per cent’. Addressing this lack of awareness among staff is therefore vital to avoid falsely accusing students and would foster a ‘culture of trust rather than detection or policing of students’.

6. Discussion

Meadows describes how complex systems react to a change in the environment in which they operate; this wide-ranging analysis applies to thermostats just as much as it does to the education or justice system.²⁷ In both legal academic tasks and legal professional tasks, critical thinking is at the core. Critical thinking is not merely a component of legal education; it is the paradigmatic foundation upon which both legal education and the legal system itself are constructed. When GenAI threatens to bypass this core function, interventions that strengthen critical evaluation skills protect not just academic integrity, but also the essential nature of legal reasoning itself. Following Meadow’s analysis, the system can most effectively be influenced not by policing, which operates at the lower-leverage level of rules and supervision, but by engaging at a paradigm level, by creating a culture of trust where students are actively involved in the review and critique of AI-generated texts.²⁸

Our approach to developing guidelines was based on a process of co-creation or students as partners. In literature, these terms have a slightly different meaning, with co-creation referring to ‘a meaningful collaboration between students and staff, with students becoming more active participants in the learning process, constructing understanding and resources with academic staff’.²⁹ The term ‘students as partners’ is broader and refers to ‘a collaborative, reciprocal process through which all participants have the opportunity to contribute equally, although not necessarily in the same ways, to curricular or pedagogical conceptualization, decision making, implementation, investigation, or analysis’.³⁰ Either way, the idea is to empower students and give them a voice in the design and development of the guidelines.

Based on these two stages, our preliminary recommendations included the need for educating teaching staff and students, assisting staff to ensure that fundamental learning experiences are protected. It will be necessary to clarify reserved activities – those activities where using GenAI is not compatible with the pedagogical aim (a working group could assist with this process). Exploratory sessions critiquing GenAI outputs should be held within the curriculum (on a subject basis) as part of the technique/pedagogy. For example, a formative assignment could include an exploratory session to produce essays that can be dissected by staff and students.

6.1 Recommendations

We make the following recommendations:

- Integrate the analysis of GenAI-produced texts into formative assessments rather than a separate education program, allowing staff and students to work together analysing these outputs. This promotes mutual learning, reveals the evolving capabilities and limitations of GenAI and embeds AI literacy into core legal teaching in an agile, responsive way.
- Establish an internal faculty working group.
- Maintain continuing, iterative development of student-led guidance, leading to a culture of trust.

As discussed above, modelling is a recognised and effective technique. Taking this beyond the workshop, we recommend incorporating this kind of exercise in tutorials as part of standard teaching in any module. GenAI is a technology that cuts across specialisms.³¹ By using formative exercises in tutorials and co-critiquing AI-generated texts with teaching staff, multiple pedagogical aims are met. A sense of belonging can be fostered as students see the exercise not as a ‘risk’ but rather as a learning opportunity.³² This builds not only expertise but a sense of belonging that is particularly difficult to achieve in the digital age of education. In interviews, students reported very positive views on the exercises that took place. This exercise will

²⁷ Meadows, “Thinking in Systems,” 2, 36.

²⁸ Meadows, “Thinking in Systems,” 136, 163.

²⁹ Bovill, “Addressing Potential Challenges,” 195, 197.

³⁰ Cook-Sather, “Engaging Students As Partners in Learning and Teaching,” 2014, 6.

³¹ Ahn, “Do as I Do.”

³² Larsen, “A Sense of Belonging.”

also ensure that the students are prepared for digital skills in their working life by exposure to these technologies and building sophistication in critiquing outputs.³³

6.2 Limitations and Transferability

This study has several important limitations. A sample drawn from a single institution and specific legal specialisation (financial and fintech law) means that application across different legal domains or educational contexts should only occur with an awareness that there may be limitations. The period of the study was relatively short, and future assessment of long-term retention of critical evaluation skills should follow. Additionally, the methodology relied on qualitative analysis of student responses rather than quantitative measures that, while rich in insight, may not capture the full range of student experiences or learning outcomes.

6.3 Implications for Legal Professional Formation

The broader implications of these findings for how legal education prepares students for professional practice can be outlined briefly. Despite the concern within legal education and the legal profession as a whole with the advent of GenAI, the students placed great emphasis on originality, critical analysis, and nuanced reasoning in their interactions. This suggests that traditional legal education values remain highly relevant in an AI-integrated environment. Rather than viewing GenAI as a threat to these competencies, our results indicate that structured engagement with these tools can actually strengthen students' appreciation for and development of distinctly human analytical capabilities. This goes beyond legal education and into legal practice, where the ability to exercise judgement in complex, contextual situations remains central to legal expertise. Indeed, one student's approach illustrates the application to working life as well as academic life: "I would say think of it as an intern under you, a very hard-working intern, not a high professional".

7. Conclusion

This project focused on a core element of both the education system and the legal system: the ability to critically analyse. This was achieved by encouraging students to analyse an AI-generated text according to an existing marking rubric, reflecting on the process and co-creating guidelines for AI use in assignments. We observed that students were able to identify both surface-level and structural limitations in AI-generated legal texts. The process supported the development of AI literacy grounded in the values of academic integrity and independent thought.

A key recommendation emerging from the project is the incorporation of AI-text analysis within standard tutorial settings, involving both students and teaching staff in the critical evaluation of GenAI outputs. This approach extends the modelling techniques described by Ahn, Hu and Vega,³⁴ and is designed to support mutual learning while remaining responsive to technological developments. Embedding such exercises within the curriculum, rather than positioning them as isolated or ancillary, allows for a more agile pedagogical response that can be adapted according to the needs of specific modules or student cohorts.

This work does not aim to generalise beyond the legal domain. On the contrary, its focus is the centrality of critical, interpretative and evaluative thinking in legal education and the profession more broadly. The practices explored here are intended to strengthen these professional competencies within an evolving context.

In line with the Russell Group's principles on AI and education,³⁵ we conclude that the responsible integration of GenAI into legal education should not rest on detection or restriction; instead, it requires structured opportunities for engagement, critical reflection and collaborative scrutiny.

Acknowledgements

This research was supported by the King's College London College Teaching Fund Awards: Innovations in Teaching, Assessment and Feedback in the Age of Generative Artificial Intelligence. We are especially grateful to the following students for their contributions to the workshops, interviews and co-creation of the GenAI guidelines: Kanya Vimalanath, Sanpipat Huangswat, Mridula Nair, Alaa Alkathlan, Surya Kale, Anastasiia Zubriy, Anjali Gurumoorthy, Shreyan Srivastava, Adithya Nair and other anonymous participants. Their insights and engagement were instrumental to the project.

³³ Traester, "Pedagogy to Disrupt," 329–349.

³⁴ Ahn, "Do as I Do."

³⁵ Russell Group, "New Principles," 3.

Bibliography

- Abrami, Philip C., Robert M. Bernard, Eugene Borokhovski, Anne Wade, David I. Waddington, Tonje Persson and Caroline Levinton. "Strategies for Teaching Students to Think Critically: A Meta-Analysis." *Review of Educational Research* 85, no 2 (2015): 275–314. <https://doi.org/10.3102/0034654314551063>.
- Ahn, Janet N., Danfei Hu and Melissa Vega. "'Do as I Do, Not as I Say': Using Social Learning Theory to Unpack the Impact of Role Models on Students' Outcomes in Education." *Social and Personality Psychology Compass* 14, no 2 (2020): e12517. <https://doi.org/10.1111/spc3.12517>.
- Ardill, Nicole. "Peer Feedback in Higher Education: Student Perceptions of Peer Review and Strategies for Learning Enhancement." *European Journal of Higher Education* (2025): 1–25. <https://doi.org/10.1080/21568235.2025.2457466>.
- Bandura, Albert. 1977. "Self-Efficacy: Toward a Unifying Theory of Behavioral Change." *Psychological Review* 84, no 2: 191–215. <https://doi.org/10.1037/0033-295X.84.2.191>.
- Biggs, John B. *Teaching for Quality Learning at University: What the Student Does*, 4th ed. New York: McGraw-Hill, 2011.
- Bonawitz, Elizabeth, Patrick Shafto, Hyowon Gweon, Noah D. Goodman, Elizabeth Spelke and Laura Schulz. "The Double-Edged Sword of Pedagogy: Instruction Limits Spontaneous Exploration and Discovery." *Cognition: Probabilistic Models of Cognitive Development*, 120, no 3 (2011): 322–330. <https://doi.org/10.1016/j.cognition.2010.10.001>.
- Chan, Carol K. and Wanqiu Hu. "Students' Voices on Generative AI: Perceptions, Benefits, and Challenges in Higher Education." *International Journal of Educational Technology in Higher Education* 20, no 43 (2023): 1–32. <https://doi.org/10.1186/s41239-023-00411-8>.
- Coffey, Lauren. "Students Outrunning Faculty in AI Use." *Inside Higher Ed*, October 31, 2023. <https://www.insidehighered.com/news/tech-innovation/artificial-intelligence/2023/10/31/most-students-outrunning-faculty-ai-use>.
- Cohen, Jérémie F. and D. Moher. "Generative Artificial Intelligence and Academic Writing: Friend or Foe?" *Journal of Clinical Epidemiology* 179 (2025): 1–10. <https://doi.org/10.1016/j.jclinepi.2024.111646>.
- Cook-Sather, Alison, Catherine Bovill and Peter Felten. *Engaging Students as Partners in Learning and Teaching: A Guide for Faculty*. San Francisco: Jossey Bass, 2014.
- Dell'Acqua, Fabrizio, Edward McFowland, Ethan R. Mollick, Hila Lifshitz-Assaf, Katherine Kellogg, Saran Rajendran, Lisa Krayner, François Candelon and Karim R. Lakhani. "Navigating the Jagged Technological Frontier: Field Experimental Evidence of the Effects of AI on Knowledge Worker Productivity and Quality." 2023. SSRN Scholarly Paper. <https://doi.org/10.2139/ssrn.4573321>.
- Gonsalves, Chahna. "Generative AI's Impact on Critical Thinking: Revisiting Bloom's Taxonomy." *Journal of Marketing Education* (2024): 1–18. <https://doi.org/10.1177/02734753241305980>.
- Hyde-Vaamonde, Cari and Sunny Dhillon. "Teaching in a Transitional Space: A Reflection on Modelling Zotero in Tutorials & 5 Tips for Graduate Teaching Assistants When Adopting Open-Source Software." *Postgraduate Pedagogies* 4, no 1 (2024): 94–123. <https://doi.org/10.82191/pp.27>.
- İpek, Ziyaeddin Halid, Ali İbrahim Can Gözümlü, Stamatios Papadakis and Michail Kallogiannakis. "Educational Applications of the ChatGPT AI System: A Systematic Review Research." *Educational Process International Journal* 12, no 3 (2023). <https://doi.org/10.22521/edupij.2023.123.2>.
- Lakharni, K. R. "AI Won't Replace Humans – but Humans with AI will Replace Humans Without AI." *Harvard Business Review*, August 4, 2023. <https://hbr.org/2023/08/ai-wont-replace-humans-but-humans-with-ai-will-replace-humans-without-ai>.
- Larsen, Ana and Trixie James. "A Sense of Belonging in Australian Higher Education: The Significance of Self-Efficacy and the Student–Educator Relationship." *Journal of University Teaching and Learning Practice* 19, no 4 (2022). <https://doi.org/10.53761/1.19.4.5>.
- Lee, Chien Ching and Ming Low. "Using GenAI in Education: The Case for Critical Thinking." *Frontiers in Artificial Intelligence* 7 (2024). <https://doi.org/10.3389/frai.2024.1452131>.
- Meadows, Donella H. *Thinking in Systems*. New York: Chelsea Green Publishing, 2008.
- Naqvi, Waqar M., Rohini Ganjoo, Michael Rowe, Aishwarya A. Pashine and Gaurav V. Mishra. "Critical Thinking in the Age of Generative AI: Implications for Health Sciences Education." *Frontiers in Artificial Intelligence* 8 (2025): 1571527. <https://doi.org/10.3389/frai.2025.1571527>.
- Prakash, G. Aswathy and Vishnu Nair. "Integrating Generative AI into Legal Education: From Casebooks to Code, Opportunities and Challenges". *Law, Technology and Humans* 6, no 3 (2024): 60–79. <https://doi.org/10.5204/lthj.3640>.
- Traester, Mary, Chris Kervina and Noel Holton Brathwaite. "Pedagogy to Disrupt the Echo Chamber." *Pedagogy* 21, no 2 (2021): 329–349. <https://doi.org/10.1215/15314200-8811517>.
- Yusuf, A., Nasrin Pervin and Marcos Román-González. "Generative AI and the Future of Higher Education: A Threat to Academic Integrity or Reformation? Evidence from Multicultural Perspectives." *International Journal of Educational Technology in Higher Education* 21 (2024). <https://doi.org/10.1186/s41239-024-00453-6>.