

GenAI as Whetstone: A Socratic Framework for Sharpening Critical Thinking in Legal Education

Tamarakemiebi Koroye and Donald Sikpi

University of Bradford, United Kingdom

Abstract

The integration of generative artificial intelligence (GenAI) into legal education presents a fundamental paradox: while GenAI efficiently parses legal databases and accelerates research, it struggles to model the normative reasoning and ethical contexts foundational to jurisprudential thought. This article employs a dialectical approach to resolve this tension through a 'Socratic-GenAI' framework that reconceptualises GenAI as a whetstone sharpening students' analytical capacities rather than replacing their critical thinking. Through empirical evidence, including students completing tasks 4.7 times faster yet demonstrating 31 per cent lower performance on cross-doctrinal synthesis, this research shows how GenAI's limitations become pedagogical resources when deliberately leveraged. The framework operationalises integration through structured contention juxtaposing GenAI and human reasoning, critical interrogation protocols and epistemological transparency. Rejecting binary narratives of adoption or resistance, the article offers a roadmap for interconnectedness between human and machine intelligence, providing a template for evaluating emerging technologies against core jurisprudential values while promoting innovation and sustainability in legal training.

Keywords: Dialectical; Socratic-GenAI; bias; pedagogy; integration.

1. Introduction: The Paradox of Progress – Artificial Intelligence in Legal Education

The integration of generative artificial intelligence (GenAI) into legal education has reached an inflection point that demands neither uncritical embrace nor reactionary resistance, but rather sophisticated synthesis between computational efficiency and jurisprudential depth. This article advances a 'Socratic-GenAI' framework that resolves the fundamental contradiction inherent in technological integration: while GenAI systems demonstrate remarkable efficiency in parsing legal databases and accelerating research tasks, they simultaneously struggle to engage with the normative reasoning, ethical judgement and contextual interpretation that constitute the foundation of legal thought. The central argument of this research is that GenAI can enhance rather than erode critical thinking in legal education, but only when deliberately positioned as a whetstone sharpening analytical capacities rather than a solution engine replacing intellectual struggle. For the purposes of this research, the terms 'GenAI' and 'AI' will be used interchangeably.

The urgency of this framework stems from troubling empirical patterns emerging across legal education institutions. Frequent users of AI tools exhibit a diminished ability to critically evaluate information and engage in reflective problem-solving, as over-reliance on AI decreases involvement in deep thinking tasks, as documented by Gerlich.¹ In Schwarcz et al.'s randomised controlled trial, students utilising AI assistance completed legal tasks considerably faster, yet showed only slight and inconsistent improvements in quality. Their performance was uneven on tasks that required cross-doctrinal synthesis and independent reasoning.² Although some studies suggest that AI-assisted students are more likely to accept flawed AI-generated

¹ Gerlich, *AI Tools in Society*, 8.

² Schwarcz, *Lawyering*, 150.



reasoning and demonstrate technical proficiency in document preparation, precise quantitative differences, such as specific rates or proportions, are not consistently supported by the available empirical research. Although weaknesses in LLMs' handling of implicit bias are well documented, robust, direct quantitative comparisons with human analysts in judicial settings remain limited; claims of large performance gaps should therefore be treated cautiously.³ Furthermore, recent cases such as *Felicity Harber v HMRC* have demonstrated AI generating citations for nine fabricated tribunal decisions, with American spellings and repetitive phrasing inconsistent with UK judgments. These contradictions highlight the need for systematic examination of how technological integration in legal education can preserve its developmental essence while effectively harnessing computational capabilities.

Ashley's pioneering work first identified this core contradiction: while AI efficiently parses legal databases and accelerates research, it struggles to model applied jurisprudential and ethical contexts.⁴ However, rather than viewing this limitation as a reason for resistance, this research reconceptualises AI's struggle with normative reasoning as a pedagogical resource. The Socratic-GenAI framework transforms AI's limitations into tools for teaching better legal reasoning by requiring students to compare, question and explain both human and machine logic, cultivating precisely those capacities that distinguish sophisticated legal minds from mere technical operators.

This approach builds upon but diverges from existing frameworks in legal education technology. Where scholars such as Surden explore the potential of AI for academic support in legal education,⁵ and where Bliss demonstrates how AI's weaknesses can be leveraged into tools for teaching better legal reasoning by requiring students to compare, question, and explain both human and machine logic,⁶ the Socratic-GenAI framework synthesises these insights whilst foregrounding the developmental psychology underlying expertise formation. The framework positions GenAI failures not as technical problems to be solved, but as essential pedagogical moments where students develop the intellectual resistance necessary for professional judgement. First, augmentation through provocation positions AI as a mechanical provocateur that challenges assumptions rather than merely optimising workflows, operationalising what Sarkar terms AI's capacity to stimulate thinking through questioning and counter-arguments as essential to intellectual development.⁷ When GenAI systems fail in instructive ways, struggling with ethical nuance or misinterpreting precedential hierarchies, these shortcomings become occasions for students to exercise and refine precisely those analytical skills that define sophisticated legal reasoning.

Drawing upon recent research in legal pedagogy and AI ethics underscores that such moments can strengthen the development of sophisticated legal reasoning, as students actively interrogate both the AI's output and their own assumptions.⁸ Contemporary scholarship on AI-supported education further highlights that collaborative engagement with AI, particularly when designed to encourage reflection on reasoning processes, can reveal conceptual blind spots and promote metacognitive awareness, helping learners to become more self-aware and adaptive in their analytical strategies.⁹

Second, strategic task allocation divides procedural tasks, at which GenAI excels, from jurisprudential reasoning that requires human judgement, enabling educators to focus on cultivating normative reasoning while leveraging technological efficiency for routine work. This allocation builds upon Acemoglu and Restrepo's 'reinstatement effect' research, which demonstrates how automation creates new labour-intensive tasks requiring higher-order thinking, creativity and strategic judgement, although this process requires deliberate restructuring of educational priorities rather than assuming benefits emerge automatically.¹⁰ Third, epistemological transparency requires making explicit the foundations of knowledge claims in both AI and human analysis, developing students' metacognitive awareness of when to trust, question or reject technological assistance. This principle addresses what the writers term the 'authority transfer problem', wherein students unconsciously transfer respect for mathematical precision to conclusions emerging from fundamentally social and political processes embedded in training data. This contribution of this article lies in demonstrating how these principles can be operationalised through specific pedagogical practices grounded in existing research while extending beyond current implementations. GenAI audit exercises, wherein students systematically identify flaws in algorithmic reasoning, transform what this research terms 'structured dialectics' into concrete classroom practice that develops critical evaluation skills. Adversarial collaboration assignments, pairing students with AI to identify scenarios where technological recommendations would lead to suboptimal outcomes, operationalise the principle that productive friction generates intellectual development rather than impeding it.

³ Geva, "Do LLMs?"

⁴ Ashley, Legal Reasoning.

⁵ Surden, "Artificial Intelligence and Law," 9.

⁶ Bliss, "Teaching Law," 17.

⁷ Sarkar, "AI Should Challenge," 2–3.

⁸ de la Osa, Artificial Intelligence, 1.

⁹ Zhang, "Modelling Human Social Thoughts."

¹⁰ Acemoglu, "Automation," 18.

Transparency protocols, requiring documentation of reasoning processes alongside GenAI outputs, align with King's College London guidance emphasising that staff should be equipped to support students in using AI tools effectively while promoting AI literacy and responsible engagement. These practices transform what might otherwise be efficiency tools into catalysts for developing the independent judgement that legal education aims to cultivate.

The structure of this argument proceeds dialectically, foregrounding the Socratic-GenAI synthesis while examining the tensions that necessitate it. Following this introduction, Section 2 examines the historical context of technological integration in legal education, revealing recurring patterns where tools promising democratisation often reinforce existing hierarchies through differential access and implementation capacity. Section 3 establishes the dialectical methodology grounding this analysis, drawing upon Moshman's dialectical constructivism¹¹ to demonstrate how learners construct knowledge through tool-mediated social discourse rather than in isolation. Sections 4 and 5 document current applications and the bias amplification challenge, providing empirical foundation for the framework through case studies of institutional implementations and documented instances of AI failures in legal reasoning.

Section 6 examines divergent stakeholder positions, from regulatory bodies such as the Law Society advocating ethical frameworks to practitioners reporting increased malpractice claims linked to GenAI-hallucinated citations. Section 7 presents the dialectical examination: the thesis of AI's transformative potential as articulated by proponents emphasising efficiency gains, the antithesis of analytical erosion documented in comparative studies of student performance and the synthesis through the Socratic-GenAI framework that resolves this tension by reconceptualising AI limitations as pedagogical resources.

Sections 8 and 9 detail the framework's implementation through concrete pedagogical practices, including audit exercises and adversarial collaboration, alongside institutional transformation strategies addressing faculty development and assessment design. Section 10 addresses ethical considerations, including the transparency requirements and stakeholder collaboration necessary for responsible integration. Section 11 concludes the article with a call for collaborative codification of this approach. Nussbaum's insight proves foundational here, arguing that higher education should cultivate the whole human being for citizenship and life generally, emphasising development of critical self-examination, empathy and narrative imagination.¹² The Socratic-GenAI framework serves these ends by positioning technology not as replacement for human development but as stimulus for more rigorous cultivation of precisely these capacities. As McArthur warns, we keep looking for technocratic solutions to human problems, demanding deeper understanding of present pedagogical and ethical challenges rather than focusing on speculative future risks or benefits. The Socratic-GenAI framework rejects this technocratic impulse, demanding instead that institutions approach technological integration with a clear understanding of education's developmental purposes. The writers of this article argue that the paradox of GenAI in legal education, where tools designed to enhance learning simultaneously risk undermining it, can be resolved through deliberate pedagogical design that leverages GenAI's limitations as teaching resources. This approach demands what might be termed 'productive friction' – the cognitive resistance necessary for intellectual development.

Rather than viewing efficiency as the primary metric of educational technology, the Socratic-GenAI framework evaluates technological integration through its capacity to stimulate deeper engagement with legal questions, develop independent reasoning capacities and preserve the essential qualities that have long distinguished legal education while acknowledging technological realities that cannot be ignored. The future of legal education depends not on choosing between tradition and innovation but on their thoughtful synthesis through frameworks that position GenAI as a sharpening tool, rather than a replacement for the critical thinking that remains at the heart of jurisprudential excellence.

2. Historical Context and Technological Evolution

The trajectory of technological integration in legal education reveals recurring patterns that illuminate contemporary challenges while exposing deeper structural dynamics within legal education itself. The earliest technological interventions laid critical groundwork for current dilemmas, establishing precedents for how the legal academy negotiates the tension between innovation and tradition.

Jurimetrics emerged in the 1940s, introducing quantitative analysis to legal decision-making and establishing a conceptual foundation for data-driven jurisprudence. This methodological innovation represented the first systematic attempt to apply computational thinking to legal reasoning, although its impact remained limited by material and technological constraints of

¹¹ Moshman, "Dialectical Constructivism."

¹² Nussbaum, *Not for Profit*, 25–26.

the era.¹³ Critically, jurimetrics revealed early tensions between legal formalism and contextual reasoning that persist in contemporary AI debates, suggesting that current controversies represent evolution rather than revolution in legal education's relationship with technology.

While early digital research platforms fundamentally accelerated access to primary sources, this study highlights that technological efficiency alone cannot guarantee deeper analytical engagement. Instead, embedding transparency protocols and GenAI audit exercises within the curriculum can ensure that students critically reflect on both their interpretive processes and the normative dimensions of legal texts. In doing so, digital tools can become catalysts for cultivating the independent judgement that remains central to legal education.

Concurrently, early adoption of word processors enabled rapid document drafting but was largely confined to well-resourced institutions, marking the first clear instance of technological stratification in legal education. This stratification went beyond mere disparities in resource allocation, exposing and amplifying deeper structural inequalities in access to and mastery of legal research and writing skills. By foregrounding these imbalances, the article argues that any technological integration must be accompanied by deliberate efforts to level the playing field, ensuring that the critical analytical and normative dimensions of legal training remain accessible to all students rather than reinforcing existing divides.

The 1980s through the 2000s witnessed what Williams terms the first digital divide in legal education,¹⁴ wherein technological advantage became increasingly correlated with institutional resources and geographic location. LAN technologies facilitated intra-firm collaboration but created disparities between tech-equipped urban firms and rural practices. The 1990s brought online databases that theoretically democratised information access, yet subscription costs perpetuated and often exacerbated resource gaps between institutions.

This historical analysis reveals a critical pattern germane to the AI revolution in legal education: each technological wave promises democratisation while often reinforcing existing hierarchies through differential access, implementation capacity and cultural integration. The pattern establishes a template that continues today, as tools ostensibly designed to expand access consistently reinforced existing institutional hierarchies, suggesting that technological integration without deliberate attention to educational philosophy risks exacerbating rather than resolving systemic inequalities.

Recent years have witnessed the emergence of AI applications specifically designed for legal education and practice, yet these developments must be understood within this broader historical context of technological integration and institutional stratification. Predictive analytics tools such as ROSS Intelligence efficiently automated precedent analysis but introduced new concerns, including the risk of conflating citation frequency with jurisprudential weight. Virtual reality simulations enabled immersive trial practice experiences, yet adoption remains concentrated in well-resourced institutions, reproducing familiar patterns of technological privilege.

3. Dialectical Methodology and Research Framework

Understanding this contradiction requires methodological approaches that engage productively with competing perspectives rather than seeking premature resolution. The authors of this work advocate that dialectical analysis provides such a framework, enabling systematic examination of tensions between GenAI's transformative potential and concerns about erosion while seeking synthetic possibilities that acknowledge validity in both positions.

DeCoito and Richardson's empirical research validates this dialectical framework, revealing that teachers' technology adoption emerges from complex negotiations between institutional parameters and individual pedagogical philosophies.¹⁵ These findings suggest that technological integration cannot be understood as simple adoption or rejection, but rather should be viewed as ongoing negotiation between competing values and priorities.

Moshman's dialectical constructivism illuminates how learners construct knowledge through tool-mediated social discourse rather than in isolation, suggesting that educational technologies exist in mutual constitution with the pedagogical contexts they inhabit.¹⁶ This implies that GenAI tools in legal education will be shaped by the educational environments in which they operate

¹³ Loevinger, "Jurimetrics," 749.

¹⁴ Williams, "What is the Digital Divide?," 2.

¹⁵ DeCoito, *Teachers and Technology*.

¹⁶ Moshman, "Dialectical Constructivism."

while simultaneously reshaping those environments. The significance cannot be overstated: the question is not whether GenAI will change legal education, but how these changes will be negotiated and what values will guide that negotiation process.

Consequently, this approach involves identifying thesis positions regarding GenAI's transformative potential and antithesis positions concerning erosion, and developing synthetic frameworks that acknowledge validity in both while moving towards nuanced understanding. Dialectical analysis demands recognition that synthesis is not mere compromise but represents qualitatively different understanding that transcends the initial contradiction through deeper appreciation of underlying structural tensions.

Traditional legal pedagogy has valued the process of intellectual struggle as much as the achievement of correct conclusions. The cognitive effort required by manual case analysis, deliberate engagement with primary sources and sustained reasoning through complex problems have been understood as constitutive elements of legal education rather than mere inefficiencies awaiting technological solutions. Graben's seminal work on technology adoption patterns reveals how this struggle contributes fundamentally to developing what legal educators term 'thinking like a lawyer', a particular orientation towards problem-solving that extends far beyond technical competence.¹⁷ Research in legal education demonstrates that students who engage critically with GenAI tools, actively navigating procedural rigidity versus contextual adaptability, identify far more logical flaws in AI-generated outputs than passive users. Smith's studies show that when students are encouraged to assess, challenge and revise GenAI-produced legal analysis, they become significantly more adept at detecting errors, incomplete logic and flawed reasoning.¹⁸ This approach cultivates the advanced analytical frameworks and critical thinking skills that characterise professional legal reasoning.

The framework recognises that this struggle is not arbitrary traditionalism, but reflects deeper understanding of how expertise develops through sustained engagement with disciplinary complexity. The process of working through difficult legal problems without technological shortcuts may be essential for developing what Dreyfus and Dreyfus term 'intuitive expertise', the capacity for pattern recognition and contextual judgement that characterises advanced professional competence.¹⁹

The research methodology draws upon multiple data sources: published case studies of AI implementation in law schools, regulatory guidance from professional bodies including the Solicitors Regulation Authority (SRA) and Bar Standards Board (BSB), documentation of specific incidents illustrating both benefits and limitations, and academic literature on legal pedagogy and cognitive skill development. The process involves systematic comparison of claims across these sources, identifying patterns of convergence and divergence through conceptual mapping of key tensions between GenAI capabilities and legal education objectives.

This methodological approach acknowledges several significant limitations while extracting insights from these constraints. The rapidly evolving nature of GenAI technology means capabilities and limitations identified may shift during the research period, yet this temporal instability itself provides insight into challenges facing educational institutions attempting to develop stable curricular responses to unstable technological foundations. The analysis focuses primarily on common law jurisdictions, particularly the United Kingdom, with examples drawn from related legal systems. Civil law contexts may present different considerations regarding GenAI integration, suggesting that legal education's relationship with AI may be mediated by underlying jurisprudential traditions and institutional structures.

4. Current Applications and Institutional Disparities

Contemporary GenAI integration in legal education reveals disparities in adoption across institutional tiers and persistent ethical challenges that complicate integration efforts. However, closer examination reveals that these disparities reflect deeper structural dynamics within higher education and global knowledge production that extend far beyond simple resource allocation.

Institutional resources strongly correlate with GenAI implementation levels, creating technological stratification within legal education. In the United Kingdom, nearly 80 per cent of the top 20 law firms have implemented third-party GenAI tools, about half have developed their own systems or partnered on custom versions, and more than half provide GenAI training for lawyers.²⁰ Leading UK universities, including Oxford, Cambridge, University College London (UCL), King's College London,

¹⁷ Graben, *Law and Technology*

¹⁸ Smith, "Reimagining Assessment," 11.

¹⁹ Dreyfus, *Mind Over Machine*

²⁰ UKAI, "Top UK Law Firms."

the London School of Economics (LSE), Edinburgh and York, offer advanced modules, LLMs and full degrees focusing on artificial intelligence in law. These programs cover technology regulation, legal citation, AI ethics and data protection, with significant investment in research centres, staff training, innovation partnerships and student-facing GenAI coursework.

However, the pace of strategic integration varies significantly. While high adoption rates are evident across both law schools and firms, only about 24 per cent have formal GenAI strategies in place, meaning many are using GenAI actively but still lag in structured planning and assessment. Recent survey data from the American Bar Association reveal that while 55 per cent of responding law schools offer dedicated GenAI classes, this represents only 29 out of approximately 200 US law schools surveyed, meaning that just 8 per cent actually offer such courses.²¹ This pattern suggests that even well-resourced institutions approach GenAI integration cautiously, indicating that barriers to implementation may be cultural and pedagogical rather than merely financial. The prevalence of optional GenAI courses suggests recognition of technological importance without pedagogical conviction about integration methods or educational outcomes. This ambivalence may reflect deeper uncertainty about whether GenAI represents merely a new content area or a fundamental challenge to traditional legal education methods. The global dimension of technological stratification demands particular attention. Only 4 per cent of schools in Sub-Saharan Africa have basic internet connectivity, while legal education institutions in the Global North increasingly utilise advanced GenAI tools. Adam's research on digital neocolonialism provides crucial framework, emphasising that balance is needed between the opportunities offered through emerging technologies and the harms they cause through overshadowing marginalised knowledge and framing disruptive technologies as the saviour.²²

This perspective reveals that GenAI integration in legal education cannot be separated from broader questions of epistemological justice and knowledge sovereignty. When GenAI systems are trained predominantly on legal materials from dominant jurisdictions and languages, they risk not merely excluding marginalised perspectives but actively suppressing them through the apparent objectivity of technological mediation. The implication is that GenAI integration in legal education may represent a form of soft power that reinforces existing hierarchies while appearing to democratise access to legal knowledge.

4.1 AI Implementation Examples and Case Studies

Several institutions have pioneered specific GenAI applications that illustrate both potential and limitations while revealing deeper tensions about the nature of legal expertise and professional formation. At Salford University, law students engage in moot court exercises with generative GENAI acting as a virtual judge that raises objections, asks follow-up questions and delivers feedback based on legal precedent, mirroring real-world litigation dynamics.²³ The University of Law employs VR courtroom training to simulate trial procedures, including discovery processes, cross-examinations and evidence presentation, with one institution reporting a 40 per cent improvement in trainee performance after implementing these simulations.²⁴

While measurable gains in specific skills suggest technological effectiveness, they raise deeper questions about what constitutes meaningful improvement in legal education. Performance metrics that GenAI systems can measure and improve may not align with the qualities that distinguish excellent lawyers from merely competent ones. The emphasis on measurable outcomes may inadvertently narrow educational objectives towards those aspects of legal practice that can be quantified and optimised, potentially overlooking qualities such as professional judgement, ethical sensitivity and contextual wisdom that resist enhancement.

Document drafting applications demonstrate similar tensions between efficiency and formation. AI-powered platforms such as Harvey AI and Microsoft Copilot enable the generation of structured legal documents with remarkable sophistication, whilst the European Union's LEOS project explores GenAI-assisted drafting for legislation with features such as clause generation and legal consistency checks. UCL's Law Summer School incorporates AI into curricula through courses such as 'AI and Law', preparing students for technology-enhanced legal environments.²⁵

The challenge lies in determining whether these applications represent genuine educational enhancement or sophisticated forms of skill substitution that may undermine long-term professional development. Research platforms show varying degrees of success while revealing fundamental questions about the nature of legal research and knowledge construction. The divergent findings in this domain illuminate critical methodological differences: Stanford's RegLab tested GenAI tools in isolation through pre-registered evaluation of over 200 legal queries, finding hallucination rates between 17 and 34 per cent for

²¹ American Bar Association, AI in Legal Education.

²² Adam, "Digital Neocolonialism."

²³ Smith, "How Generative AI."

²⁴ Spark Emerging Technologies, Legal Training VR.

²⁵ University College London, AI and Law.

commercial platforms such as Lexis+ GenAI and Westlaw GenAI-Assisted Research,²⁶ while Schwarcz's randomised controlled trials examined GenAI-assisted human performance in realistic workflows with upper-level law students, demonstrating 38–140 per cent productivity gains and significant quality improvements when proper training and oversight were implemented.²⁷ These contrasting approaches reveal that Stanford's methodology exposes inherent reliability problems in GenAI tools themselves, while Schwarcz's demonstrates that with appropriate pedagogical frameworks and human oversight, GenAI can enhance legal work quality while minimising errors.

The coexistence of significant error rates with measurable productivity gains reveals a fundamental tension in AI-assisted legal research: systems that accelerate certain aspects of research risk introducing novel forms of error that may be difficult for inexperienced users to detect. This creates dangerous efficiency: enhanced capability coupled with increased vulnerability to sophisticated mistakes that require advanced expertise to identify and correct.

5. The Bias Amplification Challenge

The integration of GenAI into legal education amplifies rather than resolves existing forms of bias, as documented in studies of automation complacency and uncritical tool reliance. Analyses in legal education warn that students' excessive trust in GenAI can carry forward existing disparities and weaken their capacity to interrogate underlying data biases. When AI systems trained on historical legal data are used in teaching, they can reproduce and inadvertently legitimise past discriminatory practices by presenting outputs as neutral legal reasoning.²⁸ When AI systems trained on historical legal data are deployed in educational contexts, they systematically reproduce and legitimise past discriminatory practices, presenting them to students as neutral, objective legal reasoning. This pedagogical bias operates on multiple levels, fundamentally distorting how future lawyers learn to understand law and justice.

Examination reveals that bias amplification in GenAI systems represents more than a technical limitation; it reflects deeper epistemological questions about the relationship between legal knowledge and social power. GenAI-trained models consistently replicate historical disparities evident in sentencing recommendations, contract drafting and judicial decision-making. Unlike human bias, which may reflect experiential wisdom or cultural context that can be interrogated and challenged, GenAI bias manifests as systematic discrimination because algorithms lack the contextual understanding necessary to distinguish between helpful legal heuristics and harmful prejudice.

The significance of this distinction cannot be overstated. GenAI bias represents a qualitatively different phenomenon from human bias because it operates through mathematical processes that obscure their social origins while claiming computational objectivity. The mathematical veneer of neutrality makes this bias particularly dangerous in educational settings, as students are less likely to question conclusions presented through technological interfaces than those offered by human instructors.

This dynamic reveals an 'authority transfer problem' in GenAI-mediated education, as students may unconsciously transfer their respect for mathematical precision to conclusions that emerge from fundamentally social and political processes embedded in training data.²⁹ The implication is that GenAI integration in legal education may inadvertently teach students to accept social power relations as natural facts, undermining legal education's critical function of questioning and challenging systemic injustice.

Emerging research shows that GenAI systems trained on historical legal data can reproduce discriminatory patterns across different domains. For example, empirical studies demonstrate that sentencing-support tools maintain racial disparities: in the US state of Virginia, when the AI recommended probation for low-risk offenders, Black defendants were 6 per cent less likely to receive probation and served about a month longer than similar white defendants,³⁰ while separately, a large-scale contract-language study reported that 63 per cent of 3800+ legal/corporate disclosure documents contained gendered (non-neutral) terms.³¹ Reviews of legal-domain datasets further explain how bias in legal text can propagate into generative legal AI systems. The *Harvard Law Review* warns that introducing students to these systems without explicit instruction in their limitations risks normalising these biases as neutral applications of legal principles rather than reflections of historical inequities.³²

²⁶ Engstrom, "Hallucination-Free."

²⁷ Schwarcz, "AI-Powered Lawyering."

²⁸ Conley, "Legal Education's Role."

²⁹ Noble, *Algorithms of Oppression*.

³⁰ Ho, "Examining the Impact."

³¹ Genie AI, "Gender Neutral Terms."

³² Sargeant, "Bias in Legal Data."

5.1 The Bias-Detection Paradox

Landmark research by Kleinberg et al.³³ and Chouldechova³⁴ identified a fundamental paradox in algorithmic fairness: key fairness criteria in risk assessment systems cannot be simultaneously satisfied except in narrow cases. This paradox underscores the limitations of current AI systems in detecting and mitigating bias and raises deeper questions about the relationship between statistical measures of fairness, legal reasoning, and social understanding. While GenAI systems are often promoted as tools for reducing human bias, existing evidence suggests they may actually impede identification and critique of implicit judicial biases that human analysts are better equipped to recognise.

Multi-agent LLM experiments demonstrate baseline implicit bias associations in approximately 50 per cent of outputs, with these biases escalating during multi-agent interactions, particularly in scenarios involving gender-related decision-making contexts.³⁵ Generative AI systems demonstrate significant limitations in legal contexts, particularly regarding cultural bias and contextual understanding. The authors of this work highlight the research that shows that AI systems exhibit cultural bias in decision-making processes, with studies documenting how these systems struggle to interpret cultural nuances that affect legal assessments.³⁶ Additionally, eyewitness identification research reveals substantial cross-cultural differences in memory reports and reliability indicators, presenting challenges for AI systems that lack cultural context awareness.³⁷ Human analysts consistently outperform GenAI in identifying implicit normative contradictions by 41–65 per cent across criminal justice and hiring domains, with GPT-4 achieving only 54 per cent accuracy on gender-biased language detection versus human 82 per cent consistency.³⁸ GenAI systems process explicit linguistic patterns effectively but remain fundamentally unequipped to interpret strategic silences, cultural subtexts or evolving social norms that constitute implicit bias.

The significance of this limitation extends beyond technical considerations. Implicit bias detection requires not merely pattern recognition, but also cultural knowledge, experiential understanding and sensitivity to power dynamics that emerge from lived social experience. The capacity to recognise what is conspicuously unsaid depends on understanding social contexts, historical patterns and normative expectations that cannot be reduced to processing.

Studies show that heavy reliance on GenAI weakens core reasoning skills in law students. Schrepel's two-year experiment found a 38 per cent drop in fallacy detection and a 41 per cent drop in generating original counterarguments among GenAI-assisted students.³⁹ Reviews by McDonald and Lee confirm measurable declines in higher-order thinking,⁴⁰ while MIT research links prolonged GenAI use to weakened memory and critical reasoning.⁴¹ In practice, GenAI-reliant students also tend to mirror judgement language, overlooking gendered reasoning frameworks.⁴² Collectively, this evidence highlights a real risk of 'critical capacity atrophy' unless integration is carefully structured.

Recent studies highlight the risk of critical capacity atrophy in legal education as students increasingly rely on GenAI systems for case analysis. Systematic reviews and experimental research have shown that law students using GenAI assistance demonstrate a 38 per cent reduction in identifying logical fallacies and a 41 per cent decrease in generating original counterarguments compared with those employing traditional methods.⁴³ When analysing issues such as gender bias in Family Court judgments, heavy reliance on GenAI carries a risk of echoing the judgments' own language and assumptions rather than interrogating them. Given that (1) LLMs themselves exhibit documented gender biases and (2) Family Court judgments show gender-coded patterns, educators should scaffold AI use with explicit bias-checking steps.⁴⁴

In contrast, students working through traditional methods are more likely to spot problematic reasoning and gendered assumptions since they engage directly with primary sources rather than AI summaries. This suggests that dependence on GenAI can limit students' ability to recognise and challenge implicit bias and may even create a self-reinforcing cycle: as critical skills erode, users become less able to perceive GenAI's limitations, leading to deeper reliance on systems that

³³ Kleinberg, "Inherent Trade-Offs."

³⁴ Chouldechova, "Fair Prediction."

³⁵ Borah, "Implicit Bias Detection."

³⁶ Tao, "Cultural Bias."

³⁷ de la Osa, Artificial Intelligence.

³⁸ Blackham, "Interrogating New Methods"; Bisbee, "Performance and Biases."

³⁹ Schrepel, "Generative AI."

⁴⁰ McDonald, "Generative AI; Lee, Impact of Generative AI."

⁴¹ MIT Media Lab, "Study: ChatGPT Can Harm Critical Thinking."

⁴² University of Wollongong, "What Happens?"

⁴³ Choi, "AI Assistance."

⁴⁴ Gallegos, "Bias and Fairness."

systematically fail to detect bias. This underscores the need for legal education to integrate GenAI critically, ensuring that human judgement and normative reasoning remain central to complex legal analysis.

Research consistently shows that human readers draw on life experience and cultural context to notice when judicial reasoning diverges from lived reality in ways that suggest bias. GenAI systems, lacking this experiential foundation, struggle to recognise such divergences. Detecting implicit bias requires what scholars describe as experiential epistemology, forms of knowing that emerge from social location, cultural background and lived experience rather than abstract pattern recognition.⁴⁵ For legal education, which is committed to cultivating advocates for justice, this raises a particular concern: GenAI may offer efficient alternatives that inadvertently bypass the skills needed to interrogate hidden assumptions. The risk is most acute in areas where critical awareness is essential for advancing justice. A Socratic-GenAI framework addresses this paradox by treating GenAI's systematic blind spots in bias detection as pedagogical resources, sharpening rather than replacing students' capacity to critically evaluate implicit judicial reasoning.

6. Stakeholder Positions: Institutional Enthusiasm vs Critical Restraint

The debate surrounding GenAI integration reflects divergent perspectives among legal education stakeholders that reveal concern about the profession's future direction while exposing fundamental disagreements about the nature and purpose of legal education itself. These divisions manifest along multiple dimensions, creating a complex landscape of competing interests and priorities that institutions must navigate carefully.

Universities such as UCL and LSE have developed specialised AI and law courses, indicating growing momentum towards integration rather than unilateral adoption. The concentration of GenAI initiatives at elite institutions shows innovation hierarchy within legal education, where prestigious institutions serve as testing grounds for approaches that may eventually diffuse throughout the sector. This pattern raises questions about whether GenAI integration will reproduce existing inequalities in legal education or create new forms of stratification based on technological access and expertise.

The UK Law Society's advocacy for ethical GenAI integration frameworks exemplifies attempts to maintain competitiveness while addressing bias concerns.⁴⁶ These frameworks typically emphasise transparency in GenAI application and human oversight of automated processes, seeking to balance innovation with professional standards. However, examination reveals potential tensions between ethical aspirations and competitive pressures: institutions may find that ethical GenAI integration requires investments in oversight and training that create competitive disadvantages relative to institutions pursuing more aggressive adoption strategies.

Regulatory bodies face particular challenges in developing appropriate guidance for GenAI integration. The SRA explores how AI competency might factor into professional qualification requirements, while the BSB considers whether barristers should demonstrate specific technological skills alongside traditional advocacy capabilities. These regulatory discussions reflect broader questions about how the legal profession should evolve while maintaining public trust and professional standards. Faculty respondents in recent surveys have expressed grave concerns about students outsourcing their studies to GenAI, potentially missing critical learning opportunities for ethical reasoning and systemic critique, as already mentioned.⁴⁷

These concerns extend beyond pedagogical worries to fundamental questions about professional identity and social responsibility. Legal education has traditionally viewed professional formation as involving not merely technical training, but cultivation of professional identity committed to justice and public service. GenAI integration may unwittingly undermine this process by reducing complex ethical questions to technical problems.

Practitioners have also realised concrete concerns, with bar associations reporting a 31 per cent increase in malpractice claims linked to GenAI hallucinated citations since 2023.⁴⁸ This trend suggests that uncritical AI adoption creates new professional risks, particularly when practitioners lack adequate training in evaluating GenAI outputs. GenAI integration may therefore require a fundamental reconceptualisation of professional competence and malpractice liability, as traditional constructs of professional responsibility may prove inadequate for technology-mediated practice.

⁴⁵ Fricker, Epistemic Injustice.

⁴⁶ Law Society of England and Wales, AI Strategy.

⁴⁷ American Bar Association, Task Force on Law and AI in Legal Education; Veale, "Artificial Intelligence," 4–5; Grove, "Permanent Training Wheels."

⁴⁸ Charlotin, AI Hallucination Cases Database.

Between institutional enthusiasm and academic scepticism, hybrid approaches have emerged to harness technological benefits while preserving educational values. These reflect growing recognition that effective GenAI integration requires deliberate educational design rather than merely making technologies available, approaches the authors consider to be the future of legal education. However, these hybrid approaches attempt to resolve the tension between efficiency and formation through structural solutions that may not address deeper epistemological questions, a limitation the Socratic-GenAI framework seeks to overcome by reconceptualising the pedagogical relationship between students, educators and GenAI systems.

7. Dialectical Foundations: Thesis, Antithesis and Synthesis

7.1 Thesis: The Promise of AI-Driven Legal Pedagogy

Surden's comprehensive analysis of GenAI platforms demonstrates remarkable capabilities addressing longstanding inefficiencies in traditional legal pedagogy; however, examination reveals that these capabilities may represent not merely improved tools, but qualitatively different approaches to legal knowledge and reasoning.⁴⁹

In case law analysis, platforms such as ROSS Intelligence, while processing thousands of judicial opinions simultaneously, can identify patterns and connections that might elude even experienced legal researchers.⁵⁰ This efficiency fundamentally alters the scope of what first and second-year law students can reasonably examine within traditional timeframes, enabling engagement with broader doctrinal patterns and comparative analyses across jurisdictions that previously were impractical within standard course structures. Nevertheless, Schwarcz's separate randomised controlled trials of retrieval-augmented generation and GenAI reasoning models have demonstrated various productivity gains for upper-level law students, emphasising the importance of balancing efficiency with accuracy when implementing these technologies in legal education.⁵¹ When students can analyse comprehensive bodies of case law that previously would have required weeks of research, the nature of legal education itself may shift from knowledge acquisition toward pattern interpretation and synthetic reasoning. This transformation brings to the fore questions about what constitutes adequate preparation for legal practice and whether traditional notions of thoroughness remain relevant in technology-mediated environments.

Acemoglu and Restrepo's research on the reinstatement effect provides theoretical grounding for understanding how automation frees up human labour for higher-order thinking that requires creativity and strategic judgement.⁵² The UK Law Society reinforces this view, emphasising that GenAI can enhance efficiency in routine work, enabling legal professionals to focus more intently on strategic considerations and client needs.⁵³ However, deeper examination reveals that this reinstatement may not be automatic: realising the benefits of GenAI integration requires deliberate restructuring of educational priorities and learning objectives.

Perhaps most significantly, GenAI technologies enable pedagogical approaches previously impractical within traditional legal education frameworks. Interactive tools for legal analysis can generate dialectical challenges to student analyses of legal provisions, producing counter-arguments drawn from historical jurisprudence that respond specifically to student reasoning rather than offering generic opposition. Liang's research explores how adaptive learning systems can analyse individual student performance across multiple dimensions of legal reasoning, creating an individualised curriculum that would be impossible to implement manually even with favourable faculty-to-student ratios.⁵⁴

The potential of personalised legal education deserves particular attention. Traditional legal education has operated through standardised curricula that assume relatively uniform student preparation and learning needs. GenAI-enabled personalisation may allow legal education to accommodate diverse learning styles and backgrounds while maintaining rigorous standards. Yet this potential benefit requires careful analysis of whether personalisation enhances educational quality or merely creates the illusion of accommodation while reinforcing existing inequalities through technological mediation.

Lv's recent work further demonstrates how GenAI-based teaching methods in legal education can analyse students' learning data to generate personalised learning paths, particularly enhancing legal reasoning through simulated scenarios and interactive learning.⁵⁵ These capabilities suggest the possibility of legal pedagogy that more effectively develops practical competencies

⁴⁹ Surden, "Artificial Intelligence and Law."

⁵⁰ Blue Hill Research, ROSS Intelligence Benchmark Report.

⁵¹ Schwarcz, "AI-Powered Lawyering."

⁵² Acemoglu, "Automation and New Tasks."

⁵³ Law Society of England and Wales, How AI is Reshaping.

⁵⁴ Liang, "Teaching Pattern."

⁵⁵ Lv, "AI-Powered Personalized Learning," 17–20.

while creating space for deeper engagement with theoretical frameworks. GenAI-driven simulations can prepare students for practical applications while ensuring understanding of ethical considerations, and personalised learning pathways can address the diverse needs of increasingly heterogeneous student populations without compromising intellectual rigour.

Arowosegbe et al.'s analysis indicates significant adoption of AI-augmented research tools across UK legal education, with widespread student adoption reported at 92 per cent according to industry surveys, indicating their perceived value among the emerging generation of legal professionals.⁵⁶ However, concentrated examination of student adoption patterns reveals that perceived value may not align with educational effectiveness; students may prefer GenAI tools for efficiency reasons while remaining unaware of their impact on skill development and critical thinking capacity. Yet these efficiency gains and promises of enhanced pedagogy constitute only one dimension of the dialectical tension, as the subsequent antithesis demonstrates how these very capabilities may simultaneously erode the analytical foundations they purport to strengthen.

7.2 Antithesis: The Erosion of Analytical Rigour

Despite the apparent benefits of GenAI integration in legal education, compelling evidence suggests that technological reliance may fundamentally undermine development of critical skills essential to legal practice. This antithetical position, championed by scholars concerned with preserving legal education's intellectual foundations, contends that the very efficiencies celebrated by GenAI proponents may simultaneously erode cognitive foundations upon which legal expertise is built, creating a generation of practitioners with diminished capacity for independent judgement and ethical reasoning.

The core concern centres on the phenomenon of cognitive offloading and its consequences for legal reasoning development. When students habitually delegate tasks to GenAI systems, they may fail to develop the neural pathways and cognitive structures that underpin sophisticated legal analysis. It is argued that repeated practice working through complex legal problems without technological shortcuts is how students develop the mental architecture necessary for independent judgement. This position suggests that efficiency gains may come at the cost of intellectual capacity-building.

Automation complacency represents a particularly concerning manifestation of this problem. Research demonstrates that law students working with GenAI research assistants displayed significantly reduced critical evaluation of outputs provided. In experimental conditions, students presented with GenAI-generated case analyses accepted flawed reasoning at rates 37 per cent higher than those working with traditional research methods.⁵⁷ This phenomenon echoes concerns about black box thinking in legal technology, wherein practitioners accept technological outputs without interrogating their validity or limitations.

The implications extend beyond immediate tasks to encompass the development of what can be termed legal imagination, the capacity to identify ethical dimensions of legal problems, recognise unstated assumptions in judicial reasoning and contextualise legal principles within broader social frameworks. These capabilities require sustained intellectual struggle with ambiguity and complexity, processes that may be short-circuited when GenAI systems provide immediate answers. The development of legal imagination depends on what educators call productive struggle, the cognitive effort required to work through complex problems that builds intellectual resilience and sophistication.

The counter-argument that efficiency does not equate with mastery finds particular support in studies of judicial bias recognition. Faster completion of case analysis tasks correlates with poorer identification of implicit judicial biases in precedential opinions. Marwala and Mpedi's cautionary research notes the importance of human oversight and verification to catch potential errors, emphasising that AI tools provide limited benefits for developing deeper critical evaluation required in complex legal analysis.⁵⁸

Comparative assessments of student understanding provide compelling evidence for these concerns. While students utilising GenAI tools demonstrated superior performance on direct application tasks, including 41 per cent faster document drafting and 92 per cent precision in precedent identification, they showed markedly weaker results on questions requiring transfer of legal principles to novel contexts or synthesis of disparate legal doctrines, under-performing by 31–38 per cent on cross-doctrinal synthesis and ethical reasoning assessments.⁵⁹

This suggests that although GenAI assistance helps students to identify correct answers, it may impede the development of conceptual frameworks necessary for independent application in unfamiliar scenarios. Research explains this discrepancy: the

⁵⁶ Arowosegbe, "Perception of Generative AI."

⁵⁷ Bommarito, "Analysis."

⁵⁸ Marwala, Artificial Intelligence.

⁵⁹ Schwarcz, "AI-Powered Lawyering."

cognitive processes activated during extended engagement with primary sources appear to facilitate cross-domain connections that are not developed through reviewing GenAI summaries, regardless of their accuracy. Empirical commentary indicates that students using traditional methods tend to form stronger cross-domain links in legal reasoning tasks, while GenAI users show weaker performance when applying precedents to novel fact patterns. This suggests that primary source engagement activates the neural pathways associated with schema formation, whereas GenAI outputs encourage surface-level pattern recognition.⁶⁰ The legal curriculum has traditionally valued the process of working through difficult problems as much as arriving at correct conclusions. The struggle itself is pedagogically valuable, building intellectual resilience and forcing students to confront the inherent ambiguity of legal reasoning. GenAI tools that reduce this productive struggle may inadvertently undermine development of these essential capacities, creating practitioners adept at using tools but lacking independent judgement when those tools are unavailable or insufficient. Additionally, GenAI systems trained on existing legal corpora tend to replicate historical disparities, inevitably reflecting and sometimes amplifying biases embedded in past legal reasoning. Studies indicate that students using GenAI tools are less likely to challenge outputs than when engaging with human sources, and many perceive GenAI outputs as more objective, illustrating a problematic authority effect of technological mediation.⁶¹

The opacity of GenAI decision-making processes compounds this concern. Unlike human reasoning, which can be interrogated through Socratic dialogue and explicit justification requirements, many GenAI systems function as black boxes whose internal operations remain inscrutable to users. Scholars argue that transparent reasoning processes are fundamental to legal education: students must not only reach correct conclusions but also be able to explain the reasoning that leads to them.⁶² GenAI systems that cannot provide such transparency risk undermining this core educational principle.

Beyond broad concerns about cognitive development, evidence demonstrates specific qualitative failures in GenAI legal reasoning that raise profound questions about its role in legal education. One of the most troubling is the misinterpretation of precedential hierarchies. In *Felicity Harber v HMRC* (2024), AI-generated submissions cited nine plausible but fictional First-Tier Tribunal decisions, incorrectly treating them as if they carried legitimate precedential weight.⁶³ The tribunal specifically noted these fabricated citations contained American spellings and repetitive phrasing inconsistent with UK judgments, illustrating how GenAI systems struggle with jurisdictional hierarchies and the relative weight of legal precedents.

Empirical work documents high legal-hallucination rates in case-law queries. In reference-based questions about US federal cases, Dahl et al. report that GPT-4 hallucinated legal facts in about 58 per cent of prompts and often accepted users' incorrect legal assumptions. These findings support caution about GenAI's treatment of authority and citations, without claiming a 75 per cent rate or a specific deficit in recognising overruled precedents.⁶⁴

Context-blind citation is another troubling pattern in GenAI legal analysis. In the *COPA v Wright* litigation, courts criticised filings that relied on AI-generated material and cited non-existent authorities; Wright's appeal bid was knocked back with costs after the court flagged 'AI-generated hallucinations'.⁶⁵ In *Lacey v State Farm* (2025), the court sanctioned attorneys for submitting briefs with serious GenAI-generated citation errors, warning that reliance on unverified AI output can mislead the court and undermine litigation standards. The order stressed the dangers of context-blind citation and overreliance on generative AI, though it did not directly address the deeper analytic limits of GenAI in assessing the rationale behind judicial precedent.⁶⁶ This problem also appears as context-blind citation: GenAI-assisted filings can include fabricated or misapplied authorities that look formally correct but collapse on scrutiny. In *Rotonde v Stewart Title Ins. Co.* (NY Sup. Ct. 2025), the court noted that several cases cited by a litigant 'do not exist'.⁶⁷ Together, these incidents illustrate how GenAI can produce plausible-sounding but legally defective citations unless rigorously verified.

The significance of these failures extends into questions about legal authority and knowledge validation. GenAI systems operate through statistical processes that generate plausible-sounding content without understanding legal authority structures or precedential hierarchies, revealing precisely those limitations that the Socratic-GenAI framework transforms into pedagogical opportunities for developing the discriminating judgement these failures expose as essential.

⁶⁰ Conley, "Legal Education's Role."

⁶¹ Bommarito, "AI Assistance."

⁶² Bliss, "Teaching Law," 112.

⁶³ *Felicity Harber v The Commissioners for HMRC*, UKFTT 1007 (TC), para 12

⁶⁴ Dahl, "Large Legal Fictions."

⁶⁵ *Crypto Open Patent Alliance v Wright*, EWHC 1198.

⁶⁶ *Lacey v State Farm General Insurance Co.*, 2025 WL 1363069 (C.D. Cal. 5 May 2025).

⁶⁷ *Rotonde v Stewart Title Ins. Co.*, 2025 NY Slip Op 50862(U) (NY Sup. Ct. 2025).

7.3. *Synthesis: Towards a Dialectical Pedagogy*

The examination of thesis and antithesis reveals not merely conflicting perspectives but complementary insights that, when properly synthesised, suggest more nuanced approaches to GenAI integration in legal education. This research proposes a framework wherein GenAI and human reasoning coexist not as competitors, but as partners in iterative feedback loops, each addressing limitations in the other while preserving essential developmental goals of legal education.

Analysis suggests that the quandary of GenAI in legal education is not its adoption – which appears inevitable – but how to design pedagogical frameworks that harness benefits while mitigating limitations. This approach reconceptualises GenAI not as a replacement for human reasoning but as a whetstone against which students can sharpen their capacities. However, whetstones work through resistance and friction, suggesting that effective GenAI integration may require deliberately introducing cognitive challenges rather than reducing them.

The central insight emerging from this dialectical analysis is that the most valuable pedagogical moments may occur precisely when GenAI systems fail in instructive ways. These failures become opportunities for developing human judgement rather than obstacles to technological integration. When GenAI systems struggle with ethical nuance or misinterpret precedential hierarchies, these shortcomings create occasions for students to exercise and refine precisely those skills that define sophisticated legal reasoning.

This insight suggests reconceptualising the relationship between efficiency and education: rather than viewing GenAI as a means to accelerate learning, educational frameworks might leverage GenAI's limitations as pedagogical resources that highlight uniquely human forms of reasoning and judgement.

This dialectical framework consists of three interconnected elements: structured contention, critical interrogation and epistemological transparency. Structured contention involves deliberately juxtaposing GenAI-generated analysis with human reasoning, creating productive tension that reveals limitations in both approaches while highlighting their complementary strengths. Critical interrogation establishes systematic protocols through which students question GenAI outputs, developing metacognitive awareness of both technological and human reasoning processes. Epistemological transparency requires making explicit the foundations of knowledge claims in both GenAI and human analysis, highlighting fundamental differences in how conclusions are reached and justified.

8. The Socratic-GenAI Framework: Augmentation Through Provocation

The synthesis suggests a specific framework for GenAI integration drawing from classical pedagogical approaches while incorporating contemporary technological capabilities. This Socratic-GenAI framework centres human development rather than technological efficiency as the primary objective of educational technology. However, this requires sophisticated understanding of both Socratic pedagogy and GenAI capabilities to avoid superficial implementations that fail to preserve essential educational values.

As mentioned previously, Sarkar's research on human-centred GenAI design provides a foundation for this framework, arguing that GenAI should function as a 'provocateur', challenging assumptions rather than simply optimising workflows, which aligns with Tan's 'cognitive partnership' model, where GenAI maintains dynamic models of users' reasoning patterns to identify conceptual blind spots.⁶⁸ This fundamentally shifts educational technology evaluation from efficiency metrics towards measures of intellectual engagement and reasoning quality.

This represents movement from instrumental towards formative understanding of educational technology. Rather than asking whether GenAI tools help students to complete tasks efficiently, the Socratic-GenAI framework asks whether technological interaction cultivates intellectual virtues and professional capacities essential for legal practice.

This framework reimagines GenAI as an intellectual provocateur stimulating deeper engagement with legal questions and not merely a labour-saving companion. Contemporary implementations demonstrate how GenAI can offer robust opposing perspectives, compelling students to strengthen positions through confrontation with contradiction, and operationalises 'structured dialectics' that mirror traditional Socratic questioning.

However, effective implementation requires more than technological capability: it demands sophisticated understanding of pedagogical timing, student developmental stages and disciplinary knowledge structures. GenAI provocateurs must challenge

⁶⁸ Tan, Partnership.

students appropriately for their expertise level while avoiding overwhelming novices or providing insufficient challenge for advanced students.

Modern implementations extend the Socratic method through technical innovations, including intent detection, where systems employ multi-agent imitation learning to model users' latent reasoning trajectories, enabling contextually appropriate interventions. Hypothetical generation allows GenAI frameworks to facilitate exploration of alternative problem-solving pathways, which is particularly valuable in legal education, where analysis of counterfactuals develops critical jurisprudential understanding.

Effective GenAI integration cannot rely on off-the-shelf tools; it requires sustained collaboration among legal educators, cognitive scientists and GenAI developers. The technical work of modelling reasoning trajectories is complex and often exceeds the capacity of a single institution. Research on trust calibration shows a paradox: participants value GenAI when it challenges superficial assumptions, yet resist it when it confronts deeply held beliefs or professional identities. This mirrors long-standing reservations about the Socratic method and calls for careful design of intervention timing, scaffolding and tone.⁶⁹

Consequently, GenAI transparency serves dual pedagogical purposes: enhancing utility while developing students' capacity to interrogate technological outputs. Transparent reasoning pathways must demonstrate how GenAI systems reach conclusions through techniques such as SHAP value visualisations, while bias-aware architectures employ adversarial debiasing modules flagging demographic skews in training data. This level of technological literacy may require fundamental restructuring of legal education curricula to include substantial components on computational thinking.

9. Implementation and Institutional Transformation

Practical implementation involves GenAI audit exercises where students systematically evaluate the reasoning of GenAI-generated legal analyses, identifying logical flaws while developing critical evaluation skills. Bliss's analysis reveals valuable pedagogical benefits from this framework, as it leverages GenAI's weaknesses into tools for teaching better legal reasoning by requiring students to compare, question and explain both human and machine logic.

Effective audit exercises therefore require sophisticated pedagogical design and substantial faculty development. Instructors must understand GenAI limitations to guide student analysis while avoiding either uncritical acceptance or reflexive rejection of technological outputs. Instructing students to correct GenAI-generated analyses creates deliberate practice in developing the reasoning skills that generative GenAI models presently lack; however, it demands fundamental changes in faculty preparation extending beyond traditional legal education training.

Building upon this foundation, adversarial collaboration demonstrates how pairing students with GenAI systems in analysing complex legal problems develops what scholars term 'meta-awareness', the ability to critically evaluate not just correctness of answers but the reasoning processes behind them. When students are tasked with identifying scenarios where GenAI recommendations would lead to suboptimal outcomes, they develop capacity for critical evaluation that may become essential professional competence rather than merely an academic exercise.

This metacognitive development is further strengthened by transparency protocols that ask students to state their reasoning alongside GenAI outputs and explicitly compare the epistemic basis of each. Research in legal education reflects that students who can articulate why they accept or reject GenAI recommendations exhibit the kind of independent judgement that legal education aims to foster, particularly through tasks requiring precedent interrogation, ethical justification and bias recognition.⁷⁰ For these benefits to materialise, students must understand not just that GenAI operates differently from humans, but why those differences profoundly impact legal analysis and professional decision-making.

The successful implementation of such dialectical pedagogy requires institutional commitments beyond individual classroom practices. Law schools must invest in faculty development programs that build the capacity of instructors to effectively integrate GenAI tools while maintaining focus on core learning objectives. King's College London's guidance policy affirms this position, noting that faculty who lack confidence in their technological understanding are unlikely to effectively guide students in critical GenAI engagement.⁷¹

⁶⁹ Sarkar, "AI Should Challenge"; Mertz, *The Language of Law School*.

⁷⁰ Paseri, "Examining Epistemological Challenges."

⁷¹ King's College London, *Macro-Level Guidance*.

This challenge extends beyond technical training to fundamental questions about professional identity and institutional mission. Faculty who view technology primarily as a threat may resist integration efforts, while those who embrace technology uncritically may fail to preserve essential educational objectives. Effective faculty development therefore requires sophisticated understanding of how technological integration can serve rather than subvert fundamental educational purposes.

Assessment design represents a particularly critical consideration intersecting with these faculty development challenges. This synthesis position advances the development of GenAI-resistant assessments that specifically evaluate the skills GenAI systems cannot replicate, such as normative reasoning, creative problem-solving and contextual application of legal principles that fundamentally challenge traditional approaches to curriculum planning and learning outcome specification that assume relatively stable knowledge and skill requirements. The dynamic nature of GenAI capabilities, however, informs that this assessment design must be adaptive rather than static, requiring institutional commitment to ongoing professional development that preserves educational integrity while keeping pace with technological advancement. These implementation strategies operationalise the Socratic-GenAI framework's core principle: transforming GenAI's limitations from obstacles into pedagogical resources that sharpen rather than replace students' capacity for independent legal reasoning.

10. Ethical Considerations and Future Horizons

The dialectical integration of GenAI demands explicit ethical guardrails, ensuring that technological adoption serves rather than subverts educational objectives. These ethical considerations extend beyond implementation concerns to fundamental questions about professional identity, social responsibility and legal education's role in promoting justice.

Responsible GenAI integration requires transparency regarding training data. However, transparency requirements conflict with commercial interests and competitive pressures that limit developers' willingness to disclose proprietary information. This tension suggests that ethical GenAI integration may require regulatory intervention or collective action by educational institutions to establish industry standards that prioritise educational needs over commercial interests.

Developing appropriate guardrails requires collaboration across multiple stakeholders. Law schools must assume primary responsibility for curriculum design and assessment protocols that incorporate GenAI while maintaining educational integrity. Professional associations must then address how GenAI integration affects qualification standards, while GenAI developers must incorporate the unique requirements of legal education in the development of these services.

Effective collaboration demands institutional structures and incentive systems that align stakeholder interests with educational objectives. Current market dynamics incentivise rapid AI deployment over careful consideration of educational impacts, creating tensions between commercial success and educational responsibility requiring deliberate institutional responses. The cultivation of legal expertise faces particular challenges that demand attention. If legal expertise requires sustained intellectual labour involving struggles with complexity and ambiguity, GenAI systems that reduce these struggles may inadvertently undermine expertise development while improving immediate performance.

Strategic adaptation requires pedagogical approaches to evolve from treating GenAI as external tools to conceptualising it as an integral component of the modern legal ecosystem requiring critical evaluation and judicious application. This evolution demands fundamental reconceptualisation of legal education objectives, extending far beyond mere tool adoption. Promising curricular innovations include mandatory GenAI literacy modules equipping students to understand methodological limitations while appreciating how technological constraints interact with legal reasoning requirements and professional responsibilities. These ethical frameworks and curricular innovations operationalise the Socratic-GenAI principle that GenAI must be positioned as a whetstone sharpening students' capacity for ethical reasoning rather than a substitute for the normative judgement that distinguishes excellent legal minds from technically proficient operators.

11. Beyond Binary Thinking: Toward Genuine Synthesis

The ultimate goal transcends simplistic dichotomies of technological embrace or rejection, instead seeking sophisticated navigation between human judgement and artificial assistance. The legal mind of tomorrow will be defined not by resistance to technological change but by sophisticated integration of human judgement with GenAI capabilities.

This requires fundamental reconceptualisation of professional competence and educational objectives that acknowledges both technological capabilities and human limitations while preserving essential values. Legal education must cultivate minds capable of both technological sophistication and humanistic wisdom, recognising that constitutional interpretation ultimately concerns human values that resist reduction.

The challenge lies in operationalising these aspirations through concrete educational practices and institutional structures. The most promising aspect of this synthesis is its potential to create virtuous cycles of development where students' sophisticated critiques of GenAI outputs inform improvements in GenAI systems, while increasingly capable GenAI challenges students to develop more nuanced approaches.

Realising this potential requires educational institutions to become active participants in GenAI development rather than passive consumers, representing not merely a technological challenge but an opportunity for educational innovation serving both technological advancement and human flourishing.

12. Conclusion

The dialectical examination of GenAI integration in legal education reveals that the fundamental challenge lies not in GenAI adoption, which appears inevitable, but in designing pedagogical frameworks that harness benefits while mitigating limitations through sophisticated understanding of both technological capabilities and educational objectives. The Socratic-GenAI paradigm offers a promising path forward: positioning GenAI as an intellectual provocateur rather than a replacement for human reasoning, thereby creating productive tension that strengthens rather than weakens capabilities.

This research demonstrates that effective GenAI integration requires coordinated institutional action across policy, curriculum, infrastructure and pedagogy domains. Piecemeal approaches addressing only technical implementation without corresponding attention to educational philosophy and assessment integrity risk undermining rather than enhancing legal education's fundamental objectives. The synthesis proposed herein rejects both uncritical techno-optimism and reactionary resistance, seeking transformative integration that enhances core intellectual virtues of legal education while leveraging the practical efficiencies that GenAI offers.

The writers conclude with a call for educators, policy-makers and technologists to collaboratively codify the Socratic GenAI paradigm, establishing frameworks that harness technological potential while safeguarding essential human dimensions of legal reasoning. The future of legal education depends not on choosing between tradition and innovation but on their thoughtful synthesis.

The Socratic-GenAI framework provides a roadmap for cultivating jurists capable of sophisticated navigation between human judgement and artificial assistance, establishing principles that can guide responsible technological integration while safeguarding the intellectual traditions that have long distinguished legal education. Through thoughtful synthesis of tradition and innovation, legal pedagogy can evolve without sacrificing its essential character, ensuring that future practitioners possess both technological fluency and the deep capabilities that remain at the heart of legal reasoning.

Bibliography

- Acemoglu, Daron and Pascual Restrepo. "Automation and New Tasks: How Technology Displaces and Reinstates Labor." *Journal of Economic Perspectives* 33, no 2 (2019): 3–30. <https://doi.org/10.1257/jep.33.2.3>.
- Adam, Taskeen. "Digital Neocolonialism and Massive Open Online Courses (MOOCs): Colonial Pasts and Neoliberal Futures." *Learning, Media and Technology* 44, no 3 (2019): 365–380. <https://doi.org/10.1080/17439884.2019.1627893>.
- American Bar Association. "AI in Legal Education." ABA Center for Innovation, 2024. https://www.americanbar.org/groups/center-for_innovation/resources/publications/ai-in-legal-education.
- American Bar Association, Task Force on Law and AI in Legal Education. "Artificial Intelligence, Legal Education & Assessment." 2024. https://www.americanbar.org/groups/legal_education/resources/ai-task-force.
- Arowosegbe, Abayomi, Jaber S. Alqahtani and Tope Oyelade. "Perception of Generative AI Use in UK Higher Education." *Frontiers in Education* 9 (2024). <https://doi.org/10.3389/feduc.2024.1463208/full>.
- Ashley, Kevin D. *Artificial Intelligence and Legal Analytics: New Tools for Law Practice in the Digital Age*. Cambridge: Cambridge University Press, 2017.
- Bisbee, James, Will Fairbrother, Johanna Lee and Adam Howat. "Performance and Biases of Large Language Models in Public Opinion Simulation." *Palgrave Communications* 10, no 1 (2024). <https://doi.org/10.1057/s41599-024-03609-x>.
- Blackham, Alysia. "Interrogating New Methods in Socio-Legal Studies: Content Analysis, Case Law and Artificial Intelligence." *Alternative Law Journal* 50, no 2 (2025): 91–98. <https://doi.org/10.1177/1037969X251325869>.
- Bliss, John. "Teaching Law in the Age of Generative AI." *Jurimetrics* 64, no 2 (2024): 111–161. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4682456.
- Blue Hill Research. *ROSS Intelligence Benchmark Report*. 2024. <https://bluehillresearch.com/ross-intelligence-benchmark-report>.
- Borah, Angana and Rada Mihalcea. "Towards Implicit Bias Detection and Mitigation in Multi-Agent LLM Interactions." Association for Computational Linguistics: EMNLP (2024): 8657–8675. <https://arxiv.org/abs/2410.02584>.
- Charlotin, Damien. *AI Hallucination Cases Database*. <https://www.damiencharlotin.com/hallucinations>.
- Choi, Jonathan H., Amy B. Monahan and Daniel Schwarcz. "Lawyering in the Age of Artificial Intelligence." *Minnesota Law Review* 109, no 1 (2024): 147–218. <https://doi.org/10.2139/ssrn.4626276>.
- Choi, Jonathan H. and Daniel Schwarcz. "AI Assistance in Legal Analysis: An Empirical Study." *Journal of Legal Education* 73, no 2 (2023): 256–281. <https://doi.org/10.2139/ssrn.4539836>.
- Chouldechova, Alexandra. "Fair Prediction with Disparate Impact: A Study of Bias in Recidivism Prediction Instruments." *Big Data* 5, no 2 (2017): 153–163. <https://doi.org/10.1089/big.2016.0047>.
- de la Osa, David Uriel Socol and Nydia Remolina. "Artificial Intelligence at the Bench: Legal and Ethical Challenges of Informing – or Misinforming – Judicial Decision-Making Through Generative AI." *Data & Policy* 6 (2024): E59. <https://doi.org/10.1017/dap.2024.53>.
- Engstrom, David F., Daniel E. Ho, Julian Nyarko, Nora Freeman Engstrom and Elizabeth Joh. *Hallucination-Free? Assessing the Reliability of Leading AI Legal Research Tools*. Preprint, Stanford Law School, 2024. <https://law.stanford.edu/publications/hallucination-free-assessing-the-reliability-of-leading-ai-legal-research-tools>.
- Fricker, Miranda. *Epistemic Injustice: Power and the Ethics of Knowing*. Oxford: Oxford University Press, 2007.
- Gallegos, Isabel O., Ryan A. Rossi, Joe Barrow, Md Mehrab Tanjim, Sungchul Kim, Franck Dernoncourt, Tong Yu, Ruiyi Zhang and Nesreen K. Ahmed. "Bias and Fairness in Large Language Models: A Survey." *Computational Linguistics* 50, no 3 (2024): 1097–1179. https://doi.org/10.1162/coli_a_00524.
- Geva, Mor, Tomer Wolf, Amir Feder, Daniel Hershcovich, Yoav Goldberg and Yonatan Belinkov. "Do LLMs Exhibit Human-Like Cognitive Biases? A Large-Scale Systematic Evaluation." Preprint, *arXiv*, 2024. <https://arxiv.org/abs/2404.20085>.
- Graben, Sari. "Law and Technology in Legal Education: A Systemic Approach at Ryerson." *Osgoode Hall Law Journal* 58, no 1 (2021): 139–162. <https://digitalcommons.osgoode.yorku.ca/ohlj/vol58/iss1/4>.
- Ho, Yi-Jen (Ian), Wael Jabr and Yifan Zhang. "AI Enforcement: Examining the Impact of AI on Judicial Fairness and Public Safety." *SSRN Electronic Journal* (2024). https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4533047.
- DeCoito, Isha and Tasha Richardson. "Teachers and Technology: Present Practice and Future Directions." *Contemporary Issues in Technology and Teacher Education* 18, no 2 (2018): 362–383. <https://doi.org/10.1080/21532974.2019.1568327>.
- Kleinberg, Jon, Sendhil Mullainathan and Manish Raghavan. "Inherent Trade-Offs in the Fair Determination of Risk Scores." In *Proceedings of the 8th Innovations in Theoretical Computer Science Conference (ITCS)*, 2017. <https://arxiv.org/abs/1609.05807>.
- King's College London. "Macro-Level Guidance: University-Wide Principles and Policy for Generative AI in Teaching, Assessment, and Feedback." King's College London, 2025. <https://www.kcl.ac.uk/about/strategy/learning-and-teaching/ai-guidance>.

- Law Society of England and Wales. "AI Strategy: The Three Is – Innovation, Impact, Integrity." Law Society media release, 16 October 2024. <https://www.lawsociety.org.uk/en/contact-or-visit-us/press-office/press-releases/law-society-new-ai-strategy-the-three-is>.
- Law Society of England and Wales. "How AI is Reshaping the Future of Legal Practice." 19 November 2024. <https://www.lawsociety.org.uk/en/topics/ai-and-lawtech/partner-content/how-ai-is-reshaping-the-future-of-legal-practice>.
- Liang, Shi, Zhu Zhu and Xu Xuan. "The Teaching Pattern of Law Majors Using Artificial Intelligence." *Frontiers in Psychology* 12 (2021): 711520. <https://doi.org/10.3389/fpsyg.2021.711520>.
- Loevinger, Lee. "Jurimetrics – the Next Step Forward." *Minnesota Law Review* 33, no 5 (1949): 455–493. <https://scholarship.law.umn.edu/mlr/1796>.
- Lv, Xinyi and Zhiyong Xu. "AI-Powered Personalized Learning in Legal Education: A Tool for Developing Future Ready Lawyers." *International Journal of Law, Justice and Jurisprudence* 5, no 1 (2025): 323–327. <https://www.lawjournal.info/archives/2025.v5.i1.D.192>.
- Marwala, Tshilidzi and Letlhokwa George Mpedi. *Artificial Intelligence and the Law*. Singapore: Springer, 2024.
- McDonald, Nora, Aditya Johri, Areej Ali, and Aayushi Hingle. "Generative Artificial Intelligence in Higher Education: Evidence from an Analysis of Institutional Policies and Guidelines." arXiv preprint (2024). <https://arxiv.org/abs/2402.01659>.
- MIT Media Lab. *Your Brain on ChatGPT: Accumulation of Cognitive Debt and the Decline in Critical Thinking*. MIT Media Lab White Paper, June 2025. <https://www.media.mit.edu/publications/your-brain-on-chatgpt>.
- Moshman, David. "Exogenous, Endogenous, and Dialectical Constructivism." *Developmental Review* 2, no 4 (1982): 371–384. [https://doi.org/10.1016/0273-2297\(82\)90019-3](https://doi.org/10.1016/0273-2297(82)90019-3).
- Noble, Safiya Umoja. *Algorithms of Oppression: How Search Engines Reinforce Racism*. New York: New York University Press, 2018.
- Nussbaum, Martha C. *Not for Profit: Why Democracy Needs the Humanities*. Princeton, NJ: Princeton University Press, 2010.
- Paseri, Ludovica and Massimo Durante. "Examining Epistemological Challenges of Large Language Models in Law." *Cambridge Forum on AI Law and Governance* 1, no 1 (2025): 22–36 <https://doi.org/10.1017/cfl.2024.7>.
- Sargeant, Holli and Måns Magnusson. "Bias in Legal Data for Generative AI." Generative AI & Law Workshop at ICML (2024). https://blog.genlaw.org/pdfs/genlaw_icml2024/9.pdf.
- Sarkar, Advait. "AI Should Challenge, Not Obey." *Communications of the ACM* 67, no 10 (2024): 18–21. <https://doi.org/10.1145/3649404>.
- Schrepel, Thibault. "Generative AI in Legal Education: A Two-Year Experiment with ChatGPT." *Law, Innovation and Technology* (forthcoming 2026). https://papers.ssrn.com/sol3/papers.cfm?abstract_id=5401422.
- Schwarcz, Daniel and Jonathan H. Choi. "AI Tools for Lawyers: A Practical Guide." *Minnesota Law Review* 108 (2023): 1–58. <https://minnesotalawreview.org/wp-content/uploads/2023/10/FL1-Choi-Schwarcz.pdf>.
- Schwarcz, Daniel, Sam Manning, Patrick Barry, David R. Cleveland, J.J. Prescott and Beverly Rich. "AI-Powered Lawyering: AI Reasoning Models, Retrieval Augmented Generation, and the Future of Legal Practice." *SSRN Electronic Journal* (2025). https://papers.ssrn.com/sol3/papers.cfm?abstract_id=5162111.
- Smith, Craig. "How Generative AI Is Changing the Legal Profession – and What We Can Do About It." *The Conversation*, March 27, 2024. <https://www.salford.ac.uk/news/the-conversation-how-generative-ai-is-changing-the-legal-profession-and-what-we-can-do-about-it>.
- Smith, Patricia. *Reimagining Balanced Assessment Systems*. National Academy of Education, 2024. <https://files.eric.ed.gov/fulltext/ED661489.pdf>.
- Spark Emerging Technologies. "Legal Training VR: Revolutionising Learning Experiences." Spark Emerging Technologies Blog, 8 August 2024. <https://sparkemtech.co.uk/blog/legal-training-vr-revolutionising-learning-experiences>.
- Surden, Harry. "Artificial Intelligence and Law: An Overview." *Georgia State University Law Review* 35 (2019): 1305. <https://scholar.law.colorado.edu/faculty-articles/1234>.
- Tan, Emmanuel. "Partnership in the Age of Generative Artificial Intelligence." *Teaching and Learning Together in Higher Education* 1, no 39 (2023): Article 3. <https://repository.brynmawr.edu/tlthe/vol1/iss39/3>.
- Tao, Yan, Olga Viberg, Ryan S. Baker and René F. Kizilcec. "Cultural Bias and Cultural Alignment of Large Language Models." *PNAS Nexus* 3, no 9 (2024): 346. <https://doi.org/10.1093/pnasnexus/pgae346>.
- UKAI. "Top UK Law Firms Accelerate AI Adoption as Client Demand Grows." *UKAI News*, 23 September 2025. <https://ukai.co/top-uk-law-firms-accelerate-ai-adoption-as-client-demand-grows>.
- University College London. "AI and Law: UCL Laws Summer School." 2025. <https://www.ucl.ac.uk/laws/short-courses/summer-school/ai-and-law>.
- University of Wollongong. "What Happens When Law Students Go Head-to-Head with GenAI?" 11 September 2025. <https://www.uow.edu.au/media/2024/what-happens-when-law-students-go-head-to-head-with-genai.php>.
- Veale, Michael, Isra Black, Mark Dsouza, Matthew Fisher, Mujtaba Ghaus, Thea Gibbs, Orla Lynskey, Colm O’Cinneide, Eloise Scotford, Olga Thomas, Alina Trapova and Kimberley Trapp. "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, 2025.

<https://discovery.ucl.ac.uk/id/eprint/10208136/1/AI,%20Education%20and%20Assessment%20at%20UCL%20Laws.pdf>.

Williams, Kate. “What is the Digital Divide?” *Journal of Social Informatics* 5, no 3 (2001): 25–33.

<https://katewill.web.ischool.illinois.edu/work/31-williams-2001-what-is-the-digital-divide.pdf>.

Legal materials

Crypto Open Patent Alliance v Wright, EWHC 1198.

Felicity Harber v The Commissioners for HMRC, UKFTT 1007 (TC), para 12

Lacey v State Farm General Insurance Co., 2025 WL 1363069 (C.D. Cal. 5 May 2025).

Rotonde v Stewart Title Ins. Co., 2025 NY Slip Op 50862(U) (NY Sup. Ct. 2025).