APPENDIX C: Cowork in Practice: A Research Library Analysis

The analysis below comes from a single Cowork session.

That is not something a regular chat session handles well. In chat, you upload files one at a time, the context window accumulates, and by the time you have loaded a substantial collection, the model is losing track of what came first. Cowork sidesteps this by accessing the folder directly and working through documents iteratively, rather than trying to load everything into a single prompt.

From that one session, Cowork produced a five-theme synthesis with supporting evidence drawn from across the collection and cross-cutting observations about the field as a whole. The session ran for about twelve minutes.

How to Run a Similar Task

Before you start: you need Claude Desktop installed (Mac or Windows) and a paid Claude subscription — Pro, Max, Team, or Enterprise. Cowork is not available through the browser or the mobile app.

Step 1. Open Claude Desktop.

The app shows three tabs at the top: Chat, Cowork, and Code. Click Cowork.

Step 2. Start a new task.

You will be prompted to grant folder access before Cowork takes any action.

Step 3. Designate your folder.

Navigate to the folder you want Cowork to work with and select it. Cowork reads and writes only within folders you explicitly authorize. If you have sensitive files you do not want included, move them to a separate folder before the session. Do this first before you start rather than trying to manage mid-session.

Step 4. Describe your goal.

For this example, the prompt was close to: *"Please analyze the contents of this folder and produce a written summary of the major themes and patterns across the collection. Include specific supporting examples from the documents."*

Be specific about what you want. Naming a format (a report, a memo, an annotated reading list) and a rough length saves you a revision cycle. Cowork handles open-ended requests, but it will make its own choices about structure and depth if you leave them unspecified — and its choices may not be yours.

Step 5. Review Cowork's plan.

Cowork reads the folder, proposes an approach, and surfaces that plan before taking any action. Read it. If the plan is heading somewhere unhelpful — wrong scope, wrong angle, wrong output format — redirect it before you approve. This is your leverage point.

Step 6. Approve and monitor.

Once you approve, Cowork runs. The interface shows what it is reading and doing in real time. You can interrupt and redirect at any point.

Step 7. Review the output.

Cowork will save the deliverable to your designated folder. Read it critically. The output below is a solid synthesis — and it is also mine to verify, correct, and use as I see fit. Cowork does not know what you already know, and it does not check its own work against your judgment.

A Note on What This Is Not For

Cowork is in research preview, and Anthropic says so directly. It is not designed for regulated workloads, it does not appear in compliance exports or audit logs, and a prompt injection vulnerability (where malicious content embedded in a document could redirect Cowork's actions mid-session) has been identified and partially addressed but not fully resolved. Keep Cowork away from client files, clinic materials, externship documents, and anything subject to confidentiality obligations.
It works well with your own research, course materials, and publicly available sources. See below for its output based on my prompt.

AI in Legal Education

Patterns and Insights Across 110 Sources

A Thematic Analysis of the Legal Education Research Collection

---

March 2026

Korin Munsterman

About This Collection

This analysis spans 110 files organized across 28 thematic subfolders, representing a working research library on the intersection of generative AI and legal education. The collection draws from law review scholarship, practitioner surveys, institutional policy documents, journalism, pedagogical frameworks, and course materials. It ranges in date from early 2024 through March 2026, capturing both the initial wave of GenAI-in-law-school commentary and the more recent reckoning with what that wave has — and has not — produced.

File types include 82 PDFs, 14 Word documents, 12 text transcripts, one Excel file, and one presentation. Thematic folders cover adoption and the workforce gap, agentic AI, classroom policies, assignments and rubrics, the bar exam, cheating and academic integrity, grading, AI for legal research, libraries, scholarship and authorship, surveys, training, and use cases — alongside a growing body of original course materials built around the AI Fluency Framework.

The date range runs from approximately February 2024 to March 2026, with a notable clustering of scholarship added in late February 2026 (focused on disclosure, plagiarism, and AI authorship) and the most recent addition — a TaxProf Blog post on agentic AI and law school deans — dated March 2, 2026. The collection thus captures a field in rapid motion, with its center of gravity visibly shifting from "should we integrate AI?" to "what happens now that we can't avoid it?"

Theme 1: The Law School Response Gap

Industry demand for AI-competent graduates continues to grow. Law school course offerings are shrinking.

The most consistently documented pattern across this collection is a widening mismatch between what the legal profession expects from new graduates and what law schools are actually teaching. Multiple independent data sources converge on the same finding: despite rising employer demand for AI competency, law schools have not increased — and in some measures have actually reduced — their AI-related course offerings.

Bloomberg Law's analysis of successive Path to Practice surveys provides the sharpest evidence. Between the 2023–2024 and 2024–2025 school years, faculty reporting that their schools offered classes on the legal implications of AI dropped from 73% to 51%. Faculty reporting practical AI skills courses dropped from 68% to 42%. Meanwhile, the percentage of faculty who said their institution had no AI offerings, or that they weren't sure, nearly doubled. Student participation reflects the same pattern: only 9% of students reported taking a legal technology course, 8% a practical AI skills course, and 6% an AI literacy course — all well below what course availability would predict.

Supporting Examples

Bloomberg Law / Path to Practice 2026 survey: Four in ten surveyed students reported their classes helped them gain proficiency in AI-driven legal tools. Only 11% of faculty said their law school requires professors to take AI training. Attorneys surveyed believed AI is growing faster than they can keep up — and half believe it will reduce opportunities for graduating students.

LexisNexis/ALITA survey (Asia-Pacific, 2025): Only 15% of surveyed lawyers believed law schools should stay the same. Nearly three-quarters called for curriculum reform specifically to "better train future lawyers in AI." Corporate in-house counsel were even more insistent, with 85% endorsing reform.

Bloomberg Law, "AI and the Coming Mismatch Between Law Schools and Law Firms" (Feb. 2026): Law school applications are up 17–18% even as admissions consultants ask the uncomfortable question: are schools admitting more students into a market that AI is already contracting? The piece documents firms planning to reduce junior associate hiring — not by laying off current associates, but by simply hiring fewer going forward.

What makes this gap notable is its persistence. The field has been talking about AI integration in legal education since at least late 2022. By 2025–2026, the conversation has produced plenty of think pieces, some elective courses, and a handful of ambitious institutional initiatives (WashU Law's publicly stated goal of becoming "the global leader in AI education and training" being the highest-profile example). What it has not produced, according to the survey data, is systemic change at the curricular level. The gap between rhetoric and course catalog is itself a theme worth naming.

Theme 2: The Disclosure Paradox

The field has settled on disclosure as the ethical answer to AI use. A cluster of recent scholarship suggests disclosure creates as many problems as it solves.

Your most recent scholarship additions — clustered in a single afternoon of downloading in late February 2026 — collectively document a problem that the emerging AI-in-legal-education consensus has not reckoned with: mandatory AI disclosure, the near-universal recommendation of ethics committees and institutional policy documents, has significant and underappreciated costs.

The empirical case is made most directly by Schilke and Reimann's "The Transparency Dilemma: How AI Disclosure Erodes Trust" (Organizational Behavior and Human Decision Processes, 2025). Across thirteen experiments involving supervisors, subordinates, professors, analysts, and investment funds, disclosing AI use consistently caused evaluators to trust the disclosing actor less than those who did not disclose — regardless of whether the disclosure was voluntary or mandatory, regardless of whether

AI involvement was already known, and regardless of how the disclosure was framed. The mechanism the authors identify is legitimacy: disclosure signals that the person needed help in a domain where self-sufficiency is expected.

The disclosure problem looks different depending on which corner of the scholarship you are reading. For legal scholars producing articles, Frazier and Rozenshtein ("Large Language Scholarship," 20 FIU L. Rev., forthcoming 2026) argue that mandatory disclosure norms are both practically unenforceable and theoretically misguided — the AI tools are part of the scholarly workflow now, and the relevant question is authorial accountability, not disclosure mechanics. For producers of creative and professional content, Noti-Victor ("Regulating Hidden AI Authorship," 111 Va. L. Rev. 139, 2025) documents that copyright law creates active financial incentives to conceal AI involvement, meaning disclosure requirements cut directly against authors' economic interests. And for students and faculty submitting academic work, Lemley and Ouellette ("Plagiarism, Copyright, and AI," U. Chi. L. Rev. Online, 2025) argue that disclosure of AI use doesn't solve the underlying plagiarism concern because the humans whose ideas AI has synthesized still receive no credit — disclosure tells the reader the author used a tool; it doesn't trace the intellectual debts that tool created.

Supporting Examples

Jabotinsky & Sarel, "Co-Authoring with an AI?" (56 Ariz. St. L.J. 187): After surveying major publishers' policies — ranging from outright bans on AI-generated text to disclosure requirements to a "mixed approach" — the authors found that law reviews occupy the most permissive (or perhaps simply the most neglectful) category: no guidelines whatsoever. Law reviews have largely not yet established standard norms for AI use in legal writing.

Schilke & Reimann, "The Transparency Dilemma" (2025): Perhaps most counterintuitively, the study found that being caught using undisclosed AI has a worse trust effect than proactive disclosure — but that proactive disclosure still produces a meaningful trust penalty compared to no AI use at all. The authors conclude that "transparency is not straightforwardly beneficial," which is not a sentence that appears anywhere in institutional AI policy documents.

Korin's own research conversation (Feb. 27, 2026): A research session documented in the Scholarship folder explored whether universities have created policies on GenAI and academic tenure — probing the question of whether AI-assisted scholarship should be evaluated differently for promotion purposes. The research found that institutional policies on AI in scholarship remain sparse and inconsistent, and that the analogy to judicial disclosure rules (Bluebook citation requirements for AI use) is both apt and underdeveloped.

The practical implication for legal education is real. If educators encourage students to disclose AI use as a matter of professional habit — a reasonable and defensible pedagogical goal — they are simultaneously asking students to accept a trust penalty in professional environments that have not yet normalized that disclosure. Teaching disclosure as a value means preparing students to navigate the friction that comes with it.

Theme 3: What Should Actually Be Taught — And Nobody Agrees

Faculty agree students need AI education. Faculty disagree about what that means, and most feel underprepared to provide it.

Across survey data, scholarship, and practitioner commentary, a clear consensus emerges that legal education must change in response to AI. What sits underneath that consensus — agreement about content, sequence, and method — is considerably murkier. The collection documents at least three competing views of what AI education in law school should actually do, none of which has carried the day.

The first view holds that AI education should be primarily about tool proficiency: students need to know how to use Harvey, Westlaw's AI features, contract review tools, and the general-purpose models that will show up in firm workflows. This view is most prominent in practitioner surveys and the employer-demand literature. It is the view that generates headlines about "students graduating without knowing how to use AI."

The second view — more prevalent in the pedagogical scholarship — holds that tool proficiency is the wrong target, because the tools change too fast. What students need is critical evaluation: the capacity to identify when AI output is wrong, incomplete, or inappropriate, and to maintain professional judgment even when AI is handling the drafting. Bliss's "Teaching Law in the Age of Generative AI" (Jurimetrics, forthcoming) exemplifies this view, arguing for AI-integrated teaching that develops students' evaluative skills alongside — not instead of — foundational legal reasoning.

The third view, represented by the "Gut Instinct in the Age of AI" piece (Above the Law, July 2025), questions whether the first two views adequately account for what is actually lost when junior lawyers stop doing unglamorous foundational work. The worry is not that students will graduate unable to use AI, but that they will graduate unable to exercise the judgment that comes from having done the work AI is now doing for them.

Supporting Examples

Schrepel, "GenAI in Legal Education: A Two-Year Experiment with ChatGPT" (Aug. 2025): Structured training (explicit instruction in how to use AI effectively) outperformed both prohibition and unstructured exposure in year one of the experiment. By year two, the advantage had diminished as students arrived with greater baseline familiarity. The study recommends against one-size-fits-all faculty mandates and in favor of professor-level experimentation — a finding that validates pluralism but does not resolve what the content of good AI training actually is.

Bliss, "Teaching Law in the Age of Generative AI" (Jurimetrics): A national faculty survey found that even professors who strongly agreed students need AI preparation reported feeling "uninformed about this technology and unsure how to proceed." The article concludes that the limiting factor in AI integration is not faculty resistance but faculty capacity — professors need training before they can provide it.

The AI Fluency Framework (Teaching AI Fluency course, Korin Munsterman): The 4D Framework — Delegation, Description, Discernment, Diligence — represents a structured answer to the "what to teach" question. By organizing AI fluency around competencies rather than tools, it sidesteps the problem of tool obsolescence and provides a transferable framework applicable across disciplines and practice areas. The Teaching AI Fluency transcripts document four distinct pedagogical entry points: linear, non-linear, focused (single-D), and the "two loops" model that distinguishes strategic from tactical AI decision-making.

Theme 4: Academic Integrity in a State of Policy Chaos

Prohibition doesn't work. Detection doesn't work. A coherent integrity framework hasn't arrived yet.

This collection contains more documents directly concerned with academic integrity than any other single topic: seven individual law school policy documents, a general AI policies framework from a leading scholar at Emory, a 2025 anti-cheating guide for higher education, and multiple scholarship articles addressing plagiarism, copyright, and the definitional challenges AI creates for conventional academic honesty standards. The sheer volume is a symptom. Law schools know they have an integrity problem. They have not agreed on how to address it.

The policy documents themselves — from Columbia, Stanford, Duke, UChicago, William & Mary, Loyola, and USD — reveal the range of institutional responses: prohibition, limited disclosure requirements, context-specific permissions, and various hybrid approaches. Sag's "AI Policies for Law Schools" (2024) provides the most systematic critique of this landscape, arguing that policies which attempt to prohibit AI use entirely suffer from three interlocking failures: they are pedagogically unsound (preparing students for a world that no longer exists), practically unenforceable (detection tools are unreliable and equity concerns loom large), and professionally counterproductive (discouraging the supervised AI use that students will need to demonstrate on day one of practice).

The plagiarism question is particularly thorny for legal education, and the collection captures why. The "Language Models, Plagiarism, and Legal Writing" piece makes the deliberately provocative argument that law has always tolerated — indeed required — extensive copying from templates, prior work, and opposing counsel's language. Teaching students to condemn AI-assisted copying while training them in a profession that routinely engages in professionally sanctioned copying is, the author argues, inconsistent. Lemley and Ouellette's more measured analysis distinguishes copyright infringement (which AI use typically does not constitute) from plagiarism (which it often does, in the specific sense of passing off others' ideas without attribution) — a distinction that most institutional policies paper over.

Supporting Examples

Sag, "AI Policies for Law Schools": Proposes a framework that distinguishes prohibited uses (copying AI text verbatim, submitting AI-generated citations without verification), permitted uses with disclosure (editorial assistance and grammar review), and freely permitted uses (study aids, comprehension support). Argues that effective policies must be course-specific and focused on professional preparation, not prohibition.

Cyphert, "Generative AI, Plagiarism, and Copyright Infringement in Legal Documents" (25 Minn. J.L. Sci. & Tech. 49, 2024): Analyzes the specific IP issues that arise when lawyers use LLMs in practice — including whether GenAI output can expose attorneys to copyright liability if models reproduce copyrighted text verbatim. Concludes that while outright copyright liability is unlikely, the risk is not zero, and lawyers need to understand the technical parameters of the tools they use.

AI Assignments and Rubrics (Assignments folder): The rubric design document in your collection embeds integrity considerations into assessment design rather than treating them as a separate enforcement problem. It recommends requiring process artifacts (prompt logs, AI use statements), designing assessments that make AI use visible and productive rather than hidden and penalized, and restructuring rubric criteria to evaluate judgment alongside output quality.

Theme 5: The Agentic AI Inflection Point

The field has been preparing for AI as a tool. Agentic AI operates more like a junior associate — and that changes the calculus entirely.

The most recent documents in the collection signal a shift in the underlying frame of the AI-in-legal-education debate. The first wave of GenAI commentary — represented by much of the 2024 material here — addressed AI as a productivity tool that augments lawyer work. The emerging frame, visible most clearly in the Rodriguez TaxProf Blog piece from February 2026 and implicitly in the "Mismatch" piece from the same month, positions AI not as a tool but as a potential replacement for the junior-associate tier of legal work.

Rodriguez's analysis is precise about the distinction. Generative AI follows a simple structure: prompt to output. Agentic AI follows a more complex chain: goal, plan, actions, feedback, revised plan, result. The practical difference is that an agentic system can be tasked with constructing a non-disclosure agreement that meets all jurisdictional requirements and accomplishes the client's objectives — not just drafting a clause when prompted. The human moves from being "in the loop" (monitoring each step) to "on the loop" (setting the goal and reviewing the final output). Rodriguez frames this as a dilemma for law school deans who have been acting as AI cheerleaders: the technology they have been celebrating may now be positioned to displace the entry-level roles their graduates are competing for.

Supporting Examples

Rodriguez, "Law School Deans as Cheerleaders, and the Delicate Marketing Dance, Part I: The Case of Agentic AI" (TaxProf Blog, Feb. 22, 2026): The core dilemma Rodriguez identifies is institutional: law school marketing has tied AI positively to graduate success ("AI will help your career"). Agentic AI complicates that narrative because it genuinely threatens to reduce the number of entry-level legal positions available to graduates. Deans who have been enthusiastic about GenAI now face a harder messaging challenge.

"AI and the Coming Mismatch Between Law Schools and Law Firms" (Bloomberg Law, Feb. 10, 2026): Documents the two camps among law firm leaders: those who expect AI to shrink junior hiring significantly (citing that 80% of first-year associate work — discovery, document review — is already automatable), and those who expect AI to generate more work just as email did. The piece notes that applications to law school are still rising sharply, meaning students are entering the pipeline without clear visibility into how the exit will look by the time they graduate.

Bar Examiner Letter from the Chair (NCBE, Summer 2025): Signals how even the credentialing infrastructure is being forced to reckon with AI. NCBE is studying whether AI might assist with item development and scoring for the NextGen UBE, while insisting on a "human-in-the-loop" approach and committing to at least one year's advance notice before any operational AI use. The parallel to legal education is direct: gatekeeping institutions are trying to integrate AI without destabilizing public trust in the standards they are designed to protect.

What makes this theme important for the collection as a whole is that it suggests the previous four themes may be oriented toward the wrong problem. If agentic AI is genuinely approaching the capacity to handle a substantial portion of entry-level legal work, then the question is not just how to integrate AI into legal education — it is what legal education is for in a world where the early-career developmental pipeline that has defined law practice for decades may be contracting.

Cross-Cutting Observations

Several patterns cut across all five themes and are worth naming explicitly.

The research-to-practice lag. The scholarship in this collection is substantially ahead of the institutional policies. Articles published in 2024–2025 document problems — with detection, with disclosure, with prohibition, with the skills gap — that most law school policies written in the same period have not yet acknowledged. The most sophisticated institutional policy frameworks (Sag's template, the assignment rubric design in your collection) are the exception, not the rule.

The faculty capacity constraint. Multiple sources identify the same bottleneck: faculty support the goal of AI integration but do not feel equipped to implement it. Only 11% of schools require faculty AI training. The Bliss survey and the Schrepel two-year study both

document that structured training produces better outcomes than unstructured exposure — which means the limiting factor in AI integration is not student resistance but faculty preparation. This points toward professional development as the highest-leverage investment law schools could make.

Your own work is filling a gap the research identifies. The AI Fluency Framework, Teaching AI Fluency course materials, and the AI Fluency in Disciplines resources in this collection represent exactly the kind of competency-based, transferable, faculty-deployable framework that the broader literature says is missing. The research-synthesis skill in your tools notes that the CALI book audience is legal educators who need practical implications, not just findings — and your materials are built precisely for that audience.

The field is still arguing about whether to integrate. A surprising number of documents in this collection are still making the foundational case that AI should be part of legal education at all. By 2026, with AI deeply embedded in every major legal research platform and firm workflow, that argument should be settled. That it is not — that "only 15% believe law schools should stay the same" is treated as a data point worth reporting in a 2025 survey — is itself a measure of where the field is.

Suggested Next Steps

Based on the patterns above, the following files in the collection deserve deeper reading as priority sources for book chapters, scholarship, or course development:

For the disclosure paradox chapter: Schilke & Reimann (The Transparency Dilemma); Noti-Victor (Regulating Hidden AI Authorship); Lemley & Ouellette (Plagiarism, Copyright, and AI); Frazier & Rozenshtein (Large Language Scholarship). These four together constitute the most recent and most provocative cluster in the collection on this topic.

For the skills gap / what to teach chapter: Schrepel (GenAI in Legal Education: A Two-Year Experiment); Bliss (Teaching Law in the Age of Generative AI); Bloomberg Law Path to Practice 2026; Sag (AI Policies for Law Schools). These provide the strongest combination of empirical evidence and practical framing.

For the agentic AI chapter (if writing one): Rodriguez (Law School Deans and Agentic AI); Bloomberg Law "AI and the Coming Mismatch"; NCBE Letter from the Chair. These three read together tell a coherent story about gatekeeping institutions caught in the shift from generative to agentic AI.

For assessment design: The AI Assignments and Rubrics docx in the Assignments folder; Cope et al. (Grading Machines: Can AI Exam-Grading Replace Law Professors?); Use Cases/Show Students Smarter Ways to Use AI. These have the most direct classroom application.

End of Report