



Deposited via The University of Sheffield.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/id/eprint/243243/>

Version: Published Version

Article:

Jones, E.J. (2026) Holism in a digital age: Bildung, generative Artificial Intelligence and goal alignment in legal education. Law and Method. ISSN: 2352-7927

<https://doi.org/10.5553/LaM.2026.17>

Reuse

This article is distributed under the terms of the Creative Commons Attribution-NonCommercial (CC BY-NC) licence. This licence allows you to remix, tweak, and build upon this work non-commercially, and any new works must also acknowledge the authors and be non-commercial. You don't have to license any derivative works on the same terms. More information and the full terms of the licence here:

<https://creativecommons.org/licenses/>

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.

ARTICLE

Holism in a Digital Age: *Bildung*, Generative Artificial Intelligence and Goal Alignment in Legal Education

*Special Issue Bildung in Academic Legal Education, Hedwig van Rossum & Bald de Vries (eds.)**

Emma Jane Jones**

Abstract

*The introduction of generative AI ('genAI') offers significant challenges and opportunities for legal education. However, to date, discussion has focused largely on the levels of error ('hallucinations') found within genAI-generated results, the potential for academic misconduct through plagiarism when students rely on genAI in assessment writing and (in more limited cases) its applications as a teaching tool. This article takes a novel approach by drawing on the field of AI safety to explore the wider issue of goal alignment, considering what goals genAI systems currently have within legal education and the normative goals which should shape future genAI use in this field. It argues that it is vital for legal educators to engage with this issue to ensure that genAI is used, and continues to be used, in a way which supports the overall objectives of legal education and avoids narrow and reductionist approaches. To facilitate this, the article proposes that the philosophical tradition of *bildung* be adapted as a conceptual framework to guide the determination of the goals shaping genAI's future development and applications within legal education.*

Keywords: artificial intelligence, *bildung*, goal alignment, legal education, normative frameworks.

1. Introduction: Legal Education in a Digital Age

To ensure that AI systems have a positive impact, it is not enough to focus on mitigating risks of harmful behaviour. Rather, we need to instill beneficial

* The author would like to thank Professor Anthony Bradney and Dr Gareth Bramley for their valuable feedback on earlier drafts of this work. Any errors remain entirely the author's own.

** Dr. Emma Jones is Senior Lecturer in Law and Director of Student Wellbeing for the School of Law at Sheffield University, UK.

goals in AIs and ensure that these align with our values and ethical principles. (Hendrycks, 2025, Chapter 6.1)

In November 2021, the launch of OpenAI's ChatGPT model heralded the introduction of generative artificial intelligence (genAI) as a widely available form of digital technology.¹ However, the concept of artificial intelligence (AI) has existed for much longer, for example, Turing's seminal 1950 article discussed the question 'Can machines think?' (Turing, 1950) and the term itself was captured at a 1956 conference at Dartmouth College, New Hampshire (Haenlein & Kaplan, 2019). Definitions of AI vary, but it can broadly be defined as 'systems that display intelligent behaviour by analysing their environment and taking actions – with some degree of autonomy – to achieve specific goals' (European Commission, 2019, Sheikh et al., 2023). GenAI has similar diffuse meanings, with differences between technical and public understanding (Ronge et al., 2025). However, key features include its '(multi-) modality' (use of text, images and sound), high levels of user 'interaction', the 'flexibility' of one system able to accomplish a range of tasks and its ability to enhance 'productivity' (Ronge et al., 2025, p. 32). These features are commonly underpinned by a Large Language Model ('LLM') which is trained on vast datasets (known as Big Data) to enable the system to predict the appropriate next word, generating an output (Surden, 2023).

Since the introduction of ChatGPT and popularization of genAI, a myriad of books, articles and social media posts have researched, theorized, empirically tested, debated and speculated on its consequences for legal education (see, for example, Ajevski et al., 2023; Ma & Hou, 2021). ChatGPT's launch has been a catalyst for urgent discussion around topics such as assessment and academic integrity, with case studies of the integration of genAI into assessment formats or, conversely, the 'AI-proofing' of assessments becoming increasingly widespread (see, for example, Veale et al., 2025). A series of seemingly AI-generated errors arising within legal proceedings has also highlighted the presence of so-called 'hallucinations' in the outputs of law-related genAI, such as fictitious case citations (Dahl et al., 2024). There have also been some innovative attempts to adopt genAI as a teaching tool within law schools (see Ryan, 2020; Schrepel, 2025). Each of these aspects of genAI is relevant and important. However, much of the discussion of such practical and pedagogic implications fails to explicitly consider broader issues of AI safety.

The field of 'AI safety' refers to work 'aimed at preventing or mitigating harms from AI systems' development and deployment' (Harding & Kirk-Giannini, 2025, p. 1518). With AI becoming increasingly powerful and the realistic prospect of human-level AI (commonly termed Artificial General Intelligence (or AGI)) being reached within the next decade, this field is vital to ensure that AI (including genAI) can be harnessed in beneficial ways while avoiding both catastrophic incidence and the myriad of lesser but still significant harms which it can generate (on GAI, see Chandra Bikkasani, 2025). As genAI becomes increasingly embedded in learning and teaching environments, including law schools, this article asks what the

1 <https://chatgpt.com/>.

current goals of genAI are in the field of legal education, whether they are sufficient and appropriate and, if not, how they can be reconfigured. It argues that it is vital to explicitly apply AI safety concepts to enable the consideration of broader harms that may emerge, beyond issues of academic misconduct and AI-generated errors. It surfaces largely undiscussed issues relating to goal alignment which could have significant impacts on the future of legal education. Overall, this article argues that the current goals of genAI in legal education are overly narrow, instrumental and reductivist. In contrast, it proposes utilizing the concept of *bildung* as an approach which better aligns with the wider objectives, richness and value of legal education.

The concept of goal alignment is a core aspect of AI safety (Johnsen, 2024). In a narrow sense, it can refer to whether an AI system meets the goal(s) of its creators and operators (Gabriel & Ghazavi, 2022). More broadly, it refers to the need for AI systems to be in line with human values and/or preferences, so that such systems contribute beneficially and usefully to human existence (Kirk et al., 2024). This latter notion involves technical aspects, to ensure that a system is ‘tuned’ appropriately (Firt, 2023). However, it also involves normative aspects, in deciding which values or preferences should be enshrined within the goals of AI systems (Gabriel, 2020). At its broadest, the normative process of goal alignment draws on deeply philosophical questions around preferred values systems and how these should be selected and applied when developing AI systems (Gabriel, 2020). This article will focus on such normative questions, albeit limited to the specific context of legal education.

The article makes a novel contribution to the literature on genAI in legal education by highlighting and exploring the neglected issues around AI safety in this arena. It explores the goals currently driving the development of genAI use cases within legal education, identifying that these typically focus on speed, accuracy and productivity. It argues that these goals are insufficient to align with the overall objectives of legal education and are vulnerable to forms of proxy gaming by AI models in a way which could be detrimental to the accomplishment of such objectives. To proactively address these issues, the article proposes that the educational philosophy of *bildung* be used as a conceptual framework which can be applied to the future development of genAI in legal education, guiding the formation of suitable goals and ensuring their appropriate alignment with the broader objectives of legal education.

2. GenAI and Goal Alignment

Before moving on to consider *bildung* and its application to genAI in legal education, it is important to understand the relevance of goals and goal alignment within genAI systems generally. Here, the notion of a ‘goal’ refers to the objective, or set of objectives, which an AI system is designed to achieve, which in turn will be determined by the interests, preferences, desires, morals and values of its creator or users (Shen et al., 2024). At its broadest, AI can be seen as seeking to achieve, or even surpass, human intelligence (Bignami, 2022). However, in the contemporary age of narrow or weak AI, the goals of individual AI systems will often be dependent

on the specific context in which they are being utilized, from safely driving a car to winning at chess (Babu & Banana, 2024).

Within the AI safety literature, the ‘alignment problem’ is recognized as a significant issue in relation to genAI (Ngo et al., 2025). It refers to the need to ensure ‘that AI systems pursue goals that match human values or interests rather than unintended and undesirable goals’ and the difficulties in achieving this both technically and normatively (Dahl et al., 2024). GenAI models such as ChatGPT (sometimes referred to as general purpose AI models) are built for multiple uses and may perform tasks they have not been specifically trained for (Triguero et al., 2024). They will therefore have broad goals; for example, OpenAI’s GPT-5 chatbot advertises itself by stating ‘Get answers. Find inspiration. Be more productive’ (ChatGPT, n.d.). Conversely, others are domain-specific, more narrowly customized to the needs of a particular target market or user group, meaning that they will have specific goals relating to that market which they are designed to achieve (Uuk et al., 2025). One example of this is the legal AI assistant Harvey AI (Harvey, n.d.). Its website states that it is ‘designed for legal and professional services’ to allow lawyers to ‘elevate their craft’, ‘navigate complexity’ and ‘focus on high-value work’ (Harvey, n.d.).

The nature of genAI means that it is not always easy for a system’s specific goals to be achieved, for example, in the well-known example of Microsoft’s chatbot Tay, it was intended to entertain 18 to 24-year-olds in the US (Lee, 2016). However, it was swiftly influenced into expressing extreme and highly offensive views (Lee, 2016). Often, this is due to the Big Data which the LLMs have been trained on, which reflects contemporary and historical biases within society (see, for example, An et al., 2025; Bolukbasi et al., 2016). A related concept to goal alignment is that of specification gaming, where an AI system complies literally with the specified goal, but does not meet its intended outcome, as a result of a misspecification of the goal itself which provides loopholes to be exploited (Krakovna et al., 2020; Tlaie, 2024). Krakovna et al. give a range of examples of this phenomenon, including several where the intended outcome is for an AI agent to win a game, but it is not explicitly stated that this has to be while following the rules of the game, leading to ingenious forms of cheating (Krakovna et al., 2020).

The concepts of goal alignment and specification gaming demonstrate that it is important to thoroughly consider the goals of any AI system, identifying and defining them with care to (as far as is possible) avoid misalignment and specification gaming. Applying this to legal education, it demonstrates the importance of considering what goals genAI systems in the legal sphere are aligned to and how appropriate these are in this context.

As discussed previously, an AI system’s goals can be very broad. This can potentially lead to them being vague and ambiguous or open to a wide range of different interpretations (specification gaming). It is therefore common for AI developers to identify proxy goals when developing an AI model. In other words, goals which are clear and achievable in a way which allows their attainment (or otherwise) to be measured and evaluated (Hendrycks, 2025). This results in a system being attuned to quantifiable goals which are aligned with wider aims. For example, an AI model could have ‘promoting justice’ as its overall aim. However,

‘justice’ is a huge and complex concept which it would be virtually impossible to metricize and measure as a stand-alone concept. Therefore, the developers of the AI model may seek to set one or more proxy goals which can be used to measure and evaluate the system’s progress towards its overall aim. This could be a goal such as ‘providing legal advice to people who are experiencing a family law issue and cannot afford legal advice’. Such a proxy goal defines and narrows the parameters of the AI system and enables its utility to be assessed. At the same time, setting such proxy goals can be problematic; for example, the system may meet the proxy goal by providing a thousand users with legal advice, but it may not meet the overall aim if that legal advice is inaccurate and reliance on it leads to users not achieving just outcomes in cases. This example describes the notion of proxy gaming, where ‘a model optimizes for a proxy in a way that produces undesirable or even harmful outcomes as judged from the idealized goal’ (Hendrycks, 2025, Chapter 3.3.2). Such issues often arise as pressures to optimize performance increase (Hendrycks, 2025).

Therefore, alongside considering the goals of AI systems in the legal sphere, it is necessary to consider whether there are proxy goals which are governing these systems’ development and use and ask whether issues of proxy gaming may arise.

3. GenAI in Legal Education

At present, the empirical evidence on how genAI is used in legal education is relatively sparse. Publicly available versions of LLMs, such as ChatGPT and Google Gemini, are commonly accessed by students across higher education, with a 2024 global survey indicating that 86% of students were regularly using AI (Digital Education Council, 2024). There is currently nothing to suggest that law students deviate significantly from the general student population in their level of usage. There are also a number of case studies of legal educators incorporating its use within teaching and learning, sometimes encouraging students to access such systems (see, for example, Serra, 2025). A small-scale survey of 54 predominantly UK-based law students indicated that 73% were using AI for study and/or work, with 65% of those using a free version of chat GPT (Andreeva & Savova, 2024; see also Buhari Bello et al., 2025). This accumulated data provides some evidence that general purpose AI systems are being utilized by law students to assist them in legal research, writing and other aspects of study. Such chatbots have wide-reaching stated goals, for example, as mentioned previously, ChatGPT states that it ‘helps you get answers, find inspiration and be more productive’.² Google Gemini states ‘Get help with writing, planning, brainstorming, and more’.³ Google’s Notebook LLM, seemingly targeted specifically at students and researchers, advertises itself as ‘your research and thinking partner’ which will help you to ‘learn faster and deeper’, ‘present with confidence’ and ‘unlock your creative potential’ (Google, 2025). Microsoft Copilot states that it is ‘a conversational, AI-powered assistant

2 Words taken from the web search preview of <https://chatgpt.com> (retrieved 15 September 2025).

3 Words taken from web search preview of <https://gemini.google.com/app> (retrieved 15 September 2025).

that helps boost productivity and streamline workflows by offering contextual assistance, automating routine tasks and analysing data' (Microsoft, 2025). While universities and law schools may enter into agreements to provide law students with institutionally funded licensed versions, potentially offering greater guardrails in terms of data exploitation, these appear to retain the same goals as the publicly available systems, with the differences lying in more restricted uses of data that is inputted (to avoid it being used for training), better data protection, more support and potentially greater citation of source materials (Webb, 2024).

There are also an increasing number of law-specific genAI systems, often developed by existing legal publishers, which may be adopted within law schools. For example, Thomson Reuters' Co-Counsel Legal, launched in January 2025, which promises accuracy, speed and security in legal research, analysis and drafting (Thomson Reuters, 2025). Lexis+ AI from LexisNexis also promises similar benefits, referring to it as 'so reliable, fast, and well-informed, it feels like your protégé' (LexisNexis, 2025). Harvey AI, which has recently announced a UK expansion from law firms to law schools (Harvey Team, 2025), refers to 'Rapid research. Grounded results' (Harvey, 2025). More broadly, while not yet specifically aimed at legal education, Anthropic's Claude for Legal suite of plugins has the tagline 'Speed up contract reviews, NDA triage, and compliance workflows' (Claude, 2026).

What appears to be common across general and law-specific systems is the focus on speed, often linking it to productivity and efficiency with the seeming message of 'work faster, do more'. Perhaps unsurprisingly, the law-specific LLMs also focus on levels of accuracy. A number of high-profile legal cases involving fictitious case citations, seemingly produced via genAI have clearly demonstrated the pitfalls of so-called 'hallucinations' in AI outputs within legal settings.⁴ Therefore, to build trust in such systems, it is inevitable that their accuracy will be emphasized. One distinction is that the general purpose AI systems also include some references to creativity, something notably missing from the law-specific variants. However, given the ubiquity of speed and accuracy as goals of systems likely to be commonly used by law students, these appear to be the key aspects to consider.

The question then arises as to whether speed and accuracy are the actual overall goals of genAI systems or whether they are proxy goals which have been introduced to help measure the achievement of broader overall goals. In fact, general AI systems do appear to have such broader (if somewhat nebulous) stated goals; for example, OpenAI (creator of ChatGPT) states that its mission is to 'ensure AI benefits all of humanity' (OpenAI, 2025). Google states that Google Gemini's aim is to be 'the most helpful AI assistant', performing a range of roles from 'personal assistant' to 'coder' and 'researcher' (Gemini, n.d.). Microsoft refers to AI as 'transforming how people work and helping organizations operate more efficiently' (Microsoft, 2025).

When considering law-specific genAI systems, it could be argued that the overall aim of these is to in some way produce 'better' law students or lawyers. For

4 See, for example, the UK case of *Ayinde v. London Borough of Haringey and Al-Haroun v. Qatar National Bank* [2025] EWHC 1383 (Admin).

example, in a recent blog post publicizing Harvey AI's partnership with four UK law schools, the quotations from legal academics involved in the adoption seem to focus largely on its potential for fostering AI literacy among students and staff, with references to developing practical skills and ethical insights (Harvey Team, 2025). This could be read in tandem with several of the system providers' websites, which seem to focus on the potentially transformative impact of AI on the legal profession (see, for example, Thomson Reuters, 2025). It could be that they are seeking to engage law students with AI to provide them with the skills to use it on their (potential) entry into the legal profession. However, the end goal of such transformation still appears to focus largely on speed and accuracy – freeing time for lawyers and law firms to maximize profits and focus on higher-value tasks (Thomson Reuters, 2025).

The remainder of this article will focus on law-specific systems of genAI, while suggesting that similar issues with goal alignment also arise with broader non-law-specific genAI models. When considering law-specific models, there are two alternative perspectives at play. First, it could be argued that speed and accuracy are the ultimate goals of the law-specific models. If so, the question arises as to whether speed and accuracy as goals align with the values, and subsequent objectives, of legal education. This article argues that these goals do not align as they are too narrow and reductionist to be appropriate. Second, it could be argued that speed and accuracy are proxy goals for the wider overall aim of enhancing productivity and efficiency. This article argues that this goal only amplifies the misalignment demonstrated through applying the goals of speed and accuracy. Even if a very generous wider interpretation of such an ultimate goal is given (for example, viewing enhancing productivity as a way to reduce legal costs or increase levels of pro bono work), there is a strong argument that the phenomenon of proxy gaming is occurring. There is little focus on such wider goals and instead much emphasis on speed and accuracy. Even if these were not the intended ultimate goals, they have become central to the adoption and assessment of such AI systems at the expense of any wider intended goals.

In any event, whether viewed as goals or proxy goals, the current focus on speed and accuracy, and/or productivity and efficiency, does not accord with the aims of legal education. This article therefore critiques these goals before moving on to propose the concept of *bildung* as a holistic alternative which can be applied to shape more appropriate goals for law-specific genAI systems.

4. Critique of Speed, Accuracy and Productivity as Goals

Each of the four goals of speed and accuracy, and/or productivity and efficiency, share several significant flaws in the context of legal education, all of which impede the pedagogic development of students. Key concerns arise around both the impact on skills development and the influence on content and levels of contextualization. None of these flaws should be viewed as relating solely to genAI. Instead, they also reflect the broader neo-liberalization of law schools and the ways in which students are encouraged (whether explicitly or impliedly) to engage instrumentally with

legal education (Balan, 2023; De Hingh & Van der Linden, 2025; Thornton, 2014), often involving a focus on high levels of academic attainment and the prioritization of employability (Nicholson & Johnston, 2021). However, it is suggested that the current goals of genAI models applied within legal education have the potential to exacerbate and magnify such wider trends in a range of potentially harmful ways.

4.1 *Skills Development*

The easy availability of AI tools, particularly in legal research, can have a reductive effect on skills development. Law students are commonly expected to develop both legal and wider academic skills during their legal education. In the UK, the Quality Assurance Agency identifies one of the distinguishing features of the law degree as ‘not only being aware of existing law, but also developing an expansive and critical understanding, identifying potential gaps and discussing its future potential’ (Quality Assurance Agency, 2023). Other international documents also emphasize the breadth of law as a subject, with the Council of Europe stating that ‘Legal education, including programmes of continuing education, should seek to strengthen legal skills, increase awareness of ethical and human rights issues, and train lawyers to respect, protect and promote the rights and interests of their clients and support the proper administration of justice’ (Council of Europe, 2000). Even heavily doctrinal interpretations of law require the application of a range of skills to research, analyse and evaluate key words, phrases and rules – a point illustrated by the variety of approaches taken to statutory interpretation in both common law and civil law settings (Carney, 2015).

In contrast, if a law student is researching the law on a specific topic, it is likely that a genAI system focused on speed and accuracy, and/or productivity and efficiency, will (depending on the prompting received) simply produce a short summary of the current law focused on the relevant legislation and/or the most recent leading case. Assuming there are no errors, a student can obtain from this a quick overview which may well suffice for many tasks within legal education, such as citing the source identified as an authority in an essay. However, they will not have undertaken the slower but pedagogically more fruitful process of identifying the relevant case law, scanning it, synthesizing it with their existing knowledge and analysing and evaluating the outcome. It is this experiential process which arguably develops a richer range of legal and academic skills through active participation in the process of learning (Boyle, 2003; Turner et al., 2018). The focus on quickly and efficiently providing a ‘right’ answer overlooks the forms of nuanced analysis, evaluation and critique involved in comparing and contrasting different sources, arguments and examples to reach the required result.

Such specificity and speed in an answer could also remove the pedagogic benefits of law students engaging with liminality. In other words, the sense of disorientation and disconnection which can be experienced by law students when their preconceptions around the law are challenged and they are confronted with new and challenging concepts (Field & Meyer, 2020). This experience of liminality is crucial to students’ transition to a new statement of knowledge, where they grasp core disciplinary concepts (Field & Meyer, 2020). Field & Meyer note that:

the experience of liminality and dissonance, in relation to learning legal threshold concepts such as legal reasoning, is a critical aspect of the epistemological change necessary for successful learning on the path to becoming and thinking like a lawyer because learning success is what results from the 'dissonance' and liminal state. (Field & Meyer, 2020, p. 53)

Extracting information from genAI in a way which prioritizes speed could lead to law students being unwilling or unable to engage in this uncomfortable but vital process. An AI system's almost immediate and (relatively) reliable generation of the 'right' answer could militate against such forms of deeper and more complex engagement by providing an appealing 'short cut', avoiding the need to undertake an uncomfortable transition.

It could be argued that the type of skills and concepts commonly associated with legal education (and still championed in sector-wide guidance, such as the QAA Subject Benchmark Statement for Law) (Quality Assurance Agency, 2023) are becoming outdated (see, for example, Bonilla Maldonado, 2023; Goldsworthy, 2020). The incorporation into the law curriculum of a range of digital literacies and tech-related topics has been proposed as a way to 'future proof' legal education (see, for example, Thanaraj & Gledhill, 2022). However, there is little to suggest that such digital literacies and tech-related content should replace the wide range of other legal and academic skills which are developed through a law degree. There is not a binary choice to be made between fostering appropriate engagements with genAI and developing a broad range of other skills. There may be debates over what balance can, and should, be struck between digital and other skills, but this does not equate to a suggestion that more traditional legal and academic skills should be disregarded altogether. Therefore, the potential of a genAI system to erode such skills must be acknowledged as demonstrating a flaw within its goals (for an example of an article addressing this point via a human-AI agency rubric, see Boetje et al., 2026).

4.2 *Content and Context*

A reliance on AI systems focused on speed and accuracy, and/or productivity and efficiency, is also likely to narrow the scope of the content which law students engage with. Law is often described as appealing to technology companies because it is a rules-based system of the type which allows AI to analyse information and identify patterns within clear parameters (Halliwell & McGlinchey, 2024). However, this fails to acknowledge the important, complex and messy ways in which law interacts with individuals, other institutions and society as a whole, for example, the practical impacts of specific laws on millions of people's lives as highlighted through socio-legal studies in areas from housing and employment to reproduction and end-of-life care (for a useful discussion of the scope of socio-legal work, see Wheeler, 2020). These aspects of law span the sociological and political, with important discussions provided by smaller bodies of critical legal, feminist and

queer scholarship.⁵ Such aspects and critique will not easily be illuminated through a rigid focus on extracting the ‘right’ answer as efficiently as possible.

Rather, the demands of speed and accuracy suggest a narrow approach which eliminates possible ambiguities, contextual discussion and wider multidisciplinary perspectives. This could result in a narrowly focused summary extracting only doctrinal legal principles. This issue can be illustrated through the example of conducting legal research through case judgments. When conducting the research through an AI tool with such goals, students are far less likely to be referred to dissenting judgments within relevant cases or academic critique and commentary which explores the narrative arc and nuances of the judgments, the ways in which the legal principles involved can potentially be applied and distinguished or the opportunities and challenges this brings to the shape and contours of the legal field being researched. In terms of speed, this type of detail would take longer for the model to process (assuming relevant content was within its training data in the first place) and would need to be extricated through a detailed series of prompts, which may lead to disjointed and piecemeal sharing of such additional detail (even if a student were astute enough to recognize what might be missing in the first place). In terms of accuracy, such detail could be seen as unnecessary, irrelevant or even potentially misleading. The same issues would arguably arise with any form of legal research and potentially other uses of genAI in legal education, for example, generating a draft of an essay, advice letter or policy brief.

Even if prompts are designed to extract more contextualized responses, the predictive nature of genAI means that such broader responses are likely to regress to the mean and reflect the focus of the bulk of existing work on a topic, rather than showcasing more creative and less well-explored perspectives and responses (Xie & Xie, 2026). One example of this could be the role of affect and emotions within all aspects of the law. While a body of law and emotions scholarship has developed since the publication of the seminal text *The Passions of Law* (Bandes, 2000), discussion of many aspects of affect and emotions remain ignored or even stigmatized in law school settings (Jones, 2023). A genAI system with narrow goals is unlikely to surface this important but marginalized work, leaving it invisible to many law students. Vital aspects of narrative imagination, drawing on the intersections of affect and cognition, could remain unexplored, diminishing the experience and insights of students (Nussbaum, 1998).

It is arguable that law develops and evolves through understanding, identifying, exploring and seeking to resolve ambiguity and uncertainty. The continued existence of equity as an exercise of judicial discretion illustrates this, for example the then-Justice Denning’s creative extension of the doctrine of promissory estoppel in the case of *Central London Property Trust Ltd*.⁶ Similarly, such development is enriched and enhanced through an understanding of the context in which law functions and its multifaceted influences and impacts, from political considerations to emotional responses. Immediately being provided with a single and linear ‘right’ answer can remove these forms of exploration, iteration and

5 See, for example, <https://www.queerjudgments.org/>.

6 *Central London Property Trust Ltd v. High Tree House Ltd* (1974) 1 KB 130.

contextualization from law students' learning processes, in a way which could be seen as quashing discretion, creativity and critical thinking.

It can be argued that genAI systems used in legal education are merely a useful but limited tool, leaving legal educators scope to expand on the broader context and content of the law (and, indeed, legal and academic skills) away from the digital sphere. However, this argument fails to recognize the pervasive and fast-developing nature of AI. AI agents which 'direct and use LLMs (or foundational models) to complete particular tasks' are becoming increasingly integrated into legal settings (Gage O'Grady & Sims O'Grady, 2025, p. 253). In future, such agentic AI is likely to perform more complex and involved tasks within legal education, whether such tasks are determined by legal educators in formal learning and teaching settings or through students utilizing such platforms as part of their independent study (Reid, 2018; Buhari Bello et al., 2025; Reid, 2018; Mustapha, 2024). Allowing AI systems to evolve without first unpacking and critiquing their goals could result in the existing narrow focus of genAI systems (and the accompanying harms) being perpetuated or even worsened. Therefore, it is vital for wider conversations around goals and alignment to happen now, to feed into the way probable future developments unfold.

Given the significant detrimental impacts on pedagogic development which the current goals of AI systems contribute to, it is necessary to consider whether a different (or broader) set of goals could be devised, to ameliorate some of these harms. This article will therefore go on to consider the concept of *bildung*, arguing for its application as a conceptual framework to guide the development of appropriate alternative goals.

5. The Relevance of the Concept of *Bildung*

Bildung entails a dynamic world view that values independence of mind and spirit grounded in ecological and social interdependence (Rowson, 2019).

The term *bildung*, while Germanic in origin, has numerous interpretations and a wide range of political, cultural and educational applications, most notably those of Herder, Schiller, Goethe, Humboldt, Hegel and Schleiermacher (Herdt, 2021; Nordenbo, 2002). It also overlaps with a range of other philosophical concepts globally (Rowson, 2019; Sander, 2019). However, this article will adopt a Humboldtian interpretation of *bildung*, drawing on four key components: the centrality of holistic individual growth, the emphasis on civic contribution, the recognition of the role of language and the value it places on diversity. In taking this approach, it is worth noting the critique of Barnett who notes (disapprovingly) in relation to the terms *bildung* and 'Humboldtian university' that:

Their sacredness is secured much in the way the two terms have been positioned as markers of educational purity, in response to phenomena such as entrepreneurialism, instrumentalism, learning outcomes, and a separation of research and teaching – phenomena that in turn are positioned as profane. (Barnett, 2024, p. 45; see also Masschelein & Ricken, 2003)

While Barnett views the deification of these concepts as stifling more creative responses to the realities of 21st century higher education, in the context of genAI goals it is arguable that this level of adherence to these concepts is beneficial. It provides a basis of common understanding from which to challenge the existing goals of genAI systems in legal education and to develop alternatives.

It is also worth noting that the role of *bildung* in a digital age has recently been a subject for debate. Thomas suggests that it has been conceptualized either through its separation from the digital, emphasizing a form of unique human essence, or as a 'hybridization with digital systems' (Thomas, 2024, p. 675). This article focuses on the latter for three reasons. First, the pervasive nature of digital technology within legal education and the likelihood of genAI increasing such a trajectory would mean that excluding *bildung* from those aspects of legal education would lead to its role becoming increasingly diminished overall. Second, as discussed previously, genAI is not neutral, but instead is laden with values and biases which, while often hidden, impact broad normative questions and debates in legal education (Fershee, 2025; Madden, 2025). The type of goal alignment addressed in this article is a clear example of this socio-technical and ethical enmeshment. Third, digital technologies are shaped by the human use of them, whether as a tool, a companion or a collaborator. One current example of this is the importance of prompt engineering in the use of genAI chatbots (Qian, 2025). Therefore, to influence the role and uses of genAI, we need to understand and influence staff and student engagement and interaction with it in a way which is shaped by core values within legal education. This means it is necessary to apply *bildung* to digital aspects of legal studies, including genAI, rather than separating it from it. Indeed, it has already been operationalized successfully in relation to applications of digital media and wider digitalization projects within education (Kergel et al., 2022).

Having established the relevance of *bildung* to the digital sphere, it is therefore appropriate to consider how each aspect of the Humboldtian conceptualization of *bildung* could be applied to inform the goals of genAI in legal education.

5.1 Holistic Individual Growth

Humboldt characterizes education as holistic endeavour, focused on individual development, growth and cultivation to achieve 'the highest and most harmonious development' of a person's 'powers to a complete and consistent whole' (Von Humboldt, 1854, p. 11). This implies a focus on not only cognitive achievement but also emotional, moral and ethical growth, to develop the inner self in a neo-humanistic manner (Heidt, 2015; Zelić, 2021). Cleary suggests it draws together 'various aspects of the individual into a whole' or 'integrated character' (Cleary, 2019, p. 6). Humboldt himself identified 'freedom' as a prerequisite for this form of development, with subsequent commentators interpreting this to refer to both autonomy and agency (Von Humboldt, 1854; see also Heidt, 2015, who goes on to critique the application of these concepts within contemporary neo-liberal society).

The focus on individual development, growth and cultivation challenges the notion that speed, accuracy and productivity alone are appropriate goals for genAI

in this context. As discussed previously, such goals can actively impede the pedagogic development of law students. They do not in any way surface, let alone prioritize, wider forms of holistic individual growth and the development of the inner self. A possible argument could be made that a sense of autonomy and agency could be facilitated through genAI empowering students to conduct independent research more easily. However, conversely, genAI could be viewed as diminishing students' autonomy and agency in this regard, by depriving them of the development of broader legal and academic skills which would enable them to engage in a deeper and richer process of research and critique. Applying *bildung* as a conceptual framework requires the shaping of goals which recognize the value of such deeper and richer approaches, enabling them to be acknowledged and incorporated into genAI use.

5.2 Civic Contribution

Another key aspect of Humboldt's *bildung* is the emphasis on the civic contribution of the individual to society. Baur notes that the development inherent in *bildung* 'takes place in active examination and interaction with the environment' and that such engagement can 'trigger changes and transformations of current perspective on the self and the world' (Bauer, 1997, p. 165). The focus on the individual has arguably been the object of greater attention, but several writers have recently highlighted that the Humboldtian notion of *bildung* also has an emphasis on the need for 'reflective and responsible action in, and in interaction with, society' (Sjöström & Eilks, 2025, p. 56; see also Heidt, 2015, pp. 4-5). Indeed, a form of self-alienation through engagement for society is seen as essential for further individual growth and development, providing an opportunity for reflection and learning (Kenkies, 2022; Von Humboldt, 2000). While the dual demands of individual freedom and civic contribution are arguably not fully reconciled within Humboldt's writing (Sorkin, 1983), his recognition of both is important to the application of *bildung* as a conceptual framework to guide the development of genAI in legal education.

GenAI's decontextualizing of information within its outputs, and the form of reductionism which is exacerbated through its focus on the goals of speed, accuracy and/or productivity and efficiency, fails to recognize this form of vital connection between the individual and society. There is unlikely to be a reason to supply students with this type of context and societally embedded content where it may (on the surface at least) appear less relevant and pertinent to a law student's specific prompt. In contrast, using *bildung* as a conceptual framework could require genAI systems to make apparent broader aspects of connection, context and socio-legal interaction in a way which enriches its outputs and the experience of the law students engaging with it.

5.3 Language's Role

A further reason why a Humboldtian characterization of *bildung* is particularly pertinent to genAI systems is his focus on language, namely his belief that:

the character and structure of a language expresses the inner life and knowledge of its speakers, and that language must differ from one another in the same way and to the same degree as those who use them. Sounds do not become words until a meaning has been put into them, and this meaning embodies the thought of a community. (Binti Alias, 2016, p. 29)

In contrast to this perception of language, the predictive nature of genAI means that each word of a sentence within an output is predicted based on statistical probabilities and/or inference mechanisms, to transform the Big Data it is trained on into a novel output (Pezzulo et al., 2024). Any sense of meaning is derived from a composite of Big Data in a way which standardizes the language used. The genAI system's inference process can be viewed as a form of systemization which reflects a community's understanding by drawing on the most frequently used words within its dataset. However, the limitations of the Big Data used to train genAI models, and the inherent flaws and biases within such Big Data, ensure such a meaning is based on historical views which may not be representative of contemporary society, or may be reflective of its worst forms of prejudice and bias (An et al., 2025; Bolukbasi et al., 2016). Language becomes standardized based on a partial and limited representation of individuals and society.

This is arguably particularly likely to be the case where the goals of the genAI system are speed and accuracy, and/or productivity and efficiency. In such instances, language will be reduced to short outputs based on the need to streamline and summarize key information. It may well also absorb dated phrasings and meanings, drawn from a narrow range of legal and scholarly writing. For example, it was not until 2007 that the historical use of masculine pronouns in UK legislation was changed to gender-neutral language (Bailey, 2020). Humboldt's characterization of *bildung* provides a valuable way to expose and critique these issues, raising the possibility of determining goals which promote and celebrate diversity and representation within language.

5.4 Diversity's Value

With regard to diversity more broadly, Humboldt argued that freedom would create a 'diversity of situation' resulting in the communal sharing of 'rich, collective resources' (Von Humboldt, 1854, pp. 11-12). A close reading of Humboldt's work leads Bacher to argue that diversity is a 'necessary and sufficient condition' for *bildung* (Bacher, 2024, p. 169). However, it is arguably diversity of thought in law which is most susceptible to erosion through genAI systems focused on speed and accuracy, and/or productivity and efficiency. As discussed in relation to content and context, a focus on achieving a quick and correct answer to assist productivity is unlikely to showcase multiple perspectives and a diversity of thinking on legal issues. Combined with the limitations and biases found in Big Data, this suggests a form of pseudo-objectivity which is not only decontextualized and potentially misleading but also has a chilling effect on marginalized voices in the law. Adopting *bildung* as a conceptual framework allows for the importance of diversity to be acknowledged, amplifying such voices and providing a nuanced examination of the

law in a way which addresses and explores subjectivities and nuance (for an example of this approach, see Guerrero & Heilman, 2024).

6. Applying *Bildung* to GenAI in Legal Education

If *bildung* is to be applied as the conceptual framework through which to establish the goals of legal education that are enshrined in genAI, then it is important to consider how this can be achieved. If it is accepted that genAI's existing goals do not align with Humboldt's characterization of *bildung*, an initial question arises over how the characteristics he identifies can be operationalized in this context. For example, if it was decided that a goal of a legal genAI system should be 'to maximise the development of individual students', this could be interpreted in a wide range of ways, from a focus on employability and career development to an emphasis on well-being and flourishing, making it susceptible to specification gaming. Such a vulnerability could be partially dealt with by setting more specific proxy goals which align with *bildung*, (while recognizing the potential issue of proxy gaming). However, it can also be addressed by drawing on the body of literature which already exists on topics such as *bildung* in legal education (Nicholson & Wainwright, 2025; Van Rossum, 2023) and liberal legal education (Bradney, 2003), as well as the emerging literature on ways in which AI may be harnessed to provide 'prompts, counterarguments or simulations that challenge students to reflect more critically and ethically' (Farber, 2025, p. 165). This provides an evidence base for crafting a more sophisticated and clearly aligned set of goals.

This solution then leads to the further question of how these different goals could be determined. This section argues that there are three key, complementary approaches to doing this. First, co-creation with system developers; second, developing an appropriate benchmark (or benchmarks) for evaluation; and, third, fostering greater critical AI literacy among legal educators, law students and other key stakeholders in both higher education and law.

6.1 Co-Creation

One way to ensure that existing narrow and reductive goals are challenged and a broader conceptual framework such as *bildung* is applied is to ensure legal educators (and students) are involved as co-creators during the development of relevant genAI systems.

At present, AI systems utilized specifically within legal education are commonly developed by large legal tech companies and publishers (see the examples discussed earlier in this article). The size and scale of these enterprises militate in favour of treating lawyers, legal educators and law students wholly as consumers of AI products, rather than as any form of co-creator. This is further exacerbated by the fact that such genAI systems are commonly used in both legal education and legal practice, meaning that the goals of both are assumed to be, and are treated as, synonymous. This can be highly problematic; for example, it is perhaps understandable that speed has been focused on in a legal practice environment, given the fast-paced nature of legal work and the persistent dominance of charging

by billable hours (despite stringent critiques) (Pilon, 2021). However, there is significantly less justification for this goal retaining its prominence within legal education, given the centrality of pedagogic development in this setting.

Nevertheless, large legal tech companies and publishers clearly see the utility and value of legal education as a marketplace for their products and, as in the foregoing example of Harvey, are seeking to build connections and traction in this field (Harvey Team, 2025). This therefore provides opportunities for stakeholders in legal education to highlight the need for collaboration and co-creation and to emphasize the differing needs and goals of legal education as opposed to legal practice. Moving beyond the sales techniques of free demonstrations and trials and seeking more meaningful engagement with such companies and publishers is unlikely to be simple but could provide significant benefits in the longer term.

6.2 Benchmarks

A connected suggestion is that stakeholders within legal education collaborate to draft a series of benchmarks specifically for the evaluation of legal education focused AI. An example of this, focused on evaluating legal reasoning by LLMs (in the context of the provision of legal advice) is LEGALBENCH.⁷ The type of dialogue and discussion involved in identifying meaningful benchmarks provides opportunities to incorporate *bildung* as a conceptual framework and explore ways to operationalize its key characteristics in this setting. The eventual benchmarks created could be used to generate data and evidence demonstrating where the existing goals of genAI systems are failing to support core aspects of, and values within, legal education. Such detailed discussion and examination would also assist in supporting the development of more sophisticated and nuanced wider goals, using *bildung* as their conceptual framework, which acknowledge and address the possible issues of specification and proxy gaming.

6.3 AI Literacy

Neither of the above two proposals will be possible unless accompanied by a wider shift within legal education, which positions a critical form of AI literacy, including an awareness of AI safety issues, as key to its future. The term AI literacy itself has been conceptualized in various ways, but Almatrafi et al identify several common features across definitions, including ‘a focus on awareness and ability and a focus on social awareness and the impact of the use of AI’ (Almatrafi et al., 2024, p. 5573). Other authors emphasize the importance of a critical literacy approach when engaging with such topics, with Carrington observing that:

texts – printed or multimodal or algorithmic – are not neutral; that they have to power to shape the ways in which individuals are ‘seen’ by their communities and that this power can be harnessed by the individual to allow him/her to participate in a variety of fields. (Carrington, 2018, p. 1276; see also Leander & Burriss, 2020)

7 <https://github.com/HazyResearch/legalbench/>. See also, Guha et al., 2023.

This supports the need for AI literacy to go beyond an understanding of its uses and limitations and encompass socio-political questions and concerns, from environmental impacts to psychological harms and beyond (for insights into a range of relevant issues, see Visvizi et al., 2025). Developing such critical AI literacy has significant potential benefits for both law students and legal educators. It fosters the confidence to engage with and critique the affordances of AI and even empowers them to contribute to its future development (Hyde-Vaamonde & Keller, 2025). In particular, for the purposes of the article, it will enable them to explore and advocate for *bildung* as an appropriate conceptual framework for the development of genAI system goals.

7. Conclusion

Overall, the key benefit of *bildung* as a conceptual framework for determining relevant goals for genAI systems lies in its holistic approach to legal education. This is an approach which already has the benefit of a strong adherence within the higher education academy, providing a fundamental shared understanding, while retaining a level of diffusion which makes it malleable and adaptable (Schneuwly & Vollmer, 2018). The Humboldtian characterization's emphasis on individual development, growth and cultivation, its acknowledgement of the wider societal context, and the way in which it prizes both the meaning of language and the value of diversity are all valuable in determining appropriate goals for genAI. The combination brings a sense of coherence to legal education's pedagogical aims, while remaining sufficiently broad to adapt to different situations and circumstances. It also leaves space for valuable ongoing debates around the relationship between genAI and legal reasoning, such as whether genAI is a tool, a critical friend and to what extent it augments or replaces traditional legal methodologies and functions (Boetje et al., 2026).

Applied as a conceptual framework, *bildung* has the potential to infuse genAI systems with a form of digital holism. This has been defined as aiming at:

a smart spirituality, uniting aesthetic design, functionality, and personalization to perfect the seamless integration promised by the cybernetic prophets. (Islam et al., 2023, p. 1049)

In the context of this article, it infers a sense of harmony and unity between the goals of genAI and the skills, content and context of legal education. Achieving this is a vital counterbalance to the current goals of genAI systems in legal education which have an instrumental focus, drawing on narrow goals which fail to align with the wider objectives, richness and value of legal education.

Emma Jane Jones

Literature

- Ajeovski, M., Barker, K., Gilbert, A., Hardie, L. & Ryan, F. (2023). ChatGPT and the future of legal education and practice. *The Law Teacher*, 57(3), 352-364. <https://doi.org/10.1080/03069400.2023.2207426>.
- Almatrafi, O., Johri, A. & Lee, H. (2024). A systematic review of AI literacy conceptualization, constructs, and implementation and assessment efforts (2019-2023). *Computers and Education Open*, 6, 100173. <https://doi.org/10.1016/j.caeo.2024.100173>.
- An, J., Huang, D., Lin, C. & Tai, M. (2025). Measuring gender and racial biases in large language models: Intersectional evidence from automated resume evaluation. *PNAS Nexus*, 4(3), pgaf089. <https://doi.org/10.1093/pnasnexus/pgaf089>.
- Andreeva, D. & Savova, G. (2024). Artificial Intelligence in the legal field: Law students perspective. *arXiv preprint arXiv:2410.09937*. <https://doi.org/10.48550/arXiv.2410.09937>.
- Babu, S. V. & Banana, K. (2024). A study on narrow artificial intelligence – An overview. *International Journal of Engineering Science and Advanced Technology*, 24, 210-219.
- Bacher, S. (2024). Wilhelm von Humboldt's concept of diversity as an integrated component of his idea(l) of *bildung*. *Ethics and Education*, 19(2), 168-184. <https://doi.org/10.1080/17449642.2024.2364554>.
- Bailey, D. (2020). Breaking down gender stereotypes in legal writing. *Civil Service Blog*. <https://civilservice.blog.gov.uk/2020/01/10/breaking-down-gender-stereotypes-in-legal-writing/>.
- Balan, A. (2023). Neoliberalism, privatisation and marketisation: The implications for legal education in England and Wales. *Cogent Education*, 10(2). <https://doi.org/10.1080/2331186X.2023.2284548>.
- Bandes, S. A. (Ed.). (2000). *The passions of law*. NYU Press.
- Barnett, R. (2024). Can we please stop talking about 'Bildung' – and for that matter, too, 'the Humboldtian university'? *Journal of Praxis in Higher Education*, 6(3), 45-66. <https://doi.org/10.47989/kpdc464>.
- Bauer, W. (1997). Education, bildung and post-traditional modernity. *Curriculum Studies*, 5(2), 163-175. <https://doi.org/10.1080/14681369700200012>.
- Bignami, F. (2022). Artificial Intelligence accountability of public administration. *The American Journal of Comparative Law*, 70(1), i312-i346. <https://doi.org/10.1093/ajcl/avac012>.
- Binti Alias, A. (2016). Wilhelm von Humboldt: A critical review on his philosophy of language, theory and practice of education. *Journal of Creative Writing*, 2(2), 21-29. <https://doi.org/10.70771/jocw.v2i2.25>.
- Boetje, J., de Graaf, Z., van Ginkel, S., Smakman, M., Versendaal, J. & Barendsen, E. (2026). The human-AI agency rubric: A reflective instrument for human-AI collaboration in digital information problem solving. *SSRN*. <https://ssrn.com/abstract=6408618>; <http://doi.org/10.2139/ssrn.6408618>.
- Bolukbasi, T., Chang, K.-W., Zou, J. Y., Saligrama, V. & Kalai, A. T. (2016). Man is to computer programmer as woman is to homemaker? Debiasing word embeddings. *arXiv, arxiv.org/abs/1607.06520v1*. <https://doi.org/10.48550/arXiv.1607.06520>.
- Bonilla Maldonado, D. (2023). Legal education and technological innovation: A critical essay. *Latin American Law Review*, 10, 1-36. <https://doi.org/10.29263/lar10.2023.01>.
- Boyle, R. A. (2003). Employing active-learning techniques and metacognition in law school: Shifting energy from professor to student. *University of Detroit Mercy Law Review*, 81, 1-30.

- Bradney, A. (2003). *Conversations, choices and changes: The liberal law school in the 21st century*. Hart Publishing.
- Buhari Bello, S., Sani Abdurrahman, M. & Sani Bebeji, U. (2025). Artificial intelligence in legal education: Perceptions of ease of use and usefulness of ChatGPT among undergraduate law Students of Ahmadu Bello University, Zaria. *LexScriptio: A Journal of the Department of Jurisprudence and Public Law*, 2(1), 81-94.
- Carney, G. (2015). Comparative approaches to statutory interpretation in civil law and common law jurisdictions. *Statute Law Review*, 36(1), 46-58. <https://doi.org/10.1093/slr/hmu019>.
- Carrington, V. (2018). The changing landscape of literacies: Big data and algorithms. *Digital Culture & Education*, 10(1), 67-76.
- Chandra Bikkasani, D. (2025). Navigating artificial general intelligence (AGI): Societal implications, ethical considerations, and governance strategies. *AI and Ethics*, 5(3), 2021-2036. <https://doi.org/10.1007/s43681-024-00642-z>.
- ChatGPT. (n.d.). *Overview*. Retrieved 8 January 2026, from <https://chatgpt.com/overview/>.
- Claude. (2026). *Legal*. <https://claude.com/plugins/legal>.
- Cleary, B. (2019). Reinterpreting bildung in social pedagogy. *International Journal of Social Pedagogy*, 8(1), 3-15. <https://doi.org/10.14324/111.444.ijsp.2019.v8.1.003>.
- Council of Europe. (2000). *Recommendation No. R(2000)21 of the Committee of Ministers to member States on the freedom of exercise of the profession of lawyer (Adopted by the Committee of Ministers on 25 October 2000 at the 727th meeting of the Ministers' Deputies)*. Retrieved 9 January 2026, from <https://search.coe.int/cm?i=09000016804d0fc8>.
- Dahl, M., Magesh, V., Suzgun, M. & Ho, D. E. (2024). Large legal fictions: Profiling legal hallucinations in large language models. *Journal of Legal Analysis*, 16(1), 64-93. <https://doi.org/10.1093/jla/laae003>.
- De Hingh, A. & Van der Linden, T. (2025). Why allowing law students to use GenAI for writing assignments is a bad idea: Some reflections on the labour market orientation on HLE curriculum decisions. *European Journal of Legal Education*, 6(2), 539-572.
- Digital Education Council. (2024). *What students want: Key results from DEC's global AI student survey 2024*. <https://www.digitaleducationcouncil.com/post/what-students-want-key-results-from-dec-global-ai-student-survey-2024>.
- European Commission, High-Level Expert Group on Artificial Intelligence. (2019). *A definition of AI: Main capabilities and scientific disciplines*. https://ec.europa.eu/newsroom/dae/document.cfm?doc_id=56341.
- Farber, S. (2025). Back to the roots: Rethinking legal education in the age of artificial intelligence. *The Law Teacher*, 60(1), 163-180. <https://doi.org/10.1080/03069400.2025.2549645>.
- Fershee, J. (2025). Reimagining legal education: The role and purpose of law schools in modern society-a values-driven approach to modern legal education. *University of Toledo Law Review*, 57, 229-242.
- Field, R. M. & Meyer, J. H. F. (2020). Threshold concepts in law: Intentional curriculum reform to support law student learning success and well-being. In E. Jones & F. Cownie (Eds.), *Key directions in legal education: National and international perspectives* (pp. 142-157). Routledge.
- Firt, E. (2023). Calibrating machine behavior: A challenge for AI alignment. *Ethics and Information Technology*, 25(3), 42. <https://doi.org/10.1007/s10676-023-09716-8>.
- Gabriel, I. (2020). Artificial intelligence, values, and alignment. *Minds and Machines*, 30(3), 411-437. <https://doi.org/10.1007/s11023-020-09539-2>.

Emma Jane Jones

- Gabriel, I. & Ghazavi, V. (2022). The challenge of value alignment. In Véliz, Carissa (Ed.), *The Oxford handbook of digital ethics* (pp. 336-355). Oxford University Press.
- Gage O'Grady, C. & Sims O'Grady, C. (2025). Agentic workflows in the practice of law – AI Agents as Ethics Counsel. *Georgetown Journal of Legal Ethics*, 38, 247-272.
- Gemini. (n.d.). *Our approach to the Gemini App*. <https://gemini.google/our-approach/>.
- Goldsworthy, D. (2020). The future of legal education in the 21st century. *The Adelaide Law Review*, 41(1), 243-265.
- Google. (2025). *Notebook LLM*. <https://notebooklm.google/>.
- Guerrero, V. & Heilman, R. (2024). *Reclaiming narratives: AI as a tool for equitable voice amplification in evaluation*. https://www.fireside-consulting.com/uploads/1/4/0/1/140153715/sje_ai_practice_guide_final_12.05.2024_rh_vg.pdf.
- Guha, N., Nyarko, J., Ho, D. E., et al. (2023). LegalBench: A collaboratively built benchmark for measuring legal reasoning in large language models. *Advances in Neural Information Processing Systems*, 36, 44123-44279.
- Haenlein, M. & Kaplan, A. (2019). A brief history of artificial intelligence: On the past, present, and future of artificial intelligence. *California Management Review*, 61(4), 5-14. <https://doi.org/10.1177/0008125619864925>.
- Halliwell, D. & McGlinchey, M. (2024). *Artificial Intelligence's taking growing role in legal sector*. Pinsent Masons. <https://www.pinsentmasons.com/out-law/analysis/artificial-intelligence-growing-role-legal-sector#:~:text=How%20AI%20is%20enhancing%20efficiency,nuanced%20reasoning%20which%20lawyers%20bring>>.
- Harding, J. & Domenico Kirk-Giannini, C. (2025). What is AI safety? What do we want it to be? *Philosophical Studies*, 182, 1495-1518. <https://doi.org/10.1007/s11098-025-02367-z>.
- Harvey. (n.d.). *Harvey is AI designed for legal and professional services*. Retrieved 12 March 2026, from <https://www.harvey.ai/>.
- Harvey Team. (2025). *Harvey expands law school program to the United Kingdom*. <https://www.harvey.ai/blog/harvey-expands-law-school-program-to-the-united-kingdom>.
- Heidt, I. (2015). Exploring the historical dimensions of bildung and its metamorphosis in the Context of Globalization. *L2 Journal: An Open Access Refereed Journal for World Language Educators*, 7(4), 1-16. <https://doi.org/10.5070/L27425457>.
- Hendrycks, D. (2025). *Introduction to AI safety, ethics and society*. Routledge.
- Herd, J. A. (2021). The bildung tradition. In *Routledge encyclopedia of philosophy online*. Taylor & Francis. <https://www.rep.routledge.com/articles/thematic/the-bildung-tradition/v-1>.
- Hyde-Vaamonde, C. & Keller, A. (2025). Testing the frontier: Generative AI in legal education and beyond. *Law, Technology and Humans*, 7(3), 52-61. <https://doi.org/10.5204/lthj.4037>.
- Islam, G., Pillet, J.-C., Navazhylyava, K. & Barros, M. (2023). High-performance connections: Digital holism and communicative capitalism at HappyAppy. *Organization*, 30(5), 1046-1073.
- Johnsen, M. (2024). *AI Alignment*. Published Independently.
- Jones, E. (2023). Incorporating an affective framework into liberal legal education to achieve the development of a 'better person' and 'good citizen'. *European Journal of Legal Education*, 4(1), 27-69.
- Kenklies, K. (2022). Alienation: The foundation of transformative education. *Journal of Philosophy of Education*, 56(4), 577-592. <https://doi.org/10.1111/1467-9752.12703>.
- Kergel, D., Heidkamp, B. & Kuckuck, M. (Eds.). (2022). *Bildung in the digital age: Exploring bildung through digital media in education*. Routledge.

- Kirk, H. R., Vidgen, B., Röttger, P. & Hale, S. A. (2024). The benefits, risks and bounds of personalizing the alignment of large language models to individuals. *Nature Machine Intelligence*, 6(4), 383-392. <https://doi.org/10.1038/s42256-024-00820-y>.
- Krakovna, V., Uesato, J., Mikulik, V., Rahtz, M., Everitt, T., Kumar, R., Kenton, Z., Leike, J. & Legg, S. (2020). Specification gaming: The flip side of AI ingenuity. *Google Deep Mind*. <https://deepmind.google/blog/specification-gaming-the-flip-side-of-ai-ingenuity/>.
- Leander, K. M. & Burriss, S. K. (2020). Critical literacy for a posthuman world: When people read, and become, with machines. *British Journal of Educational Technology*, 51(4), 1262-1276.
- Lee, P. (2016). Learning from Tay's introduction. *Official Microsoft Blog*. <https://blogs.microsoft.com/blog/2016/03/25/learning-tays-introduction/Microsoft>.
- LexisNexis. (2025). *Lexis+ AI*. Retrieved 19 September 2025, from <https://www.lexisnexis.co.uk/products/lexis-plus-ai>.
- Ma, B. & Hou, Y. (2021). Artificial intelligence empowers the integrated development of legal education: Challenges and responses. *Future Human Image*, 16, 43-54.
- Madden, R. (2025). The dignity of legal education vs artificial intelligence in Australian law schools. *ANU Journal of Law and Technology*, 6(1), 7-50. <http://doi.org/10.2139/ssrn.5563519>.
- Masschelein, J. & Ricken, N. (2003). Do we (still) need the concept of bildung? *Educational Philosophy and Theory*, 35(2), 139-154. <https://doi.org/10.1111/1469-5812.00015>.
- Microsoft. (2025). *What is a Copilot?* <https://www.microsoft.com/en-us/microsoft-copilot/copilot-101/what-is-copilot>.
- Mustapha, S. (2024). The use of technology and artificial intelligence (AI) in legal education. *Fountain University Law Journal*, 1(2), 70-82. <https://doi.org/10.53704/fulaj.v1i2.546>.
- Ngo, R., Chan, L. & Mindermann, S. (2025). The alignment problem from a deep learning perspective. <https://arxiv.org/pdf/2209.00626>.
- Nicholson, A. & Johnston, P. (2021). The value of a law degree – Part 3: A student perspective. *The Law Teacher*, 55(4), 431-447.
- Nicholson, A. & Wainwright, H. (2025). Reformers or reinforcers: An exploration of how legal educators perceive their role and its relationship with the legal profession. *Legal Studies*, 45(2), 246-263. <https://doi.org/10.1017/lst.2025.12>.
- Nordenbo, S. E. (2002). *Bildung* and the thinking of *bildung*. *Journal of Philosophy of Education*, 36(3), 342-352. <https://doi.org/10.1111/1467-9752.00280>.
- Nussbaum, M. C. (1998). *Cultivating humanity: A classical defense of reform in liberal education*. Harvard University Press.
- OpenAI. (2025). OpenAI Charter. <https://openai.com/charter/>.
- Pezzulo, G., Parr, T., Cisek, P., Clark, A. & Friston, K. (2024). Generating meaning: Active inference and the scope and limits of passive AI. *Trends in Cognitive Science*, 28(2), 97-112. <https://doi.org/10.1016/j.tics.2023.10.002>.
- Pilon, A. (2021). The Billable hour: Critiques of the system and two potential solutions. *University of St. Thomas Journal of Law & Public Policy*, 15, 852-873.
- Quality Assurance Agency. (2023). *Subject benchmark statement for law*. https://www.qaa.ac.uk/docs/qaa/sbs/sbs-law-23.pdf?sfvrsn=c271a881_6.
- Reid, M. (2018). A call to arms: Why and how lawyers and law schools should embrace artificial intelligence. *University of Toledo Law Review*, 50, 411-477. <http://doi.org/10.2139/ssrn.3757672>.

- Ronge, R., Maier, M. & Rathgeber, B. (2025). Towards a definition of generative artificial intelligence. *Philosophy & Technology*, 38(1), 31-46. <https://doi.org/10.1007/s13347-025-00863-y>.
- Rowson, J. (2019). *Bildung in the 21st century – Why sustainable prosperity depends upon reimagining education*. Centre for the Understanding of Sustainable Prosperity. cusp.ac.uk/essay/m1-9.
- Ryan, F. (2020). Rage against the machine? Incorporating legal tech into legal education. *The Law Teacher*, 55(3), 392-404. <https://doi.org/10.1080/03069400.2020.1805927>.
- Sander, W. (2019). The concept of education ('Bildung') as a cultural heritage. Transcultural traditions and perspectives. *Global Education Review*, 6(4), 19-30.
- Schneuwly, B. & Vollmer, H.-J. (2018). Bildung and subject didactics: Exploring a classical concept for building new insights. *European Educational Research Journal*, 17(1), 37-50. <https://doi.org/10.1177/1474904117696096>.
- Schrepel, T. (2025). Generative AI in legal education: A two-year experiment with ChatGPT. *Law, Innovation and Technology*. <https://ssrn.com/abstract=5401422>.
- Serra, A. (2025). AI lawyering skills trainers: Transforming legal education with generative AI. *Case Western Reserve Journal of Law, Technology & the Internet*, 16, 74-132.
- Sheikh, H., Prins, C. & Schrijvers, E. (2023). *Mission AI: The new system technology*. Springer International Publishing.
- Shen, H., Kneare, T., Ghosh, R., Alkiek, K., Krishna, K., Liu, Y., Ma, Z., Petridis, S., Peng, Y.-H., Qiwei, L., Rakshit, S., Si, C., Xie, Y., Bigham, J. P., Bentley, F., Chai, J. Y., Lipton, Z., Mei, Q., Mihalcea, R., ... Jurgens, D. (2024). Towards bidirectional human-ai alignment: A systematic review for clarifications, framework, and future directions. *arXiv*, <https://arxiv.org/html/2406.09264v1>.
- Sjöström, J. & Eilks, I. (2025). The bildung theory – From von Humboldt to Klafki and beyond. In B. Akpan & T. J. Kennedy (Eds.), *Science education in theory and practice: An introductory guide to learning theory* (pp. 33-46). Springer Nature
- Sorkin, D. (1983). Wilhelm Von Humboldt: The theory and practice of self-formation (bildung), 1791-1810. *Journal of the History of Ideas*, 44(1), 55-73. <https://doi.org/10.2307/2709304>.
- Surden, H. (2023). ChatGPT, AI large language models, and law. *Fordham Law Review*, 92, 1941-1972.
- Thanaraj, A. & Gledhill, K. (Eds.). (2022). *Teaching legal education in the digital age: Pedagogical practices to digitally empower law graduates*. Taylor & Francis.
- Thomas, N. (2024). Digital bildung as semantic emplacement. *Journal of Philosophy of Education*, 58, 674-696. <https://doi.org/10.1093/jopedu/qhae063>.
- Thomson Reuters. (2025). *Co-counsel legal*. <https://legal.thomsonreuters.com/en/products/cocounsel-legal>.
- Thornton, M. (2014). Legal education in the corporate university. *Annual Review of Law and Social Science*, 10, 19-35. <https://doi.org/10.1146/annurev-lawsocsci-110413-030714>.
- Tlaie, A. (2024). Using AI Alignment Theory to understand the potential pitfalls of regulatory frameworks. *arXiv preprint arXiv:2410.19749*. <https://arxiv.org/pdf/2410.19749>.
- Turing, A. M. (1950). Computing machinery and intelligence. *Mind*, 59, 433-460.
- Turner, J., Bone, A. & Ashton, J. (2018). Reasons why law students should have access to learning law through a skills-based approach. *The Law Teacher*, 52(1), 1-16. <https://doi.org/10.1080/03069400.2016.1201739>
- Uuk, R., Gutierrez, C. I. & Tamkin, A. (2025). Operationalising the definition of general-purpose AI: Assessing four approaches. *Law, Innovation and Technology*, 17(1), 84-102. <https://doi.org/10.1080/17579961.2025.2469344>.

- Van Rossum, H. E. (2023). *Teaching for impact: Critical thinking, creative thinking and ACT responsibility as defining features of contemporary bildung in Academic Law Schools*. <https://research.vu.nl/en/publications/teaching-for-impact-critical-thinking-creative-thinking-and-act-r/>.
- Veale, M., Black, I., Dsouza, M., Fisher, M., Ghaus, M., Gibbs, T., Lynskey, O., O'Kinneide, C., Scotford, E., Thomas, O., Trapova, A. & Trapp, K. N. (2025). Artificial Intelligence, education and assessment at UCL laws: Current thinking and next steps for the UK legal education sector. *UCL Legal Studies Research Paper Series* 04/2025.
- Visvizi, A., Kozłowski, K., Calzada, I. & Troisi, O. (Eds.). (2025). *Multidisciplinary movements in AI and generative AI: Society, business, education*. Edward Elgar Publishing.
- Von Humboldt, W. (1854). *The sphere and duties of government (The limits of state action)*. John Chapman (original work published 1792).
- Von Humboldt, W. (2000). Theory of bildung. In I. Westbury, S. Hopmann & K. Riquarts (Eds.), *Teaching as a reflective practice. The German Didaktik Tradition* (pp. 57-62). Routledge (original work published 1793).
- Webb, M. (2024). Legacy post: Licencing options for generative AI. *Jisc*. <https://nationalcentreforai.jiscinvolve.org/wp/2024/10/15/licensing-options-for-generative-ai/>.
- Wheeler, S. (2020). Socio-legal studies in 2020. *Journal of Law and Society*, 47, 209-226.
- Xie, Y. & Xie, Y. (2026). When artificial intelligence makes everything similar: The risks of content homogenization. *Chinese Journal of Sociology*, 12(2), 157-172. <https://doi.org/10.1177/2057150X261419573>.
- Zelić, T. (2021). Bildung and the historical and genealogical critique of contemporary culture: Wilhelm von Humboldt's neo-humanistic theory of bildung and Nietzsche's critique of neo-humanistic ideas in classical philology and education. In M.-E. Zovoko & J. Dillon (Eds.), *Bildung and Paideia* (pp. 114-123). Routledge.