Who Learns to Think Like a Lawyer? AI-Assisted Practice and the Erosion of Professional Apprenticeship

Stevie Michelle Cline, Esq.*

Arcana Research, San Francisco, California, USA, stevie@steviecline.com

The legal profession is adopting AI tools at speed, and the conversation has focused almost entirely on whether those tools produce accurate, efficient output. A different question deserves attention: what happens to the pipeline through which legal expertise is developed? Law remains an apprenticeship profession. Competence is built through years of effortful, iterative practice under supervision. This paper argues that AI-assisted legal work disrupts three foundational mechanisms of expertise acquisition (constructive drafting, associative research, and calibration under uncertainty) and that this disruption exposes a gap in professional responsibility frameworks that presuppose the continued production of expert practitioners.

Additional Keywords and Phrases: legal expertise, deskilling, cognitive offloading, professional apprenticeship, AI-assisted work, professional responsibility

1 INTRODUCTION

Large language models now draft contract provisions, generate legal memoranda, summarize case law, and flag regulatory risks. The dominant framing of this transformation emphasizes productivity. In a recent randomized controlled trial, Choi, Monahan, and Schwarcz found that law students with access to GPT-4 completed legal tasks substantially faster, though quality improvements were slight and inconsistent [1]. Their most striking finding may be the one least discussed: the lowest-skilled participants saw the largest quality gains from AI assistance. AI made weak performers look adequate. Katz et al. demonstrated that GPT-4 passes the Uniform Bar Examination, outperforming the average human test-taker on multiple components [2]. These results are significant, but they share a premise: they treat the practicing attorney as a fixed quantity, an expert whose competence exists independent of the tools she uses. They ask whether AI is good enough to assist the lawyer. They do not ask whether the lawyer will remain good enough to supervise the AI.

That second question implicates a different set of concerns. Legal competence is not acquired through coursework alone. It is built through an apprenticeship model in which junior attorneys learn by performing cognitively demanding work (drafting, researching, analyzing) under the

** Stevie Michelle Cline, Esq. is an AI governance attorney who has served as General Counsel at four technology companies, with two successful exits. She began her career in capital markets at a major New York firm, where she worked on nearly thirty IPOs and early digital assets offerings. She is admitted to the New York State Bar, the U.S. Supreme Court Bar, and the U.S. Tax Court, and her AI law scholarship has been accepted at the International Journal of Law and Emerging Technologies and the University of Illinois Journal of Law, Technology & Policy. The views expressed in this paper are the author's own and do not represent those of any employer or client.

supervision of experienced practitioners. The errors, the redlines, the iterative refinement are not inefficiencies to be optimized away. They are the mechanism through which expertise is produced. When AI systems automate the effortful practice that develops professional judgment, the profession risks a structural erosion of its own capacity for self-reproduction.

This paper identifies three such mechanisms and argues that existing professional responsibility frameworks cannot address the resulting gap.

2 HOW LEGAL EXPERTISE IS BUILT

The dominant model of legal training is apprenticeship. Law school provides doctrinal foundations, but the core competencies of practice (issue-spotting, strategic judgment, persuasive writing, client counseling) are developed on the job over years of supervised work.

Ericsson, Krampe, and Tesch-Römer's framework of deliberate practice offers a useful lens [3]. Expertise in complex domains is built through sustained engagement with tasks that exceed the practitioner's current ability, accompanied by immediate feedback and opportunities for correction. In legal practice, this looks familiar: a junior associate drafts a memorandum, a senior attorney marks it up extensively, and the associate revises. The cognitive work is in the construction, not the consumption. An associate who has written forty draft motions to dismiss has developed pattern recognition and strategic intuition that cannot be acquired by reading forty exemplars, however excellent.

Chi's active-constructive-interactive framework reinforces this distinction [4]. Producing a work product recruits different cognitive processes than reviewing one. Construction requires the practitioner to retrieve and organize knowledge, identify gaps, and commit to positions under uncertainty. Review operates on material already structured by someone (or something) else.

The apprenticeship model also produces what I will call *doctrinal geography*: an associative map of legal concepts and authorities that experienced attorneys navigate intuitively. A senior litigator does not just know the holding of *International Shoe*; she knows how it sits relative to *Asahi*, where the circuit split lies, and which arguments will resonate with a particular judge. This map is built through the slow process of manual legal research, following citation chains, reading adjacent cases, encountering unexpected connections between bodies of law. It is a form of situated knowledge [5] that underlies the professional judgment distinguishing competent practice from mechanical rule application.

3 THREE MECHANISMS OF DISRUPTION

AI tools intervene at the points where effortful practice produces expertise.

3.1 Draft Generation as Cognitive Bypass

When an AI system generates a competent first draft of a legal memorandum, the junior attorney's role shifts from constructor to reviewer. She no longer confronts a blank page, no longer struggles with structure, no longer makes (and learns from) the characteristic errors of early practice. Instead, she evaluates a pre-formed work product she did not build.

This creates a bootstrapping problem. The ability to critically review a legal document presupposes the expertise that drafting was supposed to develop. An experienced attorney can spot when an AI-generated memorandum has missed a controlling authority or subtly mischaracterized a holding, because she has made those errors herself and been corrected. A junior attorney operating as a reviewer from the outset lacks the experiential basis for that judgment.

The Choi et al. finding makes this harder to detect: if AI disproportionately improves the output of the least-skilled practitioners, the deskilling is masked until the AI-dependent attorney must function without it or train someone else.

The risk is not that AI drafts are poor. A competent AI draft is, from a learning perspective, worse than no draft at all.

3.2 Research as Pattern Erosion

AI-assisted legal research tools return answers: relevant cases, applicable statutes, synthesized summaries. They are optimized to minimize the time between query and result. But traditional legal research builds something beyond the immediate answer.

Consider a junior attorney researching a novel question of personal jurisdiction. The manual process forces her to work through the doctrine: she reads the canonical cases, follows the evolving treatment of internet-based contacts, and begins to understand how courts in her jurisdiction have positioned themselves within ongoing debates. An AI research tool that returns the three most relevant cases and a synthesized summary delivers the answer without building the map.

Kirschner, Sweller, and Clark have cautioned that reducing cognitive load is not always pedagogically beneficial [6]. In expertise development, the cognitive effort of unguided exploration is constitutive of the knowledge being acquired. Efficient retrieval serves the expert who already possesses the doctrinal map; it deprives the novice of the process by which the map is built.

3.3 Calibration Without Reps

A critical component of legal expertise is calibration: the ability to assess the strength of one's own position and to advise clients under genuine uncertainty. This capacity is developed through repeated exposure to the gap between prediction and outcome. The motion the attorney expected to win but lost. The settlement posture that misjudged opposing counsel's resolve.

AI systems distort this process. When an AI tool presents a legal analysis, it arrives with what I will call implicit confidence: clean formatting, citation density, absence of hedging, and a completeness of structure that signals authority regardless of whether the underlying reasoning is sound. Parasuraman and Manzey's work on automation bias demonstrates that decision-makers systematically over-rely on automated recommendations, even when those recommendations are unreliable [7]. In legal practice, these interface signals mean the AI-assisted attorney may have fewer encounters with genuine, unmediated uncertainty. The AI's outputs smooth over the ambiguity that is the raw material of professional judgment.

For junior attorneys who have not yet established an independent calibration baseline, this compounds. Without sufficient experience of unassisted uncertainty, they risk developing a calibration anchored to AI-generated assessments rather than to the actual indeterminacy of legal practice.

4 THE GOVERNANCE GAP

Existing professional responsibility frameworks were built for a world in which the apprenticeship pipeline reliably produces competent practitioners. These frameworks regulate the *exercise* of expertise but do not address the conditions necessary for expertise to be *developed*.

The ABA Model Rules of Professional Conduct impose a duty of competence (Rule 1.1), requiring the legal knowledge, skill, thoroughness, and preparation reasonably necessary for the representation [10]. Rule 5.1 establishes supervisory obligations for senior attorneys. Both rules assume a background condition: the professional training pipeline will produce attorneys capable of meeting the competence standard, and supervised practice will progressively develop the judgment necessary for independent work.

Comment [8] to Rule 1.1, added in 2012, provides that a lawyer should keep abreast of changes in the law and its practice, "including the benefits and risks associated with relevant technology." The profession has created an affirmative duty to adopt AI tools while providing no framework for the expertise-development risks that adoption creates. The duty of technological competence and the conditions for professional competence are on a collision course.

If AI-assisted practice attenuates the mechanisms by which competence is acquired, these frameworks face a foundational challenge. A duty of competence loses its force if the profession's training structures no longer reliably produce competent practitioners. A duty of supervision is inadequate if the supervised work itself has been hollowed out, if the associate is reviewing AI outputs rather than engaging in the constructive practice that supervision was designed to support. And the economic incentives cut the wrong way: firms adopt AI to reduce billable associate hours and lower cost to clients, which means the financial pressures driving adoption are directly opposed to the conditions that develop expertise.

Aviation safety offers a structural parallel: Ebbatson et al. demonstrated that pilots who rely extensively on automated flight systems exhibit diminished manual flying proficiency, with performance tied to recent hands-on experience rather than total career hours [8]. The regulatory response has included guidance requiring opportunities for manual flight and periodic proficiency assessments [9]. The legal profession has no analogous mechanism.

The gap compounds over time. An AI-assisted junior attorney may perform adequately today, but she may never develop the judgment to practice independently or train the next generation. The consequences will manifest in the next cohort, when the attorneys who should be supervising were themselves produced by an attenuated pipeline. For the HCI and AI design communities, this suggests the need for expertise-differentiated design: systems that calibrate their assistance to the user's developmental stage, rather than optimizing uniformly for task completion. The challenge generalizes beyond law to any apprenticeship profession, but the legal domain is instructive because its governance frameworks make these assumptions unusually explicit.

REFERENCES

[1] Jonathan H. Choi, Amy B. Monahan, and Daniel Schwarcz. 2024. Lawyering in the Age of Artificial Intelligence. Minnesota Law Review 109, 1 (2024), 147–218.

[2] Daniel Martin Katz, Michael James Bommarito, Shang Gao, and Pablo Arredondo. 2024. GPT-4 Passes the Bar Exam. Philosophical Transactions of the Royal Society A 382, 2270 (2024), 20230254.

[3] K. Anders Ericsson, Ralf Th. Krampe, and Clemens Tesch-Römer. 1993. The Role of Deliberate Practice in the Acquisition of Expert Performance. Psychological Review 100, 3 (1993), 363–406.

[4] Michelene T.H. Chi. 2009. Active-Constructive-Interactive: A Conceptual Framework for Differentiating Learning Activities. Topics in Cognitive Science 1, 1 (2009), 73–105.

[5] Hubert L. Dreyfus and Stuart E. Dreyfus. 2005. Peripheral Vision: Expertise in Real World Contexts. Organization Studies 26, 5 (2005), 779–792.

[6] Paul A. Kirschner, John Sweller, and Richard E. Clark. 2006. Why Minimal Guidance During Instruction Does Not Work: An Analysis of the Failure of Constructivist, Discovery, Problem-Based, Experiential, and Inquiry-Based Teaching. Educational Psychologist 41, 2 (2006), 75–86.

[7] Raja Parasuraman and Dietrich H. Manzey. 2010. Complacency and Bias in Human Use of Automation: An Attentional Integration. Human Factors 52, 3 (2010), 381–410.

[8] Matt Ebbatson, Don Harris, John Huddlestone, and Rodney Sears. 2010. The Relationship Between Manual Handling Performance and Recent Flying Experience in Air Transport Pilots. Ergonomics 53, 2 (2010), 268–277.

[9] Federal Aviation Administration. 2013. Safety Alert for Operators: Manual Flight Operations. SAFO 13002. U.S. Department of Transportation, Washington, DC.

[10] American Bar Association. 2020. Model Rules of Professional Conduct. American Bar Association, Chicago, IL.