



June 16, 2026

AI and Legal Education

Dear colleagues,

As the capabilities of AI models and platforms continue to improve, their impact on legal practice is mounting. This is challenging, but also exciting for many reasons (not least of which is AI's potential to expand access to justice).

Most observers seem to agree that these changes have significant implications for legal education. What is less clear is what law schools should do in response. Some discussions of this topic have a largely negative framing, as if there is an irreconcilable tension between the traditional pedagogical goals of legal education and the emergent need to train students to take effective advantage of AI capabilities. But these goals are not in a zero-sum relationship. It is possible to pursue them both with excellence. And that's a good thing, because the lawyers most likely to flourish in the years ahead are those who *both* are adept at making wise use of AI capabilities *and* possess the same expert judgment and perspective that have always been the hallmarks of the best lawyers.

That is the spirit and overarching objective of the education pillar in our AI strategy.¹ Please read on for more details (and note that I have some specific requests for those of you who are teaching this coming year). Constructive criticism and innovative ideas for implementation are equally welcome.

Context: Not One Challenge But Three

AI poses at least three distinct challenges in the education context:

- **The skills challenge:** What AI-related knowledge and skills should we impart to students (and how should we accomplish this)?
- **The assessments challenge:** Given ubiquitous student access to AI (whether or not we provide access to particular tools), what steps should we take to preserve the integrity of our assessment measures?

¹ Our AI education strategy is the first pillar in a three-pillar overarching framework. The other two pillars concern our research mission (*i.e.*, how we contribute research and other forms of insight relating to the law of AI) and our operations (*i.e.*, how best to use AI to enhance the capabilities and workflows of our staff and faculty).

- **The educational rigor challenge:** Some argue that the ready availability of AI poses a cognitive de-skilling risk for law students, particularly insofar as some students might engage less rigorously with course materials and assignments in favor of using AI as a shortcut to a superficial level of understanding. What might be done to mitigate that risk?

A sound law school AI strategy needs to address each of these challenges thoughtfully (and must be continually reassessed given the rapid pace of change). What follows below is a snapshot of our current approach.

1. The skills challenge

The role of AI tools in legal practice has grown tremendously over the past two years. More likely than not, that role will continue to expand (though headwinds associated with increasing passthrough of the full cost of token consumption may slow things down). To be sure, there is much debate regarding the relative utility of AI for some legal use cases. But even so, years of conversations with our alumni and employers have made it clear to me that we owe it to our students to prepare them well for the AI-impacted future in which they will practice.

To discharge this duty, we aim to invest our students with a nuanced and extensible command of AI in general and AI as applied in legal practice.

Towards that end we have taken several steps. First, we have identified an array of core learning objectives under each of those headings, encompassing both foundational understandings and practical capabilities. The first level has to do with AI capabilities in general. Learning objectives under this heading include:

- Understanding the basics of AI models (what they do, how they are created, who creates them, *etc.*).
- Understanding the options for accessing various models and tools, and developing judgment about how to select an appropriate path for a particular task.
- Understanding the significance of context layers and how to make them (including portability).
- Understanding the importance of connecting AI to other tools and to data sources, and knowing how to do so (including limitations and obstacles).
- Understanding the utility of agents and how to create them.
- Understanding the signal importance of output verification.
- Understanding the cost structure of AI-related services.

The second level has to do with AI capabilities as deployed in the context of legal practice in particular. Objectives under this heading include:

- Understanding leading use cases and having a demonstrable capacity to produce verified work product in one or more such contexts.
- Understanding specific risks and limitations arising with AI use in legal practice.
- Understanding the comparative capabilities and economics of particular platforms.
- Having the ability to make (and to articulate) justifiable selections among options for using AI in specific scenarios, accounting for capabilities, risks, limitations, and costs.
- Understanding how use of AI tools may implicate the rules of professional ethics, attorney-client privilege and work-product protections, court-specific rules, and data privacy/cybersecurity concerns.

These learning objectives will evolve over time (as will the technology itself, its uses in practice, and the economics of it all).

Next, we seek to ensure our students are on a relatively level playing field with respect to their access to the relevant services.² We are in a strong position in that respect. We were in the first wave of law schools whose students received access to Harvey and Legora (leading AI application-layer platforms focused on legal practice), and there are more such application-layer licenses on the way (in addition to the offerings from incumbents Westlaw and Lexis). Equally important, we also expect to have all our students (as well as all faculty and staff) on enterprise-grade licenses with leading frontier labs this fall.

Finally, we have developed a program to deliver on those learning objectives. Our model maps the various learning objectives onto an array of both for-credit courses (including, for example, 1L Legal Research and Writing and upper level Professional Responsibility courses) and a bespoke set of practical training sessions that are not for-credit courses but rather a key component of our existing Professional Pathways Program (in which all students participate). In this way, we achieve subject-matter adaptability, schedule flexibility, and the ability to keep students engaged with AI skills training across all three years of the law school experience.

2. The assessments challenge

Assessments (particularly graded ones) are an important part of the law school experience, and it matters a great deal whether we (and employers) have confidence in the integrity of the measures we use. Unfortunately, ubiquitous student access to AI capabilities has serious implications for the integrity of our assessments in certain contexts.

² As things stand, students (like anyone else) can access the free tier of various AI models. Some but not all students can and will spend to have access to a premium tier. From that perspective, a school that provides students with access to a high-quality tier of one or more such models helps to level the playing field.

A caveat: if a professor actually *wants* students to employ AI in some fashion in relation to an assignment, paper, or exam, this concern largely drops out. Indeed, there are fascinating opportunities to make creative use of AI as part of an assessment, particularly with seminar papers and the like. Some professors, for example, are building into their assignments requirements for their students to document thoughtful AI use in connection with various phases of their research and writing.

But let's assume that it is not the professor's intention to allow AI use in connection with some particular assessment. What then?

Unfortunately, some traditionally common assessment methods are conspicuously vulnerable to rule-breaking in this sense. This is particularly true for the take-home exam format, though really any assessment that allows students to work in an unsupervised setting has this same vulnerability. A similar example arises with in-class assessments, if the rules happen to allow students access to laptop/tablet/phone functionalities outside of a sandbox imposed by exam software. Such settings also are vulnerable, notably, to the harmful perception that some are breaking the no-AI-use rule (even if no one actually is).

We have made excellent progress mitigating this risk. Before the last academic year, our faculty discussed such considerations extensively. The end result was the near-complete elimination of take-home exams, and a sharp increase in in-class exams requiring use of software limiting what can be done on the device (thus precluding resort to AI). By the same token, we've also seen a surge of interest in assessments involving class participation and live presentations (even oral exams in some cases).

Alas, a significant assessment-integrity challenge remains: our writing seminars (and any other course that calls for students to produce an academic paper through work performed outside of class hours).

One of the core pedagogical goals of the writing seminar model is to have a student engage in a long term project of brainstorming, researching, drafting, and editing, and to do so on a solo basis in most cases (with only periodic input and steering from the professor). To be sure, it has long been possible for a student to subvert expectations by seeking help from a classmate. But with ubiquitous access to AI, the temptation to get external help is significantly greater.

This scenario warrants serious deliberation on the part of any faculty member teaching a writing seminar (or who otherwise plans to have a graded paper as part of a course). We likely will hold a Chalkboard session focused at least in part on this topic, and will be eager to ensure creative solutions are shared broadly. Notably, and in that spirit, some colleagues already have responded to this particular challenge by flipping it on its head. That is, they have amended their rules to include affirmative requirements for students to document their use of AI at specific steps in the writing process, and are evaluating them at least in part on the quality of that AI use. This has the advantage of going with

the grain of the new technology and aligning evaluation with important emergent skills. For some colleagues, this may prove a good enough solution. For others who prefer to keep AI at bay in this setting, however, more creativity will be needed.

Bearing all that in mind, I have the following request:

- *If this coming year you plan to have a graded writing assignment where the work will be done at least in part outside the classroom (whether a seminar paper or otherwise), I would appreciate it you would spend time reflecting on the assessment integrity challenge this presents. And if you develop an approach or solution that might benefit your colleagues, I ask that you reach out to share it with me or Associate Dean Morse.*

Now let's turn attention to the third of the AI education challenges.

3. The educational rigor challenge

The lawyers most likely to flourish in the age of AI will be those who not only have the training to make wise use of AI capabilities in support of their practice, but who also possess the core capabilities and skills that have long distinguished excellent lawyers. Examples of those capabilities include: rigorous analytical capacity; deep perspective regarding the complex conceptual structures of various bodies of law; discernment; the ability to persuasively communicate information and arguments in contested and high-stakes settings (including those that unfold live and in real time); and judgment of the sort needed for decision-making and client-communications amidst complex (and often highly indeterminate) legal and factual problems.

And so we should ask: might ubiquitous student access to AI pose a threat to the ability of law schools to cultivate those essential capabilities?

This brings us to a topic sometimes called “cognitive de-skilling.” Cognitive de-skilling refers to the notion, as old as Plato’s *Phaedrus*, that something important in our cognitive capabilities might be lost in the exchange when a new information technology emerges to take over something humans previously had to manage themselves. Such a loss doesn’t occur with every such exchange (society should not begrudge the lost utility of the book-based Shepardizing skills I acquired in the 1990s, for example). But the possibility of a problematic tradeoff warrants analysis, especially when the emergent technology at issue directly impacts education itself.

With this in mind, consider the variety of distinct educational processes that exist in the law school setting. For the sake of simplicity, we will set aside all non-curricular aspects of legal education, and we will further narrow the focus to the context of a single course. Across the months of that one course, there are many distinct educational processes at work:

- Reading and taking notes on the materials assigned for a class session
- Participating in classroom dialogue (especially if called on in class)
- Capturing insight from other aspects of a class session
- Attending office hours to ask questions (and listen to others doing the same)
- Participating in a study group
- Outlining a course
- Conducting research and writing
- Taking practice exams
- Other forms of studying for exams
- Taking exams
- Reviewing exam results with the professor

The question is: will ubiquitous student access to AI pose a cognitive de-skilling risk in the context of some of these processes?

First, the good news. There are some clearly beneficial educational use cases for AI. It can be used to generate practice exams, for example, including ones informed by the specific professor's back catalogue of exam questions (if those are made available). AI also has tremendous potential to serve as an always-available, patient, and discreet tutor, especially if the model has access to the relevant course materials and is instructed to play the tutorial role in a Socratic manner rather than just giving answers. Indeed, the custom-tutoring use case for AI is likely to emerge as a powerful complement to office hours (though some aspects of office hours will remain uniquely valuable to students, such as the ability to develop a mentoring relationship with the professor, or to have certainty that the information provided tracks the professor's understanding).

That said, AI also can be used as shortcut that quickly and easily puts an answer into a student's hands, replacing deeper learning with a more-superficial semblance. That is, it brings with it a serious and novel risk of cognitive de-skilling as to any of the learning processes described above insofar as they take place in a setting where there is no realistic way to preclude a student from turning tasks over to AI.

In the preceding section, I discussed this very problem in relation to the integrity of assessments. My point now is that the same scenario also presents an educational rigor issue. And notably, this one arises not just with direct forms of assessment, but also with one of the core elements of the legal education process listed above: the act of preparing for class by reading and synthesizing complex material.

To be sure, students have long had access to tempting shortcuts when it comes to preparing for class. From treatises and nutshells, to insights from classmates, to whatever one might find online, there have long been available means not just to deepen understanding of the readings but also to skip doing the readings at all while still gaining at least a superficial degree of understanding.

So in what sense is the possibility of taking such shortcuts via AI a novel state of affairs? The answer is that AI will likely prove to be considerably more tempting than the traditional bypasses. This is so due to AI's comparative availability, confidentiality, and range, as well as its distinct capacity for iterative dialogue. These qualities will tempt more students more often to spend less time wrestling directly with the assigned materials (or even to skip such direct engagement outright).

I would love to be wrong about that. And, to be clear, I certainly appreciate that there are plenty of students who resist this temptation and will continue to do so, understanding that de-skilling is a serious and harmful risk. Still, schools have an obligation to take this challenge seriously, and to seek opportunities to mitigate it.

Schools can respond by redoubling efforts to emphasize to students the value of doing the hard work in the first instance themselves. We will do that. Exhortations will only go so far, however.

I would like to highlight a further way to respond. It's not a silver bullet, but I believe it will help to a degree.

The practical root of the problem here is the fact that it is not possible to preclude resort to AI reliably in an unsupervised setting. Well, there is a flip-side to that observation: *supervised* settings have become even more valuable than they used to be. Such settings are not just important for preserving assessment integrity, but also because they are the sole context in which a professor can be certain that a student is engaged in the hard but essential work of analysis and communication unmediated by AI.

Put another way: as the rigor of out-of-class learning processes is eroded by AI-enabled shortcutting, the minutes of class meetings become an ever-more valuable asset.

The potential of the controlled in-class environment won't be realized, however, if the professor at the podium is merely lecturing to the assembled students, nor if the occasional questions a professor does pose to students pass quickly and lightly. And so, if the goal is to counteract deskilling by ensuring that all students have recurring occasions in which they must work through and communicate about difficult concepts wholly on their own, faculty need to make extensive use of classroom time to engage students in sustained and rigorous dialogue.

This points towards a renewed emphasis on the Socratic mode of teaching. It also suggests that faculty should take steps to ensure that when a student is engaged in such dialogue the student is not distracted by (let alone relying upon) whatever might be taking place on their screen at that time.

Of course, many if not most faculty already use a Socratic mode of teaching at least some of the time. But not everyone does, and in any event there always is room for

improvement. That certainly would be the case for me if I were still teaching, I will admit. And so I conclude with this request for all who are teaching this coming year:

- *Please give sustained thought to how your usual teaching method relates to what I've written above, and ask yourself how you might change things to take best advantage of the precious opportunity that the classroom setting provides.*

In the emerging age of AI, our students will flourish best if they possess the fruits of *both* traditional legal training *and* a thoughtful program of AI upskilling that will help them make the best possible use of AI once in practice. We are well positioned to succeed on both fronts, and I look forward to working with each of you on this, both now and for many years to come.

- Bobby Chesney, Dean