INTRODUCING THE "HARD" PROBLEM OF ARTIFICIAL LEGAL INTELLIGENCE

JOSEPH J. AVERY*

Judges and legislators now openly acknowledge using artificial intelligence (AI) to draft opinions and shape statutes. This Essay names this emerging phenomenon Artificial Legal Intelligence (ALI)—AI-created law—and argues that ALI presents a "hard" problem: whether law produced or substantively shaped by artificial systems can maintain the normative legitimacy that attaches to human-generated law. In contrast to the "easy" problems of AI in law (e.g., transparency, accuracy, accountability) whose solutions, however difficult, are conceptually imaginable, the hard problem probes a deeper incompatibility. Even if AI replicates human legal reasoning with perfect fidelity, should its outputs still be regarded as law in the full and proper sense? This Essay groups and examines the principal arguments supporting the proposition that there is a hard problem, including those emphasizing law's essential deliberative character, its reliance on first-person moral judgment, and empirical evidence that law's perceived authority diminishes when authored by machines. The Essay concludes by proposing that ALI be recognized as a distinct field of legal scholarship with its own research agenda.

INTRODUCTION

The age of artificial intelligence (AI) in legal decision-making is here already. In September 2023, Lord Justice Colin Birss, of the British Court of Appeal, drew international attention when he publicly described using generative artificial intelligence tools to help draft judicial opinions.[1] He

 Assistant Professor, Department of Business Law & Department of Psychology, University of Miami. Ph.D., 2021, Princeton University; J.D., 2013, Columbia Law School.

[1] Hibaq Farah, *Court of Appeal Judge Praises 'Jolly Useful' ChatGPT After Asking It for Legal Summary*, THE GUARDIAN (Sep. 15, 2023, at 08:58 ET), https://www.theguardian.com/technology/2023/sep/15/court-of-appeal-judge-praises-jolly-useful-chatgpt-after-asking-it-for-legal-summary?hss_channel=lcp-1623767 [https://perma.cc/579U-

recalled asking ChatGPT for a summary of a particular area of law. "I know what the answer is because I was about to write a paragraph that said that," he said, "but it did it for me and I put it in my judgment. It's there and it's jolly useful."[2] Reflecting on the experience, Birss offered a straightforward conclusion: "It is useful and it will be used."[3] The same sentiment is found in U.S. courts. In at least two opinions, Judge Kevin Newsom has experimented with asking large language models (LLMs)[4] to opine on key terms.[5] U.S. Supreme Court Chief Justice John Roberts also formally acknowledged use of AI. When asked, "Can you foresee a day when smart machines—driven with artificial intelligence—will assist with courtroom factfinding or, more controversially, even judicial decision-making?" Roberts replied, "It's a day that's here."[6] These are dispatches from an ongoing present, not predictions about a remote future.

Despite these developments, much contemporary attention remains fixed on an array of ostensibly "easy" AI-related problems: accuracy and

6692]; *Judges in England and Wales Are Given Cautious Approval to Use AI in Writing Legal Opinions*, NBC NEWS (Jan. 8, 2024, at 10:31 ET), https://www.nbcnews.com/tech/tech-news/judges-england-wales-are-cautious-approval-use-ai-writing-legal-opinio-rcna132824 [https://perma.cc/VNA5-J7SZ].

2 Jane Dalton, *Judge Admits Using 'Jolly Useful' ChatGPT to Write Court Ruling*, THE INDEPENDENT (Sep. 15, 2023, at 19:56 BST), https://www.independent.co.uk/news/uk/home-news/chatgpt-aijudge-chatbot-ruling-b2412378.html [https://perma.cc/8QUU-W2G5].

3 Farah, *supra* note 1.

4 Harry Surden, *ChatGPT, AI Large Language Models, and Law*, 92 FORDHAM L. REV. 1941, 1942 (2024) ("LLMs are AI systems that are designed to understand and generate human language . . . .").

5 *See* Snell v. United Specialty Ins. Co., 102 F.4th 1208, 1225–26 (11th Cir. 2024) (Newsom, J., concurring) (discussing benefits and risks of consulting LLMs for ordinary-meaning analyses); United States v. Deleon, 116 F.4th 1260, 1270 (11th Cir. 2024) (Newsom, J., concurring) (reiterating the potential role of LLMs in interpretive analysis).

6 RENSSELAER POLYTECHNIC INSTITUTE, *A Conversation with Chief Justice John G. Roberts, Jr.*, at 10:32 (YouTube, Apr. 12, 2017), https://www.youtube.com/watch?v=TuZEKlRgDEg (on file with the New York University Law Review).

reliability,[7] efficiency,[8] accountability and transparency,[9] fairness and bias,[10] and security and privacy.[11] The reader may object—quite rightly—that none of these are easy in the ordinary sense. Bias, for example, has bedeviled courts, computer scientists, and legal scholars alike.[12] Yet, following philosopher David Chalmers's terminology,[13] I call them "easy" insofar as their solution space is imaginable: We know the kinds of interventions (technical, institutional, doctrinal) that could, in principle, address them.

The problem of transparency illustrates this point.[14] Black-box systems

---

[7] *See, e.g.*, Exec. Order No. 13960, 85 Fed. Reg. 78939, 78940 (Dec. 8, 2020) (requiring federal government agencies to use AI in a way that is "accurate, reliable, and effective").

[8] *See, e.g.*, Ryan Calo & Danielle Keats Citron, *The Automated Administrative State: A Crisis of Legitimacy*, 70 EMORY L.J. 797, 816–17 (2021) (describing AI's promise to streamline legal processes alongside concerns about oversight and legitimacy).

[9] *See, e.g.*, Charles Tait Graves & Sonia K. Katyal, *From Trade Secrecy to Seclusion*, 109 GEO. L.J. 1337, 1370–76 (2021) (analyzing the impact of secrecy on meaningful review of algorithmic decisions); Cary Coglianese & Lavi M. Ben Dor, *AI in Adjudication and Administration*, 86 BROOK. L. REV. 791, 830, 836 (2021) (exploring the need for transparent processes in AI-based administrative systems). *See generally* Sonia K. Katyal, *The Paradox of Source Code Secrecy*, 104 CORN. L. REV. 1183 (2019) (examining proprietary code's tension with public accountability); Rebecca Wexler, *Life, Liberty, and Trade Secrets: Intellectual Property in the Criminal Justice System*, 70 STAN. L. REV. 1343 (2018) (discussing how closed-source algorithms obstruct defendants' rights); Robert Brauneis & Ellen P. Goodman, *Algorithmic Transparency for the Smart City*, 20 YALE J.L. & TECH. 103 (2018) (highlighting the importance of open data for civic trust in AI); Natalie Ram, *Innovating Criminal Justice*, 112 NW. U. L. REV. 659 (2018) (addressing the challenges of revealing algorithmic processes in criminal justice); Danielle Keats Citron & Frank Pasquale, *The Scored Society: Due Process for Automated Predictions*, 89 WASH. L. REV. 1, 21–22 (2014) (calling for due process safeguards in automated scoring systems); Danielle Keats Citron, *Technological Due Process*, 85 WASH. U. L. REV. 1249 (2008) (advocating procedural protections for algorithmic determinations).

[10] *See, e.g.*, Sandra G. Mayson, *Bias In, Bias Out*, 128 YALE L.J. 2218 (2019) (discussing the risk that data-driven tools replicate historical prejudices).

[11] *See, e.g.*, Ari Ezra Waldman, *Power, Process, and Automated Decision-Making*, 88 FORDHAM L. REV. 613, 619–21 (2019) (emphasizing privacy vulnerabilities in automated processes).

[12] *See* Joseph J. Avery & Joel Cooper, *Racial Bias in Post-Arrest and Pretrial Decision Making: The Problem and a Solution*, 29 CORN. J.L. & PUB. POL'Y 257, 275–93 (2019) (overviewing interventions tried and proposing additional avenues).

[13] In brief, Chalmers distinguishes between "easy" problems—functional, mechanistic issues that appear solvable within existing scientific frameworks—and the "hard" problem, which addresses the elusive nature of subjective experience (qualia). David J. Chalmers, *Facing Up to the Problem of Consciousness*, 2 J. CONSCIOUSNESS STUD. 200, 200–02 (1995). I borrow this distinction to highlight a similar divide in artificial legal intelligence: Certain challenges (e.g., bias, transparency) are "easy" because we can envision solutions, while others are "hard" due to their deep conceptual and normative complexities. I borrow Chalmers's terminology by analogy; it's a useful framing here, though of course not a perfect translation of his usage in philosophy of mind. For more on this terminology, see *infra* Part I.

[14] *See, e.g.*, Brauneis & Goodman, *supra* note 9, at 176 (discussing potential solutions, beyond publicizing code, to issues of transparency in public entities' use of algorithms). In law, opacity corrodes both due process and reason-giving. *See* Elizabeth A. Rowe & Nyja Prior, *Procuring Algorithmic Transparency*, 74 ALA. L. REV. 303, 327–30 (2022) (reviewing examples of due process challenges arising from the difficulty of examining use of "algorithmic black boxes").

frustrate review because judges, parties, and the public cannot see how inputs map to outputs.[15] As a result, there are high-stakes controversies surrounding proprietary criminal-justice tools such as COMPAS.[16] Still, transparency is an "easy" problem because partial remedies exist: explanation techniques (e.g., feature-attribution and perturbation methods), auditable pipelines, process logs, model and data documentation, and regulatory duties to disclose, all supplemented by open-source efforts that permit inspection.[17] Additionally, sectoral experience in finance and health care demonstrates that techno-legal regimes can enforce baseline transparency standards at scale.[18] Granted, none of this is trivial; implementation is costly and

---

[15] *See* Walter A. Mostowy, *Explaining Opaque AI Decisions, Legally*, 35 BERKELEY TECH. L.J. 1291, 1291 (2020) (arguing that "[t]he opacity of Al is a novel challenge to accountability and due process").

[16] *See* Julia Angwin, Jeff Larson, Surya Mattu & Lauren Kirchner, *Machine Bias*, PROPUBLICA (May 23, 2016), https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing [https://perma.cc/U476-5Y6U] (describing controversies surrounding COMPAS, a case management and decision assisting software).

[17] *See* Scott M. Lundberg & Su-In Lee, *A Unified Approach to Interpreting Model Predictions*, 30 ADVANCES IN NEURAL INFO. PROCESSING SYS. 4–8(2017), https://papers.nips.cc/paper_files/paper/2017/hash/8a20a8621978632d76c43dfd28b67767-Abstract.html [https://perma.cc/EE6G-YMT8] (providing a unified framework for interpreting predictions, namely, SHapley Additive exPlanations, or SHAP); Marco Tulio Ribeiro, Sameer Singh & Carlos Guestrin, *"Why Should I Trust You?": Explaining the Predictions of Any Classifier*, 22 ACM SIGKDD INT'L CONF. ON KNOWLEDGE DISCOVERY & DATA MINING PROCS. 1135, 1135 (2016) (proposing LIME, an explanation technique that aims to explain the predictions of classifiers in an interpretable and faithful manner); *see also* Sajid Ali, Tamer Abuhmed, Shaker El-Sappagh, Khan Muhammad, Jose M. Alonso-Moral, Roberto Confalonieri, Riccardo Guidotti, Javier Del Ser, Natalia Díaz-Rodríguez & Francisco Herrera, *Explainable Artificial Intelligence (XAI): What We Know and What Is Left to Attain Trustworthy Artificial Intelligence*, 99 INFO. FUSION, no. 101805, 2023 (reviewing different approaches to explainable AI); Bryce Goodman & Seth Flaxman, *European Union Regulations on Algorithmic Decision Making and a "Right to Explanation*," 38 AI MAG., Fall 2017, at 50, 55–56 (arguing that the European Union's General Data Protection Regulation may effectively create a right to AI explanation); Miles Brundage et al., Toward Trustworthy AI Development: Mechanisms for Supporting Verifiable Claims 1 (Apr. 20, 2020) (arXiv manuscript), https://arxiv.org/pdf/2004.07213 [https://perma.cc/SX52-EQRW] (suggesting steps stakeholders can take to improve the verifiability of claims about AI systems, with emphasis on evidencing safety, security, fairness, and privacy); Michele Loi, Andrea Ferrario & Eleonora Viganò, *Transparency as Design Publicity: Explaining and Justifying Inscrutable Algorithms*, 23 ETHICS & INFO. TECH. 253 (2021) (arguing that transparency in machine learning, like explanation, can be defined at multiple levels of abstraction). *See generally* Reid Blackman & Beena Ammanath, *Building Transparency into AI Projects*, HARV. BUS. REV. (June 20, 2022), https://hbr.org/2022/06/building-transparency-into-ai-projects (on file with the New York University Law Review) (overviewing efforts at building transparency into AI).

[18] *See, e.g.*, SEC Regulation S-K, 17 C.F.R. § 229.303 (2026); Interactive Data to Improve Financial Reporting, 74 Fed. Reg. 6776 (Feb. 10, 2009) (codified at 17 C.F.R. pts. 229, 230, 232, 239, 240, 249) (establishing mandatory, standardized corporate disclosures and requiring many issuers to submit financial information in machine-readable XBRL format, enabling large-scale, technology-mediated transparency and regulatory oversight across public markets); Health Information Technology for Economic and Clinical Health (HITECH) Act, Pub. L. No. 111-5, 123 Stat. 226 (2009); Standards for the Electronic Health Record Technology Incentive Program, 42 C.F.R. §§ 495.314-316 (2026) (mandating interoperable electronic health records, audit trails,

contested. But the terrain is at least legible to the law.

The "hard" problem is different. It is about whether AI can make law *at all* while preserving the quality by which we recognize law's legitimacy and authority.[19] For undertakings like surgery or algorithmic trading, once substantive criteria are satisfied, we are content to treat an outcome as what it purports to be, regardless of whether a human or a machine performed the act.[20] With law, however, authority may rest not merely on textually valid outputs or even procedurally proper steps, but on social practices—interpretive communities, justificatory norms, and institutional pedigrees—that are bound up with human participation, engagement, and deliberation.[21] If those practices are constitutive of law's authority, then functionally identical outputs authored by an artificial system may fail to register as law in the same sense. This Essay crystallizes that possibility as the "hard" problem of artificial legal intelligence: whether law produced or substantively shaped by artificial systems can maintain the normative legitimacy that attaches to human-generated law. When AI participates in the making of law, does the product still possess the features that render law binding and authoritative, or does it instantiate a different kind of normative order that calls for distinct governance? Stated plainly: Is AI-generated law *law* as we know it?

This Essay proceeds in three Parts. In Part I, I introduce and distinguish the "easy" and the "hard" problems of artificial legal intelligence. In Part II,

---

access logs, and standardized reporting requirements for health-care providers as a condition of federal reimbursement, demonstrating enforceable baseline transparency through integrated legal and technical controls at scale).

[19] Legitimacy is how, when we perceive law and legal authorities to be fair, we operate from a place of trust, complying with the law and deferring to and cooperating with those authorities. *See* Tom R. Tyler & Jeffrey Fagan, *Legitimacy and Cooperation: Why Do People Help the Police Fight Crime in Their Communities?*, 6 OHIO ST. J. CRIM. L. 231, 234–35 (2008) (defining legitimacy in the context of public cooperation with police).

[20] *See, e.g.*, Ji Woong Kim, Tony Z. Zhao, Samuel Schmidgall, Anton Deguet, Marin Kobilarov, Chelsea Finn & Axel Krieger, Surgical Robot Transformer (SRT): Imitation Learning for Surgical Tasks (July 17, 2024) (arXiv manuscript), https://arxiv.org/pdf/2407.12998 [https://perma.cc/V64Y-HTAA] (assessing success of robotic surgery); Fatima Dakalbab, Manar Abu Talib, Qassim Nasir & Tracy Saroufil, *Artificial Intelligence Techniques in Financial Trading: A Systematic Literature Review*, 36 J. KING SAUD U. - COMPUT. & INFO. SCIS., no. 102015, 2024, at 1, 1–2 (reviewing success of AI in financial contexts).

[21] *See* Frederick Schauer, *Opinions as Rules*, 62 U. CHI. L. REV. 1455, 1465–67 (1995) (arguing that, unlike other governmental actors, courts operate through their opinions, their justifications, and their textual outputs). *See generally* Tom R. Tyler & Justin Sevier, *How Do the Courts Create Popular Legitimacy?: The Role of Establishing the Truth, Punishing Justly, and/or Acting Through Just Procedures*, 77 ALB. L. REV. 1095 (2014) (examining role of public perception on court legitimacy); Robin J. Effron, *Reason Giving and Rule Making in Procedural Law*, 65 ALA. L. REV. 683 (2014) (arguing for the addition of reason-giving requirements to civil procedural rules); Diana Kapiszewski & Matthew M. Taylor, *Compliance: Conceptualizing, Measuring, and Explaining Adherence to Judicial Rulings*, 38 LAW & SOC. INQUIRY 803 (2013) (measuring and explaining public authorities' compliance with judicial rulings).

I discuss how AI has entered the realm of law generation, and I follow Pamela Gray, from nearly 30 years ago, in calling this "Artificial Legal Intelligence" (ALI).[22] In Part III, I synthesize the leading theoretical accounts of the hard problem, and I canvass emerging empirical evidence suggesting that AI's deepening role already reshapes the way citizens experience law's authority. Finally, in the Conclusion, I propose ALI be treated as a distinct field with its own research agenda, arguing that the hard problem of ALI is neither a philosopher's parlor game nor a foregone conclusion; it is an urgent, open question that warrants sustained, collective inquiry.

I
THE EASY AND THE HARD: FRAMING THE CENTRAL PROBLEM OF ARTIFICIAL LEGAL INTELLIGENCE

Before we can say what is at stake when artificial systems create law, we need a vocabulary that tracks the phenomena we are trying to describe. The familiar catalogue—accuracy and reliability, efficiency, accountability and transparency, fairness and bias, security and privacy[23]—has dominated legal-technology scholarship for good reason. So far, the harms at issue have been concrete and relatively easy to illustrate: a misclassified document in discovery that skews settlement leverage,[24] a faulty risk score that influences bail and sentencing,[25] a breach that exposes privileged files or sealed records.[26] Yet these problems are "easy" in the only sense in which that term is tolerable here: Their solutions are at least conceivable. For the three examples just mentioned, the articles cited both present the problem and propose practical solutions. In other words, we know, in outline, what improvement would look like and which levers law and engineering can pull to move practice toward it. As a further example, consider accountability and transparency. Although politically contested, these matters are at least legible to existing law and amenable to solution by various means: audit trails, provenance tracking, human-in-the-loop constraints calibrated to risk, reporting obligations, and liability rules that map onto familiar doctrines in

---

22 PAMELA N. GRAY, ARTIFICIAL LEGAL INTELLIGENCE 305–06 (1997).

23 *See supra* notes 7–11.

24 *See* Glenn Posein & Purpose Legal, *Top AI-Related Concerns in eDiscovery—and How to Address Them*, JD SUPRA (Apr. 8, 2025), https://www.jdsupra.com/legalnews/top-ai-related-concerns-in-ediscovery-3866619 [https://perma.cc/G9F5-338F] (overviewing AI mistakes in eDiscovery and how they can be addressed).

25 *See, e.g.*, Angwin et al., *supra* note 16 (discussing the COMPAS sentencing matter); Emily Black, John Logan Koepke, Pauline T. Kim, Solon Barocas & Mingwei Hsu, *Less Discriminatory Algorithms*, 113 GEO. L.J. 53, 57 (2024) (proposing various methods for making algorithms less discriminatory).

26 *See, e.g.*, Michael Conklin, Brian Elzweig & Lawrence J. Trautman, *Legal Recourse for Victims of Blockchain and Cyber Breach Attacks*, 23 U.C. DAVIS BUS. L.J. 135 (2023) (overviewing cyber breaches and potential legal responses).

adjacent domains such as medical devices, administrative reason-giving, and autonomous vehicles.[27]

Fairness and bias occupy a liminal place in this schema. On a first pass, their solutions can also look "easy": de-bias datasets, constrain models, enforce parity conditions ex ante or ex post, add governance and oversight, publish impact assessments.[28] Yet as Sandra Mayson has argued with force, prediction drags the past into the future.[29] For example, where policing, arrest, and charging data are themselves artifacts of stratified enforcement and unequal social conditions, the model that predicts from them cannot help but inherit and amplify those structures.[30] Worse, fairness is not a single target. Equalizing false-positive rates may pull against equalizing predictive value; improving overall accuracy may worsen group-specific error; optimizing for counterfactual fairness may frustrate anticlassification constraints.[31] Said again, attempts to make one aspect of a system fairer can easily make another aspect less fair, change that reduces one kind of mistake may increase others, and efforts to correct for past inequities can clash with rules that limit how certain information may be used. Tradeoffs are not pathologies but the geometry of the space. Still, none of this demonstrates conceptual impossibility. Rather, it shows that fairness is contested and multi-objective, that technical adjustments without structural reform will underperform, and that any satisfactory approach must yoke algorithmic design to policy choices made in public and revisited over time. In that qualified sense, the problem remains "easy": arduous in practice, but not beyond imagination. We can say what success would require, even if we doubt, prudently, our capacity to deliver it amid scarce resources, misaligned incentives, and institutional inertia.

The contrast I want to draw is with a different class of questions (call them the "hard" problems of artificial legal intelligence) whose difficulty is not merely practical. Borrowing the frame from Chalmers, the "easy" problems admit mechanistic or institutional solutions; the "hard" problems present a conceptual gap.[32] For law, the gap appears when we ask whether

---

27 *See, e.g.*, U.S. GOV'T ACCOUNTABILITY OFF., GAO-21-519SP, ARTIFICIAL INTELLIGENCE: AN ACCOUNTABILITY FRAMEWORK FOR FEDERAL AGENCIES AND OTHER ENTITIES (2021) (outlining various mechanisms for transparency and risk management plans that align with liability doctrines).

28 *See* Black et al., *supra* note 25, at 104–09 (offering a wealth of techniques for debiasing).

29 *See* Mayson, *supra* note 10, at 2263–67, 2282–96 (arguing that input variables must be regulated to avoid biasing predictions on biased historical data, particularly when predictions of risk are products of structural forces).

30 *See id.* at 2251–61 (explaining how various sources of data, such as past-crime data, can reflect social disparities that may be reflected when using the data to make predictions).

31 *See id.* at 2267–77 (analyzing the shortfalls of algorithmic affirmative action).

32 *See* Chalmers, *supra* note 13, at 200–04 (defining the qualities of "easy" and "hard" problems).

AI can make law at all and still have what we recognize as law's distinctive authority. It is trivially true that AI can produce texts identical in wording and structure to opinions or statutes; judges, legislators, and lawyers have already used such systems to draft them. The question is whether those outputs are law. If the human judge signs the opinion, we treat it as law. If the same text is generated, or even reasoned to, by a machine—absent or overriding human authorship—do we (or *can* we) regard it as law in the same sense? Put differently, is there anything about the mode of production that helps constitute the object we call law?

One response, tempting but question-begging, is to smuggle the answer into the premise by defining law as an essentially human practice. That way lies circularity.[33] A more defensible route focuses on effects rather than process. So let us assume, for the sake of argument, that reasonable people may conclude that law is not an essentially human undertaking. Could those same reasonable people accept the existence of a hard problem of artificial legal intelligence? Perhaps. To see how, we must shift our perspective from process to effects: How do we perceive law? What does it look like to us? How do we respond to it?

Imagine two utterly indistinguishable Pink Lady apples, each with a mottled pink-and-green skin, crisp flesh, and a sweet-tart flavor. You pick up the first one and recognize it as edible, worth preserving, and generally desirable as food, so you wouldn't dream of casually throwing it away. Now imagine that when you pick up the second apple, you experience none of those perceptions. You feel no sense of its being food, no sense of value, and no impulse to keep it. Absent any tampering or other explanation, a reasonable observer might conclude that this second object must not, despite appearances, be an apple. After all, part of what it means to be an apple is to evoke those characteristic responses in people, an awareness of its edibility and value. If those responses are missing, then unless we can give an extraordinary account of how a real apple failed to produce them, we are left to conclude that the thing in question must not be an apple.

As Donald Regan put it, law bears a "halo": a felt, if contestable, claim to respect that organizes our response to it, shapes our mutual expectations, and sometimes moves our moral judgments.[34] There are domains (mathematics, for instance) where we are content with functional equivalence. Saul Kripke's reading of Ludwig Wittgenstein can insist that machines do not "do math" in the human sense,[35] but so long as the answers

---

33 Although, by necessity, we must entertain arguments as to why law might be essentially human. *See infra* Part III.

34 Donald H. Regan, *Law's Halo*, 4 SOC. PHIL. & POL'Y 15 (1986).

35 *See generally* SAUL A. KRIPKE, WITTGENSTEIN ON RULES AND PRIVATE LANGUAGE: AN ELEMENTARY EXPOSITION 20–25 (1982) (discussing the distinction between syntactic and semantic understanding in computational systems); Jeff Buechner, *Not Even Computing Machines*

are correct, most of us are untroubled. Law may be different. If the mode of production helps constitute the object and if our recognition of law's halo partly depends on the human deliberation, justificatory practices, and institutional pedigree by which it is wrought, then functionally identical outputs may fail to register equally as law. That is the hard problem: whether there exists an incompatibility between AI as creator and law as such, not because we cannot imagine a machine producing law, but because the machine's authorship drains the product of the very status "law" names. The remainder of this Essay takes seriously the possibility of incompatibility, not as a foregone conclusion, but as a live question about authority rather than accuracy, about legitimacy rather than latency.

II

FROM TOOLS TO CREATORS: THE RISE OF ARTIFICIAL LEGAL INTELLIGENCE

Against that backdrop of the familiar, manageable problems that have so far dominated the field, the present moment looks less like an incremental advance and more like a categorical shift.[36] For decades, "AI and law" meant litigation about technology or, at most, the insertion of technology into peripheral legal tasks. HYPO, the celebrated case-based system of the late 1980s—an early attempt to mimic how lawyers argue by comparing cases—modeled argumentative moves in trade secrets disputes with real sophistication,[37] but it remained a research artifact; no judge treated HYPO as a co-author of doctrine, and no court regarded its outputs as the reasons that justified a judgment. Risk-assessment systems such as COMPAS altered lives at sentencing and bail,[38] but they were actuarial instruments, inputs to human decision, not engines of legal reasoning or sources of legal text. They shaped how law was applied, not what law was. The line, though imperfect, felt stable: Tools could inform, even influence, but they did not create.

Commercial generative artificial intelligence, powered by large language models, ruptures that settlement. Judges now acknowledge using generative systems to draft the very paragraphs that appear in judgments. As mentioned in the Introduction, this phenomenon has already seen wide-ranging production. In the United Kingdom, Lord Justice Birss has said as

Can Follow Rules: Kripke's Critique of Functionalism, *in* SAUL KRIPKE 343, 343–67 (Alan Berger ed., 2011) (discussing Kripke and the argument that computers cannot truly do math—at least not in the robust sense that humans do).

36 *See* Erik Brynjolfsson, Alex Pentland, Nathaniel Persily, Condoleezza Rice & Angela Aristidou, *Introduction: Artificial Intelligence and Democracy in America*, DIGITALIST PAPERS, https://www.digitalistpapers.com/essays/introduction [https://perma.cc/V9WL-G2U8] (last visited Jan. 17, 2026) ("We, too, stand at technological, economic, and political crossroads that demand creative rebuilding or reinvention of new institutions.").

37 *See* Kevin D. Ashley, *Reasoning with Cases and Hypotheticals in HYPO*, 34 INT'L J. MAN-MACHINE STUD. 753 (1991) (describing and examining HYPO's functions).

38 *See* Angwin et al., *supra* note 16.

much;[39] in the United States, Judge Kevin Newsom has publicly probed LLMs' treatment of contested terms,[40] and Chief Justice Roberts has stated, with equal parts awe and anxiety, that the day of AI-assisted judging "is here."[41] Beyond the Anglophone world, German courts have moved from digitization to deployment: Systems such as OLGA and "Frauke" sort cases, extract facts, and assemble draft judgments at scale, compressing timelines in ways that change what judicial work consists of.[42] Estonia has allegedly piloted small-claims adjudication by machine.[43] China's Internet Courts have, for years, intertwined automated assistance with adjudication in e-commerce disputes.[44] None of this is mere tooling at the margins; it is the use of generative and classificatory systems to produce the materials from which doctrine is made, applied, and used to bind parties, guide officials, and signal norms to the broader public.

The legislative side shows the same drift towards AI use. A Massachusetts state senator introduced a bill drafted in part by a generative model, a bill that pointedly regulated the very class of systems enlisted to help write it.[45] Congressional offices distributed licenses for AI and issued internal guidance for its use as an opening bid in normalizing machine involvement in drafting, comparison, and revision of legislation.[46] The point

---

39 Farah, *supra* note 1.

40 *See* Snell v. United Specialty Ins. Co., 102 F.4th 1208, 1225–26 (11th Cir. 2024) (Newsom, J., concurring) (discussing benefits and risks of consulting LLMs for ordinary-meaning analyses); United States v. Deleon, 116 F.4th 1260, 1270 (11th Cir. 2024) (Newsom, J., concurring) (reiterating the potential role of LLMs in interpretive analysis).

41 RENSSELAER POLYTECHNIC INSTITUTE, *supra* note 6.

42 Eckard Schindler, *Judicial Systems Are Turning to AI to Help Manage Vast Quantities of Data and Expedite Case Resolution*, IBM (Oct. 22, 2025), https://www.ibm.com/case-studies/blog/judicial-systems-are-turning-to-ai-to-help-manage-its-vast-quantities-of-data-and-expedite-case-resolution [https://perma.cc/C39A-ZY9N] (detailing German judges' use of OLGA and "Frauke").

43 Tara Vasdani, *From Estonian AI Judges to Robot Mediators in Canada, U.K.*, LAW360 CAN. (June 13, 2019, at 11:47 ET), https://www.law360.ca/ca/articles/1748405/from-estonian-ai-judges-to-robot-mediators-in-canada-u-k [https://perma.cc/AL3Y-6XHL].

44 Tara Vasdani, *Robot Justice: China's Use of Internet Courts*, LAW360 CAN., (Feb. 5, 2020, at 11:07 ET) https://www.law360.ca/ca/articles/1750396/robot-justice-china-s-use-of-internet-courts [https://perma.cc/2937-HZWE]; Alena Zhabina, *How China's AI Is Automating the Legal System*, DEUTSCHE WELLE (Jan. 20, 2023), https://www.dw.com/en/how-chinas-ai-is-automating-the-legal-system/a-64465988 [https://perma.cc/L7C3-ASCX].

45 *See* Mohar Chatterjee, *AI Just Wrote a Bill to Regulate Itself*, POLITICO (July 19, 2023, at 16:37 ET), https://www.politico.com/newsletters/digital-future-daily/2023/07/19/why-chatgpt-wrote-a-bill-for-itself-00107174 [https://perma.cc/TKM6-LGWH].

46 *See* Bridget C.E. Dooling & Mark Febrizio, *Robotic Rulemaking*, BROOKINGS (Apr. 4, 2023), https://www.brookings.edu/articles/robotic-rulemaking [https://perma.cc/X8S2-TA67] (describing emerging uses of large language models and other AI tools in federal rulemaking); Maya Kornberg, Marci Harris & Aubrey Wilson, *Congress Must Keep Pace with AI*, BRENNAN CTR. FOR JUST. (Feb. 8, 2024), https://www.brennancenter.org/our-work/research-reports/congress-must-keep-pace-ai [https://perma.cc/639J-U2V8] (arguing that Congress is lagging in AI expertise and capacity and urging investments in staff, infrastructure, and oversight

is not that machines have displaced human signatures. It is that the generative act, the shaping of the words that become law, is now (at least in limited areas) partly a machine act, with humans increasingly in the role of curators, editors, and adopters of reasons they did not fully originate.

If we need a name for this phenomenon of AI creating law, we already have one worth rehabilitating. Pamela Gray's *Artificial Legal Intelligence*, published in 1997, surveyed the first wave of programs that tried to capture legal reasoning and, with a striking mix of restraint and audacity, envisioned law consolidating into a codified computer system of legal services.[47] The book was prescient because it saw, before most of us did, that ALI might migrate from the periphery to the core, from research aids and retrieval systems to participation in the authorship of legal reasons. If earlier decades were about applying AI to law, the present decade is about AI as a participant in law's creation.

Once we distinguish easy from hard problems, the empirical claim that AI is creating law does not slide into the old catalogue of tractable worries about accuracy, reliability, or privacy. It triggers the hard question. Perhaps nothing essential turns on authorship, and the halo[48] will survive translation to a new medium of making; perhaps the institutional scaffolding (signatures, votes, reasons offered in public) does all the normative work, no matter who or what drafts the first sentence. But perhaps not. If authorship is part of authority, if reason-giving is not merely the presentation of propositions but the social act of persons offering justifications to other persons within recognized institutions, then substitution at the point of

---

mechanisms so the institution can govern AI effectively); COMM. ON HOUSE ADMIN., SUBCOMM. ON MODERNIZATION, 118TH CONG., FLASH REPORT: ARTIFICIAL INTELLIGENCE STRATEGY & IMPLEMENTATION 2–3 (2023), https://cha.house.gov/_cache/files/4/0/4043ad12-14f9-45e5-987c-4119e86b87e4/A62A82F4AA993321DD02F10F0B96C4B0.cha-modernization-flash-report-9-14-23-99-.pdf [https://perma.cc/EMM3-CJDL] (outlining a strategic vision for AI use in the U.S. House of Representatives, including pilot programs, staff guidance, and principles for responsible adoption); Alexandra Kelley, *Senate's Top Tech Official Greenlights Research Use of Generative AI*, NEXTGOV/FCW (Dec. 19, 2023), https://www.nextgov.com/artificial-intelligence/2023/12/senates-top-tech-official-greenlights-research-use-generative-ai/392896 [https://perma.cc/3K22-DNMG] (reporting that the Senate's chief information officer authorized limited use of several generative AI systems for research and evaluation purposes); DAVID FREEMAN ENGSTROM, DANIEL E. HO, CATHERINE M. SHARKEY & MARIANO-FLORENTINO CUÉLLAR, GOVERNMENT BY ALGORITHM: ARTIFICIAL INTELLIGENCE IN FEDERAL ADMINISTRATIVE AGENCIES 16 (2020), https://law.stanford.edu/wp-content/uploads/2020/02/ACUS-AI-Report.pdf [https://perma.cc/EBC9-QX46] (surveying federal agencies' use of AI tools, mapping benefits and risks across case studies, and recommending governance frameworks to ensure transparency, accountability, and participation).

47 GRAY, *supra* note 22, at 305–06 ("The technology of artificial legal intelligence has the potential to redress the contemporary problems of inaccessibility of the law, through a computer codification of legal services; it also has the potential to transform the system of social power . . . .").

48 *See supra* note 34 and accompanying text.

authorship may matter in kind, not only in degree. The phenomenon therefore deserves its own name precisely because it forces that choice. Let's call it, in the words of Pamela Gray, "artificial legal intelligence": not intelligence applied to law, but intelligence that helps to make law, and that therefore presses on the boundaries of what we are prepared, conceptually and normatively, to count as law. On that terrain, the "easy" problems keep their urgency, but they no longer exhaust the agenda. The hard problem comes into view, and with it the need for doctrinal, empirical, and institutional analysis that does not assume, prematurely, that the halo will remain intact when the human hand that drafts is joined, and sometimes led, by a machine.

III
FOUNDATIONS OF THE HARD PROBLEM

In this Part, I group and describe the major arguments that suggest the hard problem of ALI exists. Many of these arguments predate the emergence of AI. They are, at base, arguments about what law is.[49] I group the arguments into two broad types. The first emphasizes law's essentially deliberative structure: its dependence on public reasoning, communicative justification, and the institutionalization of person-to-person accountability. The second focuses on law's moral character: its reliance on agents who can occupy a first-person standpoint, capable of recognizing, endorsing, and being bound by reasons.[50] Finally, after examining these theories, I consider empirical evidence regarding the human acceptance of AI-created law. Throughout, I raise and address counterarguments.

A. Essential Deliberation

Most scholars who are skeptical of AI's capacity to create genuine law emphasize that human deliberation is not merely instrumental to legal systems, it is constitutive of them. This theme, which supports the idea of a

---

[49] At the same time, long before the present moment and predating the emergence of AI, there were efforts to mechanize lawmaking by reducing legal reasoning to rule application and minimizing human discretion. These attempts somewhat parallel the rise of Artificial Legal Intelligence. Jeremy Bentham's eighteenth-century project—the Pannomion—remains the canonical example: a comprehensive utilitarian code designed to render judges into neutral executors of legislative will, applying law as if it were a machine. *See* GERALD J. POSTEMA, BENTHAM AND THE COMMON LAW TRADITION 420, 430–34 (Tony Honoré & Joseph Raz eds., 1986); Dean Alfange, Jr., *Jeremy Bentham and the Codification of Law*, 55 CORN. L. REV. 58, 61–64 (1969) (summarizing Bentham's principles).

[50] Are there more than these two categories? Certainly. And are these groupings too coarse, collapsing distinctions that matter? Undoubtedly. But this Essay is only a beginning. As discussed in the Conclusion, *infra*, one of my aims is to help establish the field of artificial legal intelligence and to initiate a research agenda—one that includes refining, testing, reorganizing, and supplementing the arguments surveyed here.

"hard" problem of ALI, appears across jurisprudential traditions and emerging work on AI, law, and legitimacy.

A direct articulation of this view can be found in a 2025 *Harvard Law Review* issue focused on AI: "[I]t matters to people that folks gather and make decisions about their lives—*together*."[51] This statement captures a civic conception of law as inherently participatory, grounded in the shared experience of rule-making and rule-following. The vision is not new but ancient and deeply Aristotelian. Aristotle argued that humans are political animals whose flourishing depends on their engagement in collective deliberation: "[T]he good citizen must possess the knowledge and the capacity requisite for ruling as well as for being ruled."[52] Ruling, for Aristotle, was not about domination, nor even about dictation; it was about collective behavior: "a[] deliberative contribution to the initiation and direction of common action."[53]

Jürgen Habermas famously brought the Aristotelian tradition into modern legal theory. In *Between Facts and Norms*, he presented law as the institutionalization of reasoned public discourse.[54] Like John Rawls, Habermas accepted the Kantian premise that reasoned consensus is the normative foundation for legal legitimacy.[55] Habermas linked this premise to empirical sociology through his dialogue with Niklas Luhmann and the concept of autopoiesis,[56] highlighting how law is both structurally embedded and normatively aspirational.[57] On this account, law's authority is not reducible to formal rules or outputs; it depends on stakeholders' participation in a communicative process aimed at legitimacy.

More recently, Frank Pasquale echoed the theme of human participation

---

51 *Developments in the Law—Artificial Intelligence*, 138 HARV. L. REV. 1554, 1610 (2025).

52 ARISTOTLE, POLITICS 1277b (Ernest Barker trans.).

53 DAVID J. RIESBECK, ARISTOTLE ON POLITICAL COMMUNITY 289 (2016).

54 *See* JÜRGEN HABERMAS, BETWEEN FACTS AND NORMS: CONTRIBUTIONS TO A DISCOURSE THEORY OF LAW AND DEMOCRACY 107–08 (William Rehg trans., MIT Press 1996) (1992) (arguing that legal rules stem from "norms of action," which are valid only if "all possibly affected persons could agree as participants in rational discourses"); William Rehg, *Translator's Introduction* to HABERMAS, *supra*, at ix, xx–xxiv (explaining Habermas's argument on the importance of "communicative action" to the operation of law).

55 Kevin P. Lee, *Is Law Computable? Critical Perspectives on Law and Artificial Intelligence (Hart Publishing, 2022) by Simon Deakin and Christopher Markou*, 3 AUSTRALIAN NAT'L U. J.L. & TECH., Summer 2022, at 75, 76 (book review); *see* Rehg, *supra* note 54, at xii, xx (explaining Habermas's argument with reference to "Kant's appeal to rational consensus as a regulative ideal" and "Rawls's concern with overlapping consensus").

56 Autopoiesis refers to systems that are self-creating or self-organizing. Pablo Razeto-Barry, *Autopoiesis 40 Years Later: A Review and a Reformulation*, 42 ORIGINS LIFE & EVOLUTION BIOSPHERES 543, 546–47 (2012).

57 *See* Kevin W. Gray, *The Autopoietic Turn in Habermas' Legal Philosophy*, 2014 ANCILLA IURIS 41, 49–50; *see also* Lee, *supra* note 55, at 76, 83 (discussing Habermas and Luhmann and arguing for a more capacious understanding of computation and complex systems as they impinge on law).

in his defense of a rule of persons, not machines.[58] Pasquale argued that legal automation threatens to erode the dialogic structures at the heart of law.[59] Law is not merely a set of rules to be executed; it is a system of person-to-person accountability, where reasons must be articulated intelligibly by a responsible agent.[60] Even if an AI system could follow prescribed procedures with perfect fidelity, Pasquale's worry is that something essential might still be missing.[61] Due process is not only about getting the steps right. It is about a person being heard by another person who is charged with responding to their arguments in good faith. Pasquale warned that although automation may increase accuracy or reduce error, it may do so at the cost of legitimacy and public trust.[62] The very idea of due process presupposes communication between persons, an exchange of reasons that recognizes the parties as participants in a shared normative practice.[63]

In making this claim, Pasquale builds upon the Legal Process school.[64] In Richard Fallon's conception, Legal Process focuses on four key features: procedural fairness, an assumed internal connection between law and reasonableness, reasoned elaboration linking pre-existing sources to outcomes, and judicial review as a guarantor of those features.[65] These features are not satisfied by AI systems that deliver legal outputs without interpretive engagement. A system that simply produces a result, even if accurate, bypasses the justificatory work that Legal Process views as constitutive of lawful decision-making. Code-driven dispositions, like smart contracts or automated judgments, cannot offer the kind of reasoned elaboration central to Legal Process ideals.[66] They can state a conclusion, and they can map inputs to outputs, but they cannot participate in the communicative practice through which officials explain how competing claims were understood, weighed, and resolved.

On this view, the problem is not that AI fails to comply with existing procedures; it is that existing procedures presuppose a kind of human engagement that automation cannot supply. Consider a small-claims litigant who appears in court to contest a fine. She explains why she believes the

---

58 Frank Pasquale, *A Rule of Persons, Not Machines: The Limits of Legal Automation*, 87 GEO. WASH. L. REV. 1 (2019).

59 *Id.* at 46.

60 *Id.* at 5.

61 *Id.* at 1, 3, 18, 38 (discussing how automation struggles to capture ambiguities in language, greater legal complexities, and human governance).

62 *Id.*

63 *See id.* at 13 (arguing that the scope and intensity of automation depend on human coordination in developing records of behavior with due process implications).

64 *Id.* at 44–46.

65 Richard H. Fallon, Jr., *"The Rule of Law" as a Concept in Constitutional Discourse*, 97 COLUM. L. REV. 1, 5, 18 (1997).

66 Pasquale, *supra* note 58, at 45–46.

citation was issued in error. An AI system might process her statement, apply the governing rules with perfect fidelity, and return a written disposition that is technically correct. But the structure of the hearing promises more than technical accuracy. It promises that someone will hear her, interpret her arguments, and explain how those arguments were taken into account. It also promises that a responsible agent will stand behind the reasons given, not as a matter of form but as a matter of accountability and integrity. A machine can state a conclusion, but it cannot occupy the interpersonal role that makes the process intelligible as law: a person who listens, deliberates, justifies, and assumes responsibility for the judgment. The concern, then, is not about error rates but about whether the decision can be experienced as lawful in the sense that matters to the parties who are governed by it.

In some sense, these concerns are deepened by W.B. Gallie's account of essentially contested concepts, which has come to dominate legal thinking.[67] Concepts such as justice, democracy, and the rule of law necessarily involve ongoing disagreement; their meaning is constituted through continual interpretive negotiation rather than fixed by rule-following alone.[68] As Jeremy Waldron has emphasized, law is saturated with these concepts;[69] and its development depends on the very practices of argument, justification, and contestation that Pasquale and the Legal Process tradition highlight at the level of individual decision-making. Gallie's framework thus shows that the interpersonal and justificatory dimensions of law are not merely procedural expectations but part of the larger, collective practice through which legal meaning evolves over time. When AI generates legal texts without participating in this continuing argumentative community, it does more than bypass person-to-person accountability or reasoned elaboration; it short-circuits the interpretive contestation that makes law a living, human enterprise.

To conclude this section on essential deliberativeness, I want to briefly introduce three additional theoretical concerns that, if not perfectly mapped, still help to illuminate why human-centered deliberation may be indispensable to law's legitimacy. First, the concerns about AI-created law are sharpened by the opacity of contemporary AI systems.[70] As critics have noted, many advanced AI models function as "black boxes" with inscrutable

---

67 Jeremy Waldron, *Is the Rule of Law an Essentially Contested Concept (In Florida)?*, 21 LAW & PHIL. 137, 148–49 (2002) ("In the law review literature, the use of the term ['essentially contested'] has run wild . . . .").

68 *See* W.B. Gallie, *Essentially Contested Concepts*, 56 PROCS. ARISTOTELIAN SOC'Y 167, 168, 180 (1956) (identifying democracy and social justice as essentially contested concepts); *see also* Waldron, *supra* note 67 (extending Gallie's framework of essentially contested concepts to the rule of law).

69 *See* Waldron, *supra* note 67, at 140–64.

70 *See, e.g.*, Matthew Hutson, *How Does ChatGPT 'Think'?*, 629 NATURE 986, 987–88 (2024) (describing efforts to "peer inside the black box" of AI).

inner workings.[71] While humans are opaque in some respects too, the structure of legal institutions is built to expose reasoning through open hearings, written opinions, and adversarial processes. If AI displaces those mechanisms with silent computation, it risks severing the link between legal outcomes and public justification.

Second, even more troubling is emerging evidence that AI changes how humans deliberate. Recent empirical work shows that people tend to defer to AI systems in morally complex decisions, leading to diminished personal responsibility.[72] This phenomenon of moral offloading suggests that the presence of AI can erode human engagement in legal or ethical reasoning.[73] If people stop taking ownership of legal decisions because they assume AI has already solved them, the moral community that sustains law begins to fray.

Finally, Julian Jonker offers a theoretical framework that, while not developed specifically for law, provides valuable conceptual tools for understanding how AI might erode law's normative structure. In his work on "social alignment," Jonker critiques the narrow focus on value alignment, the idea that AI systems should be engineered to output results that conform to human preferences or moral goals.[74] He argues instead that we must attend to how AI technologies alter the cooperative structures through which human interaction and meaning-making occur.[75] Social alignment concerns the ways AI systems reconfigure trust, authority, and mutual recognition within human networks.[76]

I suggest that this framework has important implications for legal theory. If we apply Jonker's insights to the legal domain, AI's integration may do more than jeopardize valid legal content—it may transform the interfaces through which law is produced and received.[77] In legislatures, for instance, AI tools that generate statutory language could displace the need for coalition-building, negotiation, or iterative compromise. The locus of norm creation shifts from human deliberation to algorithmic suggestion. In

---

71 *Id.*; *see also* Angela Luna, *The Open or Closed AI Dilemma*, BIPARTISAN POL'Y CTR. (May 2, 2024), https://bipartisanpolicy.org/article/the-open-or-closed-ai-dilemma [https://perma.cc/D6TD-9R45] (comparing open, collaborative AI systems with closed, insulated systems); Ben Brooks, *Open-Source AI Is Good for Us*, IEEE SPECTRUM (Feb. 8, 2024), https://spectrum.ieee.org/open-source-ai-good [https://perma.cc/7P2V-PGTW] (arguing that open-source AI is good on account of transparency).

72 Adriana Salatino, Arthur Prével, Emilie Caspar & Salvatore Lo Bue, *Influence of AI Behavior on Human Moral Decisions, Agency, and Responsibility*, 15 SCI. REPS., no. 12329, 2025, at 1, 7–11.

73 *Id.*

74 Julian David Jonker, *Automation, Alignment, and the Cooperative Interface*, 28 J. ETHICS 483, 84 (2024).

75 *Id.* at 484–91.

76 *Id.* at 491–95.

77 Jonker talks about this in terms of a "cooperative interface." *Id.* at 487–91.

judicial contexts, parties who suspect that an AI wrote the opinion in their case may feel that their arguments were not genuinely heard. The cooperative interface, the relational space where recognition and contestation occur, has changed. And with it, the moral authority of the legal act may be diminished.

Jonker's distinction between value alignment and social alignment thus illuminates a version of the hard problem in artificial legal intelligence. Even if AI systems generate substantively correct legal outputs, they may do so in ways that bypass or distort the human deliberative practices that confer law's legitimacy. The concern is not that AI fails to perform legal tasks well, but that it performs them too well, so efficiently and independently that the practices which make law intelligible, participatory, and normatively binding begin to fall away. AI may get the answers right, but without the human process, the answers cease to be law in the full and proper sense.[78]

The theories reviewed in this Section all converge on a shared insight: Law is not something that is achieved and done with; it is done and continually redone, through deliberation over time. Any legal system that omits this dimension risks producing something that may look like law but is not law.

B. A First-Person Deficit

A second category of objections centers not on collective action or deliberation, but rather on something more personal: consciousness and the personal point of view, expressed by Thomas Nagel's what-is-it-likeness of existence.[79] Such objections[80] would maintain that absent subjective experience, free will, or the capacity for authentic moral judgment, AI cannot engage in the kind of normatively saturated reasoning that is constitutive of legal deliberation. Consequently, AI-generated outputs in law, no matter how textually coherent or rhetorically persuasive, are best understood not as genuine legal judgments but as mere simulations, surface approximations devoid of the authority that arises from speaking of justification in the first person.

---

78 Thinking from a legal positivist perspective, we might raise a brief objection: Perhaps the authority of law turns on its formal pedigree rather than on the experiential attributes of its authors. Accordingly, AI-generated legal texts could be fully valid if produced through duly authorized procedures. Yet this objection ultimately sidesteps rather than resolves the concern at issue, for it addresses legal validity in the abstract rather than the social and deliberative practices through which law secures its legitimacy in the eyes of those it governs.

79 Thomas Nagel, *What Is It Like to Be a Bat?*, 83 PHIL. REV. 435, 436 (1974) ("But fundamentally an organism has conscious mental states if and only if there is something that it is like to *be* that organism—something it is like *for* the organism. We may call this the subjective character of experience.").

80 Joshua Davis makes this case, which I discuss in the remainder of this Section. Joshua P. Davis, *Artificial Wisdom? A Potential Limit on AI in Law (and Elsewhere)*, 72 OKLA. L. REV. 51 (2019).

At the pith of these accounts lies a fundamental claim: Moral judgment in legal contexts—especially in the face of interpretive indeterminacy, value conflict, or contested principles—requires a form of consciousness, a first-person perspective that is irreducibly human.[81] Such a perspective allows a legal agent to own reasons,[82] to be responsive to moral criticism, and to stand within the space of reasons not merely as a computational processor of inputs but also as a participant who can be held accountable.[83]

Joshua Davis exemplifies this line of reasoning by arguing that, although AI systems may demonstrate extraordinary proficiency in factual and even probabilistic domains, they nonetheless fall short in the domain of moral wisdom.[84] True moral agency, Davis contends, necessitates not merely the generation of normatively acceptable outcomes, but a standpoint from which those reasons are evaluated, endorsed, and lived.[85]

Can AI engage in this kind of authentic evaluative judgment?[86] Any critique along these lines begins not with empirical concerns about performance but with a conceptual boundary: namely, that syntactic manipulation of symbols, no matter how sophisticated, does not amount to understanding.[87] AI systems process inputs and generate outputs through rules and data structures, yet this computational operation lacks the semantic depth that characterizes human reasoning.[88] The distinction between syntax and semantics is not merely linguistic but ontological, cutting to the core of what it means to reason with, and be responsible for, normative claims.[89]

John Searle's Chinese Room thought experiment offers a vivid

---

81 *Id*. at 55.

82 By "reasons," I mean normative reasons in the philosophical sense of considerations that purport to justify an outcome and that bind an agent as one who can be criticized, persuaded, or held responsible for acting on them, rather than mere causal or informational inputs.

83 *See* Davis, *supra* note 80, at 62 ("The logic is that those who impose legal judgments on others must themselves be subject to the law. The law, in turn, applies only to conscious actors capable of moral agency."); *see also* W. Bradley Wendel, *The Promise and Limitations of Artificial Intelligence in the Practice of Law*, 72 OKLA. L. REV. 21, 42 (2019) ("The law is a means for giving the types of reasons that human moral agents owe to one another, in response to others' demands for accountability."); Kiel Brennan-Marquez & Stephen E. Henderson, *Artificial Intelligence and Role-Reversible Judgment*, 109 J. CRIM. L. & CRIMINOLOGY 137, 149–56 (2019) (arguing that, in a democracy, the entity passing judgment via the law ought to be, in turn, subject to the law).

84 Davis, *supra* note 80.

85 *Id.* at 55–56, 70–73.

86 *See, e.g.*, JOHN R. SEARLE, MIND: A BRIEF INTRODUCTION 79 (2004) [hereinafter SEARLE, MIND: A BRIEF INTRODUCTION] (describing "biological naturalism," the idea that consciousness is an emergent quality of the biological brain); *see also* JOHN R. SEARLE, THE MYSTERY OF CONSCIOUSNESS 8, 13, 161, 210–14 (1997) [hereinafter SEARLE, MYSTERY OF CONSCIOUSNESS] (arguing the same).

87 *See* KRIPKE, *supra* note 35, at 30–31 (contending that syntactic rule-following cannot, by itself, generate meaning or understanding).

88 *See* SEARLE, MYSTERY OF CONSCIOUSNESS, *supra* note 86, at 12, 110; SEARLE, MIND: A BRIEF INTRODUCTION, *supra* note 86, at 46–52, 58–71.

89 *See* Davis, *supra* note 80, at 72–73.

illustration of the distinction between syntactic processing and genuine understanding.[90] Imagine a person confined in a sealed room with no knowledge of the Chinese language. Slips of paper bearing Chinese characters are passed into the room. The person possesses a comprehensive rulebook, written entirely in English, that specifies precisely how to correlate incoming Chinese symbols with outgoing ones. By meticulously following these rules, the person assembles responses that, when passed back out, are indistinguishable from those of a native Chinese speaker. To external observers, it appears as though the person in the room understands Chinese.

But inside the room, no understanding takes place.[91] The person is merely manipulating symbols according to formal rules: syntax without semantics.[92] There is no grasp of meaning, no awareness of what is being communicated, and no capacity to interpret or endorse the content of the exchanges. Understanding requires more than rule-governed manipulation; it requires intentionality, a mind capable of aboutness.

Likewise, AI may simulate linguistic or legal competence without grasping the meaning of its own outputs. Searle's implication is stark: Computational systems, however advanced, operate in the realm of syntax without access to semantics. They do not, and cannot, understand. As such, they lack the subjective awareness necessary to recognize, evaluate, and endorse the normative content of legal arguments.

Daniel Dennett, however, offers a contrasting view, suggesting that consciousness itself may be an elaborate illusion generated by cognitive architecture.[93] On this account, what we experience as awareness or intentionality may be no more than a narrative construct, a byproduct of information processing.[94] If taken seriously, this position destabilizes the boundary between artificial and human cognition: If our own consciousness is illusory, then AI's lack of it may be less disqualifying than presumed. And yet, even under this skeptical view, a key distinction persists: Human agents inhabit the illusion, treating reasons as their own, accepting blame, praise, and moral responsibility. AI systems, by contrast, do not (at least to date, in this author's experience) profess to enter the moral space of reasons.

The jurisprudential implications of this philosophical divide (between participation in the space of reasons and the mere manipulation of reasons as data) are evident in the work of H.L.A. Hart and Ronald Dworkin. Hart's theory of the legal system acknowledges the existence of "penumbral" cases,

---

90 *See* John R. Searle, *Minds, Brains, and Programs*, 3 BEHAV. & BRAIN SCIS. 417, 417–19 (1980).

91 *Id.* at 422. For discussion of opposing views, see Josef Moural, *The Chinese Room Argument*, *in* JOHN SEARLE 214 (Barry Smith ed., 2003).

92 Searle, *supra* note 90, at 422.

93 *See* DANIEL C. DENNETT, CONSCIOUSNESS EXPLAINED 309–14 (1991).

94 *Id.*

situations in which legal rules run out, and the judge must decide not through deduction but through discretion.[95] These are moments of normative indeterminacy that require evaluative judgment, a capacity grounded in the agent's ability to treat reasons as binding.[96] It is precisely this kind of judgment that AI, lacking a standpoint from which to own reasons, cannot exercise.

Similarly, Dworkin's theory of law as interpretive practice presupposes a moral agent capable of engaging with legal principles not as static commands but as normative commitments to be justified within a broader moral narrative.[97] The process of fit-and-justification is not merely algorithmic; it requires a perspective that can assess, revise, and stand behind reasons. Legal interpretation, in Dworkin's model, is not a mechanical function but an exercise in moral responsibility. Indeed, Dworkin's theory is so human, so rooted in the world, that we may venture to call it an embodied legal philosophy.

The implications from these thinkers delineate the contours of a perhaps non-negotiable boundary. Law is not simply a system of rules but a moral practice, and moral practices require participants, that is, beings who can understand, who can express meaning, and who can be held to account. No matter how advanced its simulations, AI lacks the first-person standpoint necessary to enter this normative domain.[98] Or so the argument goes.[99]

The ontological divide emphasized here (between computational output and conscious authorship) highlights AI's alleged incapacity to meet the accountability conditions that law imposes. Though AI may generate outputs that are legally valid in form, they lack the intrinsic legitimacy conferred by aspects of human authorship: the capacity to bear responsibility, to be blamed or praised, to be persuaded by counterargument.[100] In this respect, the limit of AI in law is not simply technical or pragmatic, but metaphysical.

C. Empirical Evidence

Imagine the following scenario: A case has come before the bench. It

---

95 *See* H.L.A. Hart, *Positivism and the Separation of Law and Morals*, 71 HARV. L. REV. 593, 607 (1958).

96 *Id.* at 608.

97 *See* RONALD DWORKIN, LAW'S EMPIRE 225–75 (1986).

98 Davis, *supra* note 80, at 69–88.

99 It may be that AI has a first-person perspective and is conscious, but we just don't realize it yet because our understanding of consciousness is incomplete. Barbara Gail Montero, *A.I. Is on Its Way to Something Even More Remarkable Than Intelligence*, N.Y. TIMES (Nov. 8, 2025), https://www.nytimes.com/2025/11/08/opinion/ai-conscious-technology.html [https://perma.cc/C6U4-WL42].

100 *See* Davis, *supra* note 80, at 69–88 (arguing that because science has not fully captured first-person subjective experience—consciousness, free will, and unified self—AI lacks the moral agency and first-person judgment required for legitimate legal decision-making).

involves a statute governing the use of a public park, a seemingly minor matter until one notices that the dispute raises subtle interpretive questions about what counts as permissible use. After hearing arguments and deliberating, the court issues its opinion.

Now imagine three alternatives. In the first, the court relies solely on its own deliberation. In the second, it uses AI to assist in researching the legal issues. In the third, the court delegates the research and drafting of the opinion to an artificial intelligence system, merely approving the result. Importantly, the text of the ruling remains identical across all three conditions.

This vignette comes from a study I conducted as part of a previous work on AI.[101] It was used to explore public perceptions of what might be called law's normative halo: the sense that legal pronouncements carry a unique kind of moral force. Building atop Donald Regan's formulation,[102] the halo includes at least three facets: (i) moral authority, (ii) moral influence, and (iii) social norm signaling.[103] While not exhaustive, these dimensions capture essential ways in which law functions not merely as a system of rules but as a distinctive moral and cultural institution.

First, law often exerts moral authority.[104] Even where citizens disagree with a law's substance, they may feel a pull to treat it as legitimate, a source of guidance that warrants deference.[105] While theorists differ on whether there is a general moral obligation to obey the law, many agree that law occupies a special role in moral reasoning.[106] Regan's "halo" metaphor articulates this moral valence: We respond to legal commands not just as edicts but as reasons bearing a normative weight.

Second, law has moral influence.[107] It can shape attitudes, refine

---

101 Joseph Avery, *AI and the Erosion of Law's Moral Authority*, 50 BYU L. REV. 895, 934–54 (2025).

102 Regan, *supra* note 34.

103 Avery, *supra* note 101, at 904.

104 JOSEPH RAZ, THE AUTHORITY OF LAW: ESSAYS ON LAW AND MORALITY 3–27, 246–49 (1979).

105 *Id.*; *see also* Regan, *supra* note 34, at 15 (observing that in a reasonably just society, law is understood as morally significant, not neutral).

106 *See* Pasquale, *supra* note 58 (arguing that law's legitimacy depends on human moral judgment and accountability, not mere algorithmic rule application). *See generally* Patricia D. White, *Law and Moral Obligation*, 49 U. CHI. L. REV. 249, 256 (1982) (reviewing Joseph Raz's *The Authority of Law* and discussing the presence or absence of a special relationship between law and moral reasons); Robin Bradley Kar, *The Deep Structure of Law and Morality*, 84 TEX. L. REV. 877, 882 (2006) (arguing that law and morality are deeply and essentially intertwined); Bruce P. Frohnen, *The Irreducible, Minimal Morality of Law: Reconsidering the Positivist/Natural Law Divide in Light of Legal Purpose and the Rule of Law*, 58 ST. LOUIS U. L.J. 467, 468 (2014) (proposing a conception of law that "recognizes a spectrum of law and law-likeness encompassing the possibility of many goals and of many levels of excellence or virtue, down to the nadir of lawlessness").

107 George C. Christie, *On the Moral Obligation to Obey the Law*, 1990 DUKE L.J. 1311, 1319

intuitions, and even shift beliefs.[108] A canonical example is *Loving v. Virginia*,[109] which not only invalidated bans on interracial marriage but also catalyzed widespread changes in public attitudes.[110] Studies show that even hypothetical legal mandates—absent enforcement—can nudge individuals' stances in classic moral dilemmas.[111] Law persuades, not just coerces.

Third, law serves as a social norm signal.[112] People look to legal pronouncements to discern what society deems acceptable.[113] Legal change can alter perceptions of majority opinion. As empirical research by Margaret Tankard and Betsy Levy Paluck has shown, landmark rulings such as *Obergefell v. Hodges*[114] influence individuals' beliefs about prevailing norms—even if those beliefs turn out to be inaccurate.[115] Law sets focal points in moral discourse, even where it provokes resistance, as seen in the aftermath of *Dobbs v. Jackson Women's Health Organization*.[116]

Returning to the experimental vignette, the design aimed to assess whether the public's perception of these three features varied based on AI involvement. The results were striking. When AI was responsible for both

("The stark separation between law and morality—the assumption that the law does not affect morality—is simply untenable. In particular, I wish to assert that the prevailing public morality of any society is very definitely influenced by the law."); *see also* Bert I. Huang, *Law and Moral Dilemmas*, 130 HARV. L. REV. 659, 691, 693 (2016) (presenting empirical work indicating law's moral influence).

108 *See* Huang, *supra* note 107, at 663 (reporting evidence that the presence of law can influence intuitions and beliefs about the trolley problem, a canonical moral dilemma).

109 388 U.S. 1 (1967).

110 GRETCHEN LIVINGSTON & ANNA BROWN, PEW RSCH. CTR., INTERMARRIAGE IN THE U.S. 50 YEARS AFTER *LOVING V. VIRGINIA* 7, 24–28 (2017), https://www.pewresearch.org/wp-content/uploads/sites/20/2017/05/Intermarriage-May-2017-Full-Report.pdf [https://perma.cc/4KLP-HR53] (highlighting increases in both public acceptance and actual rates of interracial marriage).

111 *See* Huang, *supra* note 107, at 691–95 (describing experiments introducing legal mandates on the trolley problem); *see also* Kenworthey Bilz & Janice Nadler, *Law, Psychology, and Morality*, 50 PSYCH. LEARNING & MOTIVATION 101, 102–03, 107–13 (2009) (arguing that law alters moral stances through signaling and cognition, not necessarily through its enforcement).

112 *See* Robert Cialdini & Linda Demaine, *The Expanding, Lop-Sided Universe of Social Influence and Law Research*, *in* THE OXFORD HANDBOOK OF SOCIAL INFLUENCE 395, 409 (Stephen G. Harkins, Kipling D. Williams & Jerry M. Burger eds., 2016) (arguing that governments use law as a means of shaping behavior in desired ways for societal stability).

113 *See, e.g.*, Cass R. Sunstein, *Social Norms and Social Roles*, 96 COLUM. L. REV. 903 (1996) (arguing that legal pronouncements can trigger norm shifts by signaling what society deems acceptable); ERIC A. POSNER, LAW AND SOCIAL NORMS 5 (2000) (finding that concern with cooperation leads individuals to engage in imitative behaviors that Posner considers social norms).

114 576 U.S. 644 (2015).

115 Margaret E. Tankard & Elizabeth Levy Paluck, *The Effect of a Supreme Court Decision Regarding Gay Marriage on Social Norms and Personal Attitudes*, 28 PSYCH. SCI. 1334, 1341–42 (2017).

116 142 S. Ct. 2228 (2022); *see* Chelsey S. Clark, Elizabeth Levy Paluck, Sean J. Westwood, Maya Sen, Neil Malhotra & Stephen Jessee, *Effects of a US Supreme Court Ruling to Restrict Abortion Rights*, 8 NATURE HUM. BEHAV. 63, 65–68 (2024) (charting the effects of *Dobbs* on public attitudes).

research and drafting, participants perceived the opinion as significantly lower in moral authority, less capable of shaping moral beliefs, and less connected to prevailing social norms.[117] Even when AI was used solely for research, leaving human judges in control, participants still reported an erosion of law's halo, albeit to a lesser degree.[118]

The trend extends beyond adjudication. Parallel experiments involving legislative lawmaking revealed similar declines in perceived legitimacy as machine authorship increased.[119] In these experiments, the more central AI became in the creation of legal text, the more the law's perceived normative force diminished. Crucially, these shifts in perception occurred despite textual identicality. That is, participants evaluated the same legal opinion differently based solely on the mode of its production. This result suggests that something beyond content governs the law's authority in the public imagination. Authorship and, more precisely, human authorship, emerges as a constitutive feature of perceived legitimacy. These results are especially important when thinking about potential objections to the existence or threat of a hard problem of ALI.

A pragmatically grounded critique might challenge the hard problem by pointing to the absence of systemic erosion in legal authority across jurisdictions that have embraced AI. In other words, we already have AI, and the sky in the legal world has not fallen. Technologies ranging from predictive analytics and automated discovery to generative contract drafting are increasingly ubiquitous in both judicial and commercial settings.[120] Yet their integration has not been accompanied by a measurable decline in institutional legitimacy or public trust. Or, more accurately, where legitimacy deficits have been observed, they often are associated more strongly with exogenous sociopolitical variables—such as polarization and democratic backsliding[121]—than with the use of AI per se. Thus, the claim

---

117 Avery, *supra* note 101, at 939–42 (finding that respondents rated human-authored legal decisions as having greater moral authority and normative influence, with entirely human-authored opinions achieving the highest authority score).

118 *Id.* at 946–47 (finding rulings without AI were most trusted, followed by AI-assisted and then AI-automated rulings).

119 *Id.* at 951–54 (replicating the earlier experiment in legislative decision-making and finding statutes drafted without AI were most trusted, again followed by AI-assisted and then AI-automated statutes).

120 *See supra* Part II.

121 *See, e.g.*, Daniel de Visé, *The American Public No Longer Believes the Supreme Court Is Impartial*, THE HILL (Jan. 11, 2023, at 06:00 ET), https://thehill.com/regulation/court-battles/3807849-the-american-public-no-longer-believes-the-supreme-court-is-impartial [https://perma.cc/NT6T-XCPR] (noting that a majority of Americans believed that the Supreme Court let partisan views influence major rulings); John F. Harris & Ian Ward, *Left and Right Agree on One Thing: The Justice System Is Corrupted by Bias*, POLITICO (July 1, 2024, at 05:00 ET), https://www.politico.com/news/2024/07/01/justice-system-bias-supreme-court-00165991 [https://perma.cc/GYT6-PEFG] (describing the same).

that AI introduces a qualitatively distinct threat to law's legitimacy finds little support in the empirical record. As a rebuttal to this objection, I offer a note of caution: We remain at an early stage of AI adoption, particularly in legally sensitive and publicly visible decision-making contexts. Many deployments are opaque, incremental, or insulated from public scrutiny. Whether legitimacy effects emerge as AI becomes more central, transparent, and unavoidable within legal institutions remains an open question.

Another objection to the hard problem might emphasize that current AI systems operate as subordinate instruments rather than as autonomous legal agents. In prevailing practice, AI tools support—not supplant—human legal actors, functioning analogously to law clerks, precedent databases, or analytical aids. Decision-making authority remains decisively human.[122] Within this paradigm, AI's outputs are mediated by the judgment and editorial discretion of legal professionals, thereby preserving the attribution of authorship and maintaining accountability. Research indicates that careful implementation of human-in-the-loop designs can diminish the risks of cognitive distortion and other pitfalls.[123] Far from displacing human legal agency, AI may reinforce it by bolstering the procedural integrity and reliability of legal practice. Here, however, the experimental evidence discussed above largely proves the objection unfounded. Even when AI's role was formally subordinate, public perceptions of law's halo were measurably diminished.[124] These findings suggest that delegation, even short of full automation, may be sufficient to disrupt law's special role.

Taken together, the empirical findings in this Section deepen the plausibility of the hard problem. They indicate that law's normative halo may not be fully portable to machine-generated or machine-mediated legal texts, even when those texts are substantively impeccable and institutionally authorized. To be sure, this evidence represents only an initial step in what

---

[122] A.B.A. Comm. on Ethics & Pro. Resp., Formal Op. 512 (2024), https://www.americanbar.org/content/dam/aba/administrative/professional_responsibility/ethics-opinions/aba-formal-opinion-512.pdf [https://perma.cc/9ZCA-8FV8] (stating that generative AI tools may assist lawyers with tasks such as research and drafting but may not replace lawyers' independent professional judgment or ethical responsibility); LEXISNEXIS, GENERATIVE AI & THE LEGAL PROFESSION: 2023 SURVEY REPORT 12 (2023), https://www.lexisnexis.com/assets/en-hk/pdf/ai-hub-gai-legal-profession.pdf [https://perma.cc/SFA9-MFDA] (reporting that legal professionals view generative AI primarily as a tool to support human lawyers' work rather than to supplant human legal decision-making).

[123] *See generally* Helena Vasconcelos, Matthew Jörke, Madeleine Grunde-McLaughlin, Tobias Gerstenberg, Michael S. Bernstein & Ranjay Krishna, *Explanations Can Reduce Overreliance on AI Systems During Decision-Making*, 7 PROCS. ACM ON HUM.-COMPUT. INTERACTION, no. 129, 2023, at 1 (empirically showing that simple explainable AI, which presents users with model-generated rationales for outputs, is less effective at reducing overreliance than decision designs incorporating cognitive forcing functions).

[124] Avery, *supra* note 101, at 942–50.

must become a broader research agenda.[125] But it raises a foundational question that neither pragmatic reassurance nor instrumentalist accounts can fully answer: Can law retain its full character when its expressive agents are not persons? For now, the answer remains unsettled. What the data cautions against, however, is any assumption that public perceptions of legitimacy will automatically transfer from human to machine, no matter how competent the latter becomes.

CONCLUSION

In 2022, Kevin Tobia introduced the field of experimental jurisprudence, demonstrating how empirical tools could invigorate legal theory by shedding light on ordinary cognition, concept usage, and the psychological foundations of jurisprudence.[126] Tobia's project was not to replace jurisprudence but to expand its methods.[127] His call to build a more sophisticated and empirically informed account of law's conceptual terrain[128] has since catalyzed a robust interdisciplinary literature.[129]

ALI now presents a parallel imperative. The emergence of large language models, generative legal tools, and automated adjudication systems has forced the legal academy to confront a pressing question: Can machines make law? More precisely, can their outputs possess the same moral authority, persuasive force, and legitimacy that characterize human-authored legal decisions? The hard problem of ALI distills this question into its sharpest form. And just as experimental jurisprudence matured through the integration of cognitive science and empirical methods, the inquiry into ALI must develop as a field in its own right, melding jurisprudence, philosophy of mind, AI theory, and empirical social science.

The hard problem will not be settled by theoretical pronouncement alone. It demands sustained empirical engagement, doctrinal analysis, and conceptual refinement. What is needed now is a research agenda capacious enough to evaluate ALI not only in terms of what machines can do, but also

---

125 After all, the empirical basis for the hard problem remains limited. The studies discussed here are few in number, rely primarily on survey methods, and involve modest sample sizes and constrained jurisdictions, which may limit the generalizability of their normative implications.

126 Kevin Tobia, *Experimental Jurisprudence*, 89 U. CHI. L. REV. 735, 735 (2022); *see also* Roseanna Sommers, *Experimental Jurisprudence*, 373 SCIENCE 394, 394 (2021) (introducing and explaining the field).

127 Tobia, *supra* note 126, at 800–02.

128 *Id.*

129 *See generally* Joshua Knobe & Scott Shapiro, *Proximate Cause Explained: An Essay in Experimental Jurisprudence*, 88 U. CHI. L. REV. 165 (2021) (applying methods from philosophy, psychology, and linguistics to the law of proximate cause); Christophe Heintz & Thom Scott-Phillips, *Expression Unleashed: The Evolutionary and Cognitive Foundations of Human Communication*, 46 BEHAV. & BRAIN SCIS., no. e1, 2023, at 1, 13 (discussing newly emerging interdisciplinary fields, including experimental jurisprudence, and the growing bridges between law, cognitive science, and evolutionary theory).

in terms of what law is when mediated by those machines.

Several lines of inquiry are especially urgent. Large-scale, cross-cultural experiments are needed to measure public reactions to AI-generated legal decisions. These should track not just perceived fairness or accuracy, but also deeper perceptions of moral authority, normativity, and institutional trust. Comparative designs across jurisdictions and legal traditions will help determine whether AI-induced legitimacy gaps are universal or context-dependent. Research must also trace how AI tools are adopted over time within courts, legislatures, and administrative agencies. Rather than speculate abstractly, scholars should examine how AI integration interacts with procedural justice, legal pluralism, and professional norms on the ground.

Future work should clarify whether new ethical paradigms are needed to govern the use of AI in law. What forms of delegation are normatively permissible? What conditions must be met for AI authorship to retain legitimacy? Can procedural safeguards fully compensate for the absence of human judgment? Deeper theoretical analysis is also needed on the ontological and epistemic status of AI-generated law. Does legal reasoning entail irreducibly human faculties? Or can machines approximate the conditions of legal interpretation through complexity modeling, learning architectures, or probabilistic inference? Finally, ALI must be understood as a global phenomenon. Different legal cultures—common law, civil law, pluralist systems—will pose distinct challenges for AI integration. Future research must avoid jurisdictional parochialism and instead embrace a comparative lens sensitive to power, access, and systemic bias.

Together, these paths can help determine whether the hard problem is a fixed boundary or a dissolving horizon, whether machines will always remain outsiders to legal authorship, and whether legitimacy may prove more malleable than imagined. ALI, as a field, is still young. But so was experimental jurisprudence not long ago. With the same blend of conceptual rigor, empirical humility, and interdisciplinary openness, ALI can evolve from a set of isolated questions into a generative program of legal scholarship, one capable of confronting law's most urgent questions in the age of intelligent machines.