AI Exposure, Experience, and Attitudes at Columbia Law School

2025–2026 Annual Report

Prepared by Joanna Zhang, Mackenzie Berwick, and Jacob Lee

Columbia Law School

Table of Contents

Project Overview

In Fall 2025, the “Leading in AI” Initiative worked with the Columbia Law School Student Senate to design and administer the 2025–2026 AI Exposure, Experience, and Perspectives Survey. This survey is intended to gain insights into the community's sentiments and use of artificial intelligence (“AI”) tools.

The survey captures a wide range of student experiences with AI tools, spanning from daily users to those who never engage with AI. Across this spectrum, respondents express a mix of enthusiasm, caution, and concern. The data reveals strong interest in structured institutional guidance, broad agreement on the need for responsible integration, and persistent anxieties about accuracy, ethics, and the future of legal work.

Executive Summary

The 2025–2026 AI Exposure, Experience, and Perspectives Survey finds that members of the Columbia Law School community engage with AI to varying degrees, with students reporting the highest familiarity. Only 12% of JD students and 2% of LLM students indicated they were “not familiar” with AI tools relevant to the legal field. 22% of JD students and 39% of LLM students indicated they were “very familiar” with AI tools. In contrast, professors and staff reported lower familiarity, with only 19% of professors and 8% of staff identifying as “very familiar.”

I. AI Familiarity

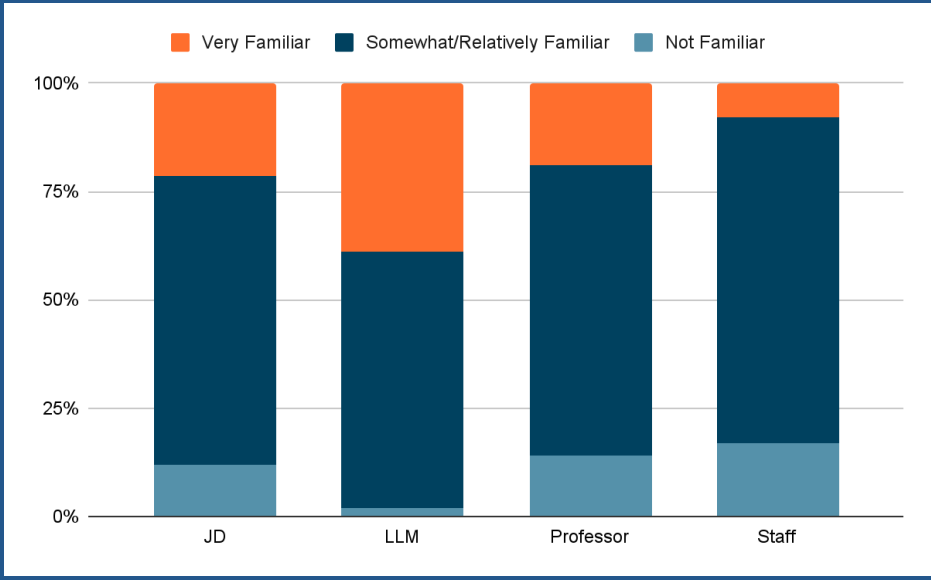

Patterns of use reflect these differences. 78% of all respondents selected ChatGPT as one of their most frequently used tools, making it a dominant platform across groups. Frequency of use also varied by group. 28% of all respondents reported daily use of AI tools. Weekly use was the most common overall, selected by 38% of respondents. LLM students reported a strong utility on AI tools, with 50% of respondents using AI tools daily.

JD and LLM students reported the broadest use of AI tools, with the most popular uses being personal tasks and academic coursework or research. Professors and staff reported more limited engagement, with 24% of professors and 17% of staff indicating they do not use AI tools in any context.

II. Respondents' Sentiment Towards AI

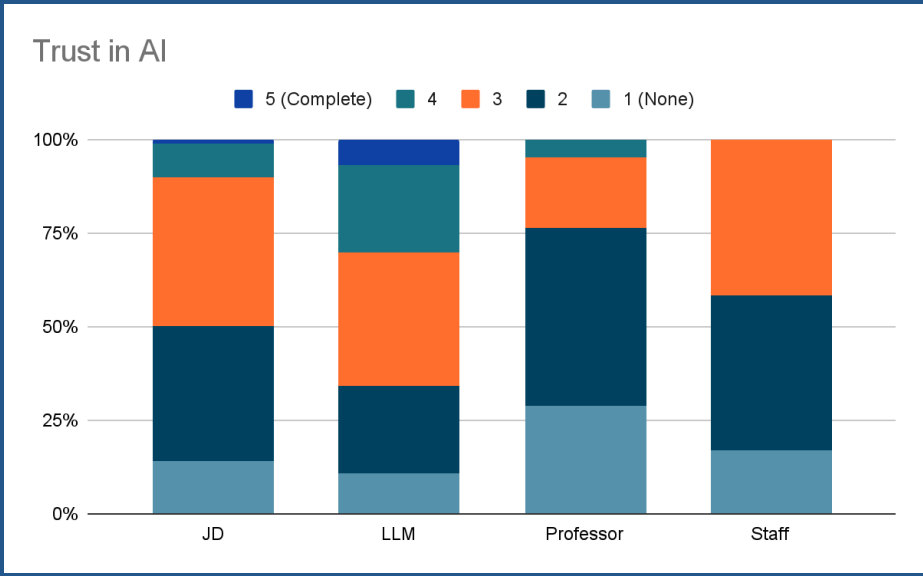

What are your biggest concern(s) with AI generally?					
	Accuracy of the AI models	Lack of critical thinking due to reliance	Transparency about whether a work product was created by a person or by AI	Privacy concerns about how information entered into AI tools is stored, used, or retained	Environmental impact due to high energy consumption by data centers
Total	84%	69%	55%	53%	45%

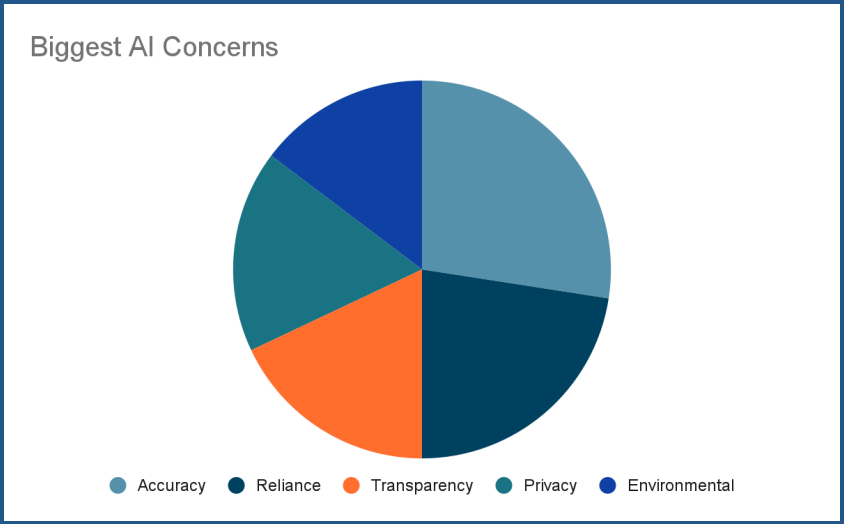

Across respondent groups, the most common sentiment toward AI was "curious but cautious." Trust in AI outputs generally fell in the middle of the response scale, with relatively few respondents selecting the highest trust categories. In the open-text responses, respondents expressed concerns related to accuracy, bias, and academic integrity; however, these themes were not quantified in the survey and therefore cannot be compared proportionally to the closed-ended items.

III. Interventions and Events

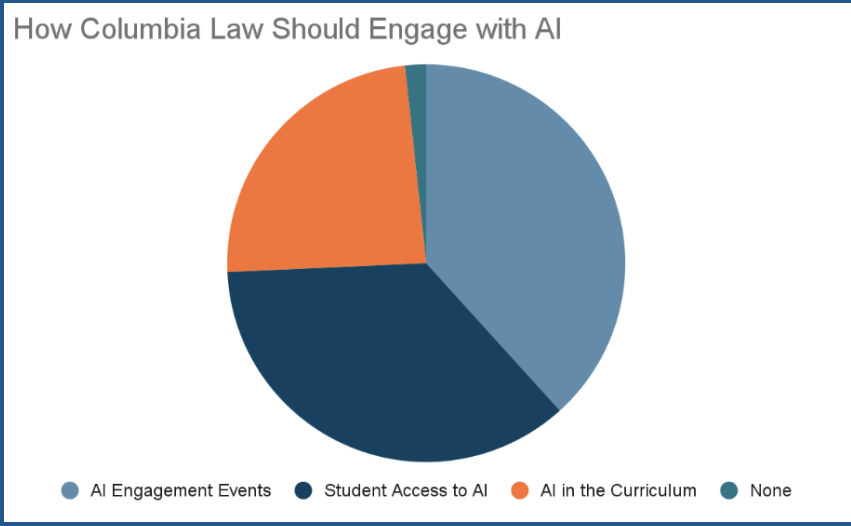

Across the community, respondents expressed interest in Columbia Law School providing additional guidance, training, and opportunities to learn about AI. While the survey results do not prescribe specific interventions, they indicate that students, faculty, and staff are seeking clearer expectations and more structured support as AI tools become increasingly present in academic and professional settings.

Literature Review

Generally, a median of "34% of adults say they are more concerned than excited about the increased use of AI."[1] However, looking globally, no country reports that more than "three-in-ten adults say they are mainly excited."[2] Within broader global, professional, and educational trends in AI adoption and critique, legal education occupies a distinct role. The duty to train the next generation of lawyers to address key questions about the integration of AI into workplaces, the judicial system, and the practice of law remains prominent.

I. Governance

Aware that "the world looks to Columbia Law School as a leader in legal education and scholarship," Columbia Law School seeks to prepare its graduates to take up that mantle.[3] Teaching students the skills of judgement, discernment, and direction is key to fulfilling that mission in the age of AI integration.

At Columbia Law School, the governance landscape involves multiple stakeholders:

- Office of the Dean – AI Task Force
- Davis Polk Leadership Initiative – "Leading in AI" Initiative
- Student Affairs Committee – Student Government
- Office of Student Services
- Office of Private Sector Careers
- Public Interest and Public Service Careers

Overlapping with this mission, the Davis Polk Leadership Initiative strives to equip students "to cultivate the lawyer leadership skills they will need to succeed and thrive at every stage of their professional development and in a wide variety of professional

[1] Poushter, J., Fagan, M., & Corichi, M. (2025). How people around the world view AI (p. 5). Pew Research Center. https://www.pewresearch.org.
[2] *Id.*
[3] Columbia Law School. (n.d.). Why Columbia. https://www.law.columbia.edu/admissions/why-columbia.

and civic settings."[4] In this nexus, the "Leading in AI Initiative" seeks to empower Columbia Law students with fluency in AI, ethical grounding, and leadership skills to navigate AI-driven disruption and develop models of good governance. The project advances urgent conversations about AI in the legal profession while promoting adaptability, collaboration, and resilience across the Columbia community and reinforcing the Law School's position as a global changemaker.

Columbia Law School's "Leading in AI" Initiative frames this moment as one requiring balance between innovation and caution, efficiency and ethics, individual responsibility and institutional policy.

II. Legal Education

In a recent report, 82% of higher education instructors reported academic integrity as a top concern for AI's impact.[5] In fact, nearly "three in four [instructors] reported needing more training to work with new technologies."[6]

Many AI detectors remain unreliable and cannot be used as the sole basis for academic misconduct decisions.[7] Researchers have concluded that detectors disproportionately misidentify samples from non-native English speakers and students with atypical writing patterns as falsified.[8]

---

[4] Davis Polk Leadership Initiative. (n.d.). About the Leadership Initiative. https://leadership-initiative.law.columbia.edu/content/about-leadership-initiative.
[5] Cengage Group. (2025, June 27). AI's impact on education in 2025. https://www.cengagegroup.com/news/perspectives/2025/ais-impact-on-education-in-2025/.
[6] *Id.*
[7] Scribbr – AI Detector Accuracy Study (2024–2025) (https://www.scribbr.com/ai-detector/best-ai-detector/) conducted a large comparative benchmark of 12 AI detectors, finding even the best detectors misclassify human writing and AI-generated text. Even premium Scribbr detectors reached 84% accuracy, meaning they still held an 16% error rate even under ideal conditions. Market competitors Copyleads AI Detector and QuillBot both demonstrated under independent evaluations that both programs still produce false positives. (https://copyleaks.com/ai-content-detector) (https://quillbot.com/ai-detector)
[8] https://kaltmanlaw.com/ai-detectors-academia/

The American Bar Association's 2024-2025 "Teaching Law in the Age of Generative AI" Report explains that AI complicates assessment design, verification of authorship, and academic integrity enforcement.[9] According to Jenny Silbiger, the president of the American Association of Law Libraries, law librarians are stepping into new roles as they promote new AI-based tools for government, law schools and law firms.[10] While other law schools have undertaken varied policies towards generative AI,[11] Columbia Law School, led by Dean and Lucy G. Moses Professor of Law Daniel Abebe, has determined to take an informed approach to preparing its students to enter a world saturated by artificial intelligence. In his own words, Dean Abebe professed that "It's our responsibility as educators here to make sure that we're empowering [students] with the tools to be successful in a world that will be increasingly driven by AI."[12]

III. Ethics and Impact

AI development generates significant environmental externalities. Training a single large model can emit more than 626,000 pounds of CO_2-equivalent,[13] and environmental scholars note that these burdens fall disproportionately on low-income and racialized communities, deepening existing environmental-justice inequities.[14]

---

[9] https://www.americanbar.org/groups/legal_education/publications/

[10] https://www.americanbar.org/groups/journal/articles/2026/law-librarians-at-the-forefront-of-using-and-educating-on-ai/

[11] See University of Chicago Law School's 2023-2025 Generative AI Policy defining generative AI, setting course-by-course discretion for AI use, and warning faculty not to rely on AI detectors for misconduct findings (https://www.law.uchicago.edu/students/handbook/generative-ai-policy). See also University of Kansas School of Law's AI for Lawyer's Course (2025) integrating AI into legal pedagogy focusing on addressing hallucinations, verification, and ethical usage (https://www.cambridge.org/core/journals/international-journal-of-legal-information/article/integrating-generative-ai-in-legal-pedagogy). See also 2024 Survey of 29 Law School Deans for the American Bar Association (https://www.americanbar.org/news/2024/ai-law-school-deans-survey/).

[12] https://www.law.columbia.edu/news/archive/faculty-experts-generative-ai-legal-education-and-future-profession

[13] Djanegara, N. D. T., Zhang, D., Badi Uz Zaman, H., Meinhardt, C., Watkins, G., Nwankwo, E., Wald, R., Kosoglu, R., Koyejo, S., & Elam, M. (2024). Exploring the impact of AI on Black Americans: Considerations for the Congressional Black Caucus's policy initiatives (p. 12). Stanford Institute for Human-Centered Artificial Intelligence.

[14] *Id*. at 12–13.

An example is Microsoft's recent 20-year agreement to purchase power from the planned restart of Three Mile Island to meet AI-driven energy needs.[15] Nuclear development has long imposed disproportionate burdens on marginalized communities, from uranium extraction on Indigenous lands to waste storage near low-income towns.[16] The Three Mile Island project risks repeating this pattern. The surrounding community has already been shaped by a legacy of nuclear trauma and they are being asked to absorb new environmental and safety risks.

Beyond the environmental risks presented, there is no question that "[technology] is never neutral. It reflects and reinforces the values of those who develop it."[17]

IV. Usage

According to Cengage's Teen and Young Adult Perspectives on Generative AI Report, 53% of surveyed students use generative AI to gather information and 51% use it for brainstorming.[18] Further, the increased use of AI in educational settings is "correlated with heightened risks to students."[19]

In the legal space, 79% of legal professionals report using AI in some capacity, reflecting a shift toward automation and augmented legal work.[20] Adoption spans a growing spectrum of tools. While Westlaw Co-Counsel and LexisNexis AI+ embed generative AI directly into major research platforms, platforms like Harvey offer large-language-model-driven analysis and contract-review systems like Kira and Luminance use machine learning to accelerate due diligence and document review.[21]

---

[15] CNBC. (2024). Constellation to restart Three Mile Island Unit 1 in 2028 under 20-year deal with Microsoft.
[16] Djanegara et al., 2024 at 12–13.
[17] *Id*. at 6.
[18] Cengage Group. (2025, June 27). AI's impact on education in 2025. https://www.cengagegroup.com/news/perspectives/2025/ais-impact-on-education-in-2025/.
[19] Center for Democracy & Technology. (2025). Hand in hand: Schools' embrace of AI. https://cdt.org/insights/hand-in-hand-schools-embrace-of-ai/
[20] Clio. (2025). AI tools and software for lawyers: Improving productivity in law firms (p. 1).
[21] *Id.*

V. Accuracy

Empirical studies show that legal AI tools hallucinate at significant rates. In fact, legal models hallucinate in 1 out of 6 or more benchmarking queries.[22] As hallucination rates climb ever higher, the promised efficiency gains are undermined. In the legal profession, this can create additional workloads as inefficient AI integration can create extra work arising from the need for verification, creating "double work."[23]

The infusion of AI-powered work into industry has invented a new term: "workslop." This term refers to AI-generated work content that "masquerades as good work, but lacks the substance to meaningfully advance a given task."[24] Employers report that their employees often use AI tools to create "low-effort, passable-looking work that ends up creating more work for their coworkers."[25] Across industries, the temporal and financial cost of this work is tangible. Juggling "workslop," employees reported spending "an average of one hour and 56 minutes dealing with each instance."[26]

VI. The Current Landscape

Students polled through Cengage's report on AI in education reported that they want to learn how to use AI "responsibly and effectively, but they need support from faculty and institutions to make that happen."[27] Still, across the literature, there is consensus that "AI will never replace the judgment and experience of lawyers."[28]

How can a law school cultivate judgment and discernment in an educational environment already saturated with the very challenges it seeks to confront?

---

[22] Stanford Institute for Human-Centered Artificial Intelligence. (2025). AI on trial: Legal models hallucinate in 1 out of 6 (or more) benchmarking queries. Stanford HAI.
[23] *Id.*
[24] Niederhoffer, K., Kellerman, G. R., Lee, A., Liebscher, A., Rapuano, K., & Hancock, J. T. (2025, September 22). AI-generated "workslop" is destroying productivity. Harvard Business Review, p. 1.
[25] *Id.* at 1.
[26] *Id.* at 4.
[27] Cengage Group. (2025). AI's impact on education in 2025. https://www.cengagegroup.com.
[28] Clio. (2025). AI tools and software for lawyers: Improving productivity in law firms. https://www.clio.com.

Findings

Following an overview of the survey's methodology and structure, this section presents a look at how members of the Columbia Law School community are responding to AI. The survey was designed to capture a wide range of experiences across students, faculty, staff, and practitioners, and its structure allows us to examine not only how respondents use AI tools, but also how they feel about AI's role in legal education and legal practice.

The findings are organized into three core areas.

First, we analyze engagement with AI tools, focusing on familiarity, usage frequency, and the contexts in which AI is being relied upon in academic, professional, and personal work. Second, we explore sentiment toward AI, capturing the excitement, caution, skepticism, and concern that shape how respondents view AI's pros and cons. Finally, we evaluate engagement with AI at Columbia Law School, highlighting the community's perspective on how the institution ought to respond.

I. Methodology

All data processing and analysis were conducted in R. As with any research enterprise, this work required us to confront a series of implicit and embedded methodological questions. By making these decisions transparent, we aim to offer clarity about the foundations of our analysis.

A. Survey Development

The survey was developed by the "Leading in AI" team in collaboration with the Columbia Law School Student Senate, the Davis Polk Leadership Initiative, Faculty Advisor and Alphonse Fletcher Jr. Professor of Law and Dean Emerita, Gillian Lester, and Associate Dean of Registration and Student Services, Andrea Saavedra. After multiple rounds of drafting and refinement, the final instrument included questions on

AI familiarity, tool usage, trust, concerns, and expectations for institutional engagement.

B. Respondent Categorization

While we recognize that these responses do not capture the full breadth of perspectives at Columbia Law School, we are confident that the data offers sufficient insight to guide interventions. A total of 277 individuals participated in the survey:

	Total Responses
Law Student (JD)	200
Law Student (LLM)	44
Professor	21
Staff	12
Total	277

The “Law Student (JD)” category includes students enrolled in the JD program. The “Law Student (LLM)” category includes students enrolled in the LLM program. The “Professor” category includes faculty, adjunct professors, and Legal Practice Workshop professors. The “Staff” category includes administrators and staff.

C. Data Cleaning

Entries that could not be reliably coded were excluded to maintain consistency and reduce the risk of misclassification.

One respondent entered “SPL” as their identification category; because the meaning could not be determined, that response was removed in accordance with the study’s exclusion criteria and is not reflected in any analysis or exhibits.

For the question asking respondents to identify the tools they most frequently used, categories indicated by fewer than ten respondents were removed from analysis. This threshold affected several AI tools (e.g., Luminance, Grok, NotebookLM, Perplexity, DeepSeek).

It should also be noted that a small number of respondents also used the open-text field to indicate that they do not intentionally seek out AI tools and instead encounter them only when embedded within other platforms. Because these responses also fell below the ten-response threshold, they were excluded under the same rule.

D. Limitations

Several limitations should be noted:

- The survey relied on self-reported data, which may introduce recall or perception bias.
- Participation was voluntary, and the sample may not fully represent all members of the Columbia Law School community.
- Group sizes were uneven, particularly among faculty and staff.
- Exclusion of categories with fewer than ten responses, while methodologically appropriate, may omit niche but meaningful perspectives.
- Open-text responses were coded by the research team; although handled conservatively, coding decisions may introduce subjectivity.

II. Survey Structure

The survey consisted of the following structure:

Part 1: AI and You

This section is intended to collect information about how you are currently using artificial intelligence tools. Most AI tools in general use today are generative AI systems, particularly Large Language Models (LLMs). An LLM is an AI system that reads and learns from vast amounts of text so it can generate responses that sound natural and intelligent.

1. Which group best describes you? (LPW instructors should also select their primary employment category)
2. [For students only] To help us capture diverse student communities and experiences, please identify which groups are you involved with on campus. (Ex. APALSA, FGP).
3. How familiar are you with the use of AI tools relevant to the legal field (e.g., legal research, document automation, predictive analytics)?
4. Which AI tools do you use most frequently? (Select up to 3 from the list below)
5. How frequently do you use AI tools?
6. In what contexts do you currently use AI tools? (Select all that apply)
7. How do you currently feel about the integration of AI into legal work and processes?
8. To what extent do you trust AI outputs (e.g., legal summaries, citations, analyses)?

Part II: AI and Us

This section is intended to gather your perspectives on how artificial intelligence should be used within the Columbia Law community and the broader legal profession.

9. In which of the following ways do you believe Columbia Law should engage in AI? (Select all that apply)
10. What is your biggest frustration with AI?
11. How do you think Columbia Law School should respond to an AI era?
12. Imagine that one of your courses adopted an open-AI tool policy for your first written assignment. Which of the following would be your greatest concern?

13. What are your biggest concern(s) with AI generally? Check all that apply.
14. Do you think the law school has provided sufficient guidance on the ethical and academic use of AI?
15. What kinds of future workshops or initiatives on AI and the law would most interest you?
16. Please feel free to share any additional comments, concerns, or thoughts below.

III. Analysis

All respondents express frustration rooted in uncertainty, particularly around AI's accuracy, norms, institutional expectations, and its broader societal direction. Analysis reveals a strong interest in workshops, faculty-supported instruction, and clear, consistent policies to reduce uncertainty and help the community navigate AI's impact.

A. Engagement with AI Tools

Across the Columbia Law School community, engagement with AI tools is widespread but unevenly distributed. Students use AI frequently and across a broad range of academic, professional, and personal contexts. Weekly and daily use is routine for a significant portion of students, and general-purpose tools like ChatGPT dominate usage patterns, outpacing specialized legal AI platforms. Faculty and staff, by contrast, report lower familiarity and more sporadic use.

Exhibit A

How familiar are you with the use of AI tools relevant to the legal field (e.g., legal research, document automation, predictive analytics)?				
	Not Familiar	**Somewhat Familiar**	**Relatively Familiar**	**Very Familiar**
Law Student (JD)	12%	28%	40%	22%

Law Student (LLM)	2%	14%	45%	39%
Professor	14%	38%	29%	19%
Staff	17%	50%	25%	8%
Total	**10%**	**27%**	**39%**	**23%**

Of 200 JD respondents, 89% reported at least some familiarity with AI tools relevant to the legal field, meaning they indicated a level of familiarity above "not familiar." LLM students, making up a smaller proportion of the total number of responses, showed a similar trend to JD students. Of the 44 LLM respondents, 98% indicated at least some level of familiarity, with 39% indicating they were"very familiar" with AI tools. Their responses ought to be contextualized by the unique composition of LLM students at Columbia Law School. Members of the LLM Class of 2026 hold citizenship from 52 countries and have an average of four-and-a-half years of professional experience.[29]

In contrast, professors and staff reported lower levels of familiarity. 38% of the professors surveyed felt "somewhat familiar" with AI tools, and 19%considered themselves "very familiar." Staff responses were similar, with just one person reporting high familiarity.

Overall, the majority of respondents (39%) placed themselves in the "relatively familiar" category. Only 10% of all participants said they were "not familiar" with AI tools.

[29] https://www.law.columbia.edu/news/archive/orientation-2025-new-cohort-joins-columbia-law. Their perspectives reflect the legal norms, regulatory environments, and professional cultures of their home jurisdictions, as well as the habits and expectations formed through prior legal practice. As a result, their responses capture not only individual attitudes toward AI, but also the broader institutional and geopolitical contexts in which they were trained and have worked.

Exhibit B

Which AI tools do you use most frequently? (Select up to 3 from the list below)									
	None	ChatGPT	Lexis	Westlaw	Claude	Copilot	Gemini	Harvey	Relativity
Law Student (JD)	11%	77%	40%	29%	13%	7%	6%	3%	3%
Law Student (LLM)	2%	89%	48%	25%	11%	20%	11%	9%	7%
Professor	19%	67%	10%	19%	38%	5%	10%	5%	5%
Staff	8%	75%	0%	8%	17%	8%	33%	0%	8%
Total	**10%**	**78%**	**37%**	**27%**	**14%**	**9%**	**8%**	**4%**	**4%**

ChatGPT stands out as the most frequently used AI tool across all groups. JD students selected ChatGPT more often than any other option. LLM students show a similar pattern, reflecting engagement with general-purpose AI assistants.

Traditional legal research tools like Lexis and Westlaw remain widely used, especially among JD students, while more specialized tools like Claude, Copilot,[30] Gemini, Harvey, and Relativity see less frequent adoption. LLM students report slightly broader use of emerging tools.

Professors and staff, by contrast, report lower usage across all tools. A number of respondents, mostly JD students, selected "None," though they remain a minority.

[30] It should be noted that the survey circulation coincided with the beginning of Columbia University's integration of and partnership with Microsoft's Copilot.

Exhibit C

How frequently do you use AI tools?					
	Daily	**Weekly**	**Monthly**	**Rarely**	**Never**
Law Student (JD)	24%	41%	13%	13%	10%
Law Student (LLM)	50%	41%	0%	7%	2%
Professor	24%	19%	14%	24%	19%
Staff	17%	25%	0%	58%	0%
Total	**28%**	**38%**	**10%**	**15%**	**9%**

Across all groups, weekly use of AI tools was the most common pattern, with 38% of respondents selecting this option. Daily use was also substantial, driven primarily by JD and LLM students.

Less frequent use (monthly or rarely) was more common among professors and staff. Professors were fairly evenly distributed across daily, weekly, monthly, and rare use. Staff responses leaned more toward rare use, with 58% selecting that option.

Only 9% of respondents reported never using AI tools, with JD students making up the majority of this group. LLM students and staff had few respondents in the “never” category.

Exhibit D

In what contexts do you currently use AI tools? (Select all that apply)						
	Personal Use	Academic Coursework or Research	Outlining or Study Preparation	Job Application or Professional Writing	Legal Writing or Citation Assistance	None
Law Student (JD)	65%	47%	45%	39%	26%	12%
Law Student (LLM)	73%	61%	59%	50%	39%	5%
Professor	57%	52%	5%	0%	29%	24%
Staff	58%	17%	0%	17%	0%	17%
Total	**65%**	**48%**	**42%**	**36%**	**27%**	**12%**

JD students reported engagement with AI across nearly every context. The most common use was personal, with 65% of students reporting that they use AI tools for personal purposes. Subsequent usages in order of popularity were: academic coursework or research, outlining or study preparation, job applications or professional writing, and legal writing or citation assistance. Only 12% of JD students reported no current use.

LLM students show similar patterns. Personal use was again the leading application followed by academic and study-related contexts. Only 5% of LLM students reported no use.

Professors and staff report more limited engagement. Among professors and staff, the most common use was for personal use followed by academic coursework or research.

B. Sentiment Towards AI

The most common response ("curious but cautious") reflects a sentiment which blends openness and hesitance across the community. Students express both the highest levels of excitement and the highest levels of concern. Trust in AI outputs remains moderate at best, with most respondents clustering around the middle of the scale and very few expressing high trust. Faculty and staff show low confidence in AI-generated legal analysis.

Exhibit E

How do you currently feel about the integration of AI into legal work and processes?					
	Excited	**Curious But Cautious**	**Neutral**	**Concerned**	**Skeptical**
Law Student (JD)	18%	37%	3%	28%	15%
Law Student (LLM)	30%	57%	2%	9%	2%
Professor	19%	29%	5%	29%	19%
Staff	0%	42%	8%	17%	33%
Total	**19%**	**40%**	**3%**	**24%**	**14%**

Across all groups, the most common sentiment toward AI in legal work is "curious but cautious." JD and LLM students tended to be more excited or curious but cautious. Professors tended to fall on either end of the spectrum. Staff tended to be more cautious or skeptical.

Neutral responses were rare, representing only 3% of total respondents.

Exhibit F

To what extent do you trust AI outputs (e.g., legal summaries, citations, analyses)					
	Not At All 1	2	3	4	Completely 5
Law Student (JD)	14%	36%	40%	9%	1%
Law Student (LLM)	11%	23%	36%	23%	7%
Professor	29%	48%	19%	5%	0%
Staff	17%	42%	42%	0%	0%
Total	**15%**	**35%**	**38%**	**10%**	**2%**

For JD students, trust levels range widely, with a skew toward distrust. Only 10% of JD students rated their trust at the highest levels (4 and 5). Of LLM students, 30% rated their trust at levels 4 or 5, while 34% rated their trust at levels 1 or 2. Among professors, 76% rated their trust at levels 1 or 2. Among staff, 58% rated their trust at levels 1 or 2. No professor or staff chose level 5, indicating complete trust.

Overall, the most common rating across all roles was a moderate 3, chosen by 38% of respondents.

Exhibit G

What are your biggest concern(s) with AI generally? Check all that apply. (Pt. 1)					
	Accuracy of the AI models	Lack of critical thinking due to reliance	Transparency about whether a work product was created by a person or by AI	Privacy concerns about how information entered into AI tools is stored, used, or retained	Environmental impact due to high energy consumption by data centers
Law Student (JD)	84%	69%	55%	54%	51%
Law Student (LLM)	89%	55%	48%	41%	11%
Professor	86%	90%	71%	57%	48%
Staff	83%	92%	67%	83%	75%
Total	**84%**	**69%**	**55%**	**53%**	**45%**

Exhibit H

What are your biggest concern(s) with AI generally? Check all that apply. (Pt. 2)					
	Bias and discrimination of AI models	Job displacement and the changing landscape of entry-level positions	Equitable accessibility due to the cost of premium AI models / Access to technology	Safety concerns, especially for children using AI	Whether AI use could compromise IP rights or attorney - client privileged information
Law Student (JD)	40%	40%	38%	38%	34%
Law Student (LLM)	25%	23%	39%	11%	11%
Professor	38%	48%	33%	57%	38%
Staff	75%	58%	33%	58%	42%
Total	**39%**	**39%**	**37%**	**36%**	**31%**

Accuracy of AI models was the most frequently selected concern overall, selected by 84% of respondents. This reflects a broad hesitation to rely on AI-generated outputs without human verification. Lack of critical thinking due to reliance on AI was one of the most frequently selected concerns overall, selected by 69% of respondents. Overall, more professors and staff were concerned about critical thinking abilities than JD and LLM students. The level of concern was elaborated in open-box comments that expressed fear that over-dependence on AI tools could erode essential analytical skills, particularly for students still developing foundational legal reasoning abilities.

C. Engagement with AI at Columbia Law School

When asked how Columbia Law School should engage with AI, respondents overwhelmingly call for more structured institutional involvement. Students want clearer ethical guidance, more learning opportunities, and integration of AI literacy into the curriculum.

Exhibit I

In which of the following ways do you believe Columbia Law should engage in AI? (Select all that apply)				
	Host events to engage faculty, practitioners, and students in conversation	**Give students access to AI platforms, in the same way we have access to Westlaw and Lexis**	**Integrate AI into the curriculum of classes**	**None**
Law Student (JD)	65%	64%	39%	4%
Law Student (LLM)	68%	70%	59%	0%
Professor	76%	43%	43%	0%
Staff	92%	58%	25%	0%
Total	**67%**	**63%**	**42%**	**3%**

Across all groups, the strongest support was for hosting events that bring together faculty, practitioners, and students, with 67% of respondents indicating this preference. Providing student access to AI platforms, similar to existing access to Westlaw and Lexis, was nearly as popular, with 63% of respondents indicating this preference. Both

JD and LLM students expressed strong support for this option. Faculty and staff were less emphatic but still supportive.

Integrating AI into the curriculum received somewhat lower but still substantial support, with 42% indicating this preference.

Very few respondents chose “None.”

Exhibit J

How do you think Columbia Law School should respond to an AI era?				
	Should ban AI in their own teaching and for students	**Should use own discretion in using AI for teaching, but not allow for students**	**Should teach relying on AI but not allow for students**	**Should teach relying on AI and also allow AI use for students**
Law Student (JD)	23%	26%	7%	45%
Law Student (LLM)	11%	16%	2%	70%
Professor	5%	33%	5%	57%
Staff	25%	17%	8%	50%
Total	**20%**	**24%**	**6%**	**50%**

Across all groups, the strongest support was for a fully integrated approach, with 50% of respondents indicating that Columbia Law should teach using AI and allow students to use it.

The least popular option overall was the idea that the school should teach relying on AI but not allow students to use it, with only 6% of respondents indicating this preference.

Professors also preferred to allow for discretion in using AI for teaching but banning AI use for students, with 33% of professors indicating this option.

Only 20% of respondents endorsed a total ban on AI access for teaching and students.

Exhibit K

Imagine that one of your courses adopted an open-AI tool policy for your first written assignment. Which of the following would be your greatest concern?				
	Unfair competition between students with more exposure / experience	**Loss of opportunity to cultivate critical skills required for lawyering**	**Lack of governance or regulation policing AI use**	**No big concern – Issues can be mitigated as they arise**
Law Student (JD)	14%	61%	12%	14%
Law Student (LLM)	18%	52%	14%	16%
Professor	5%	86%	0%	10%
Staff	0%	92%	8%	0%
Total	**13%**	**63%**	**11%**	**13%**

Across all categories, the most cited concern was the loss of opportunity to cultivate critical skills required for lawyering, with 63% of respondents selecting this option. This ratio amounts to more than four times the number who cited any other concern.

Exhibit L

Do you think the law school has provided sufficient guidance on the ethical and academic use of AI?			
	Yes	**No**	**Maybe**
Law Student (JD)	15%	62%	24%
Law Student (LLM)	32%	32%	36%
Professor	19%	62%	19%
Staff	0%	67%	33%
Total	**17%**	**57%**	**26%**

Across all groups, the clear majority of respondents felt that the law school has not provided sufficient guidance on the ethical and academic use of AI. This view was especially strong among JD students, professors, and staff.

26% of respondents selected “Maybe.”

17% of respondents believed the school had provided adequate guidance, making “Yes” the least selected option. Support for this view was limited across groups, especially among staff, none of whom felt the guidance was sufficient.

Exhibit M

What kinds of future workshops or initiatives on AI and the law would most interest you?			
	A panel of practitioners discussing how AI is affecting associate and attorney experience at big law firms, public interest employers, and in-house roles — Focusing on restructuring and workforce impacts	**A session with professors and faculty about how to better shape an education response to AI — Focusing on building skills**	**Building a discussion platform for students to share their perspectives across 1, 2, and 3L classes alongside LLMs — Focusing on the student perspective and experience**
Law Student (JD)	72%	51%	34%
Law Student (LLM)	70%	61%	41%
Professor	57%	57%	38%
Staff	67%	83%	17%
Total	**70%**	**55%**	**35%**

The most popular option was “a panel of practitioners discussing how AI is reshaping the legal workforce—particularly in big law firms, public interest organizations, and in-house roles.” This session, focused on restructuring and workforce impacts, drew interest from 70% of respondents.

The second most favored initiative was “a session with professors and faculty aimed at shaping educational responses to AI.” This workshop, focused on building skills, attracted 55% of respondents.

“A student-led discussion platform to share perspectives across JD and LLM cohorts” attracted 35% of respondents. While still significant, this initiative drew less interest overall.

Conclusions

The findings presented above allow for several data-supported inferences about how members of the Columbia Law School community are encountering and interpreting AI.

Across familiarity categories, students consistently reported higher exposure than faculty or staff. 89% of JD students and 98% of LLM students reported at least some familiarity with AI tools. Students also reported the highest "very familiar" rates: 22% of JD students and 39% of LLM students, compared with 19% of professors and 8% of staff. Given these differences, one can infer that students are driving most of the AI adoption at Columbia Law School.

50% of LLM students use AI tools daily, compared with 24% of JD students,24% of professors, and 17% of staff. Weekly use is also high among students: 41% of JD students and 41% of LLM students. Only 2% of LLM students reported "never" using AI, compared with 10% of JD students, 19% of professors, and 19% of staff. These patterns support the inference that AI is already a routine part of many students' workflows

Across all respondent groups, ChatGPT is the overwhelmingly preferred tool. 78% of all respondents selected ChatGPT as one of their top three tools. Among LLM students, usage was even higher at 89%. Given these disparities, one can infer that students rely primarily on general-purpose AI for academic and personal tasks, and that legal-specific AI tools have not yet reached widespread adoption.

Students use AI across more contexts than faculty or staff. 65% of JD students and 73% of LLM students reported using AI for personal tasks. 47% of JD students and 61% of LLM students used AI for academic coursework or research.

Although the survey did not quantify trust numerically, the distribution of responses shows that few respondents selected the highest trust categories. In fact, most clustered in the middle of the scale. Open-text responses frequently cited concerns about accuracy, bias, and hallucinations. Given that 28% of respondents use AI daily

and 38% weekly, yet trust remains moderate, one can infer that AI is viewed as useful but unreliable without human verification.

Across questions, respondents expressed strong interest in workshops, faculty-supported instruction, clearer policies, and structured opportunities to learn about AI. Because these requests come from groups already using AI frequently, one can infer that the community is not resistant to AI but is seeking clarity, consistency, and guardrails.

I. An Analysis of Open Responses: "What is your biggest frustration with AI?"

Open-text responses revealed a wide range of frustrations with AI, which cluster into several recurring themes.

A. Accuracy, Hallucinations, and Verifiability

A dominant theme across responses concerns AI's reliability. Many respondents pointed to "hallucinations" and "sycophantic response patterns" as frustrations. Several noted that the number of prompts required to obtain a correct answer "takes more time than using AI would save," particularly because "its outputs need to be checked for accuracy, so that they are not hallucinations."

These frustrations collectively contribute to concerns about "not knowing which platforms to trust." Respondents emphasized the burden of verification, with one writing: "Everything must be double/triple-checked and not simply relied on point blank. It is a great research tool, but it needs to be treated the same as search engines and other sources in that it is very useful, but not perfect. Individuals still need to be responsible for verifying all of the information in the outputs."

B. Legal-Specific Limitations

Respondents also highlighted frustrations unique to legal education and practice. Some noted that AI cannot perform Bluebooking well. Others expressed concern that

the legal profession's "whole way of handling and storing files doesn't seem optimized for AI."

Limitations in legal research were also noted, including the observation that "large language models often cannot detect if case law supports your issue or not. It is simply not intelligent enough to do more than note if it's in a similar realm as your legal question."

C. Academic Integrity, Detection, and Identity

Concerns about AI detection tools and their impact on academic integrity emerged repeatedly in the responses. One wrote: "especially in our profession where it's important to have strong writing skills, I'm concerned that people may suspect my writing is AI just because it has certain 'tells.'"

Students also described tensions between institutional messaging and classroom expectations. One respondent noted that AI-enabled productivity messaging in Legal Practice Workshops was "in conflict with Columbia Law School messaging." Another expressed frustration that "my classmates will use it even in ways that violate school or class policies—for example, Legal Practice Workshop memos and cold calls (i.e. during class when its use is explicitly banned). It's weird that its illicit use is so de-stigmatized to the point that my classmates openly flout such uses, or our TAs look the other way when AI is used in a way that brings up questions of academic honesty."

Some respondents described feeling caught between professional expectations and unclear academic norms surrounding AI use. Several expressed frustration with the absence of coherent instruction on how to use AI responsibly in law school, even as legal employers increasingly encourage its adoption. As one respondent put it, they are frustrated by "knowing firms are pushing its use but not being informed in school about how we should go about that correctly," a sentiment echoed by others who noted that "people [tell] us not to use it instead of teaching us how to use it correctly."

D. Societal-Level Risks and Existential Concerns

Some respondents articulated broader anxieties about AI's societal implications. These included concerns about AI's "deregulatedness" and "the polarity of opinion concerning its potential," as well as the effects of misinformation, disinformation, and deepfakes. One respondent warned: "Society believes the tools are far more advanced than they are, which could lead to reliance on faulty information."

Others expressed more existential fears: "I am deeply concerned that AI will take away fundamental aspects of the human experience and will threaten human society." Another noted that discourse around AI "leans academic/philosophical/economic/etc. instead of environmental/discriminatory/etc." Another described frustration with "the propaganda surrounding it."

II. Attention to Equity

Two constructed variables were developed to support an equity-focused analysis; however, they were not incorporated into the final data analysis. Although related findings were not presented in this study, these perspectives remain essential for interpreting the results and guiding future interventions.

A. First-Generation Professional (FGP) Identifier

Respondents who self-identified as FGP were tagged to allow examination of structural access patterns across groups.

Because first-generation professionals often navigate educational, financial, and professional systems without the inherited networks or institutional knowledge available to their peers, their pathways into advanced study and legal practice can be shaped by distinct structural barriers. Attuning to these differences is essential for any equity-driven analysis, ensuring that patterns in AI engagement are interpreted with an understanding of the uneven terrain from which respondents arrive.

B. Affinity-Group Membership Indicator

Respondents who indicated membership in one of Columbia Law School's student-led affinity organizations were coded into a single composite category.

Because affinity groups often serve as critical support structures for students from historically marginalized backgrounds, this coding enables examination of patterns of access, opportunity, and constraint at the group level rather than across smaller sub-populations. Attending to these dynamics is essential for an equity-minded analysis, as participation in affinity organizations can reflect shared experiences of structural barriers for navigating them.

Closing Summary

Across all categories, respondents' frustrations converge on a common theme: uncertainty. There is widespread frustration about accuracy, norms, institutional expectations, and AI's broader societal trajectory.

Students are clear about the interventions they want Columbia Law School to pursue. There is critical interest in workshops, faculty-supported instruction, and consistent policies that reduce uncertainty about when and how AI can be used. Underlying these requests is a desire for the institution to model responsible use and to ensure that all members of the community, regardless of background or prior exposure, can navigate the rapidly evolving AI landscape with confidence and clarity.

Contributions

We thank Professor Lester for her continued guidance and advice. We are also grateful for the collaboration and support of the Student Senate including Dean Saavedra, Advisor Carr, President Wolosyn, and Vice President Kerns. None of this would be possible without the opportunity provided through the Davis Polk Leadership Initiative under the leadership of Director Eunice Hong, Richard Gray, and Melanie Canela

Finally, we extend our sincere thanks to the Columbia Law School community for participating in this survey and for remaining engaged in these ongoing conversations. We are grateful for the opportunity to serve this community's interests and needs.

Sources

American Bar Association. (n.d.). Legal education publications. https://www.americanbar.org/groups/legal_education/publications/

American Bar Association. (2024). AI law school deans survey. https://www.americanbar.org/news/2024/ai-law-school-deans-survey/

American Bar Association Journal. (2026). Law librarians at the forefront of using and educating on AI. https://www.americanbar.org/groups/journal/articles/2026/law-librarians-at-the-forefront-of-using-and-educating-on-ai/

Cengage Group. (2025, June 27). AI's impact on education in 2025. https://www.cengagegroup.com/news/perspectives/2025/ais-impact-on-education-in-2025/

Center for Democracy & Technology. (2025). Hand in hand: Schools' embrace of AI. https://cdt.org/insights/hand-in-hand-schools-embrace-of-ai/

Clio. (2025). AI tools and software for lawyers: Improving productivity in law firms (p. 1).

CNBC. (2024). Constellation to restart Three Mile Island Unit 1 in 2028 under 20-year deal with Microsoft.

Columbia Law School. (n.d.). Faculty experts: Generative AI, legal education, and the future of the profession. https://www.law.columbia.edu/news/archive/faculty-experts-generative-ai-legal-education-and-future-profession

Columbia Law School. (n.d.). Why Columbia. https://www.law.columbia.edu/admissions/why-columbia

Davis Polk Leadership Initiative. (n.d.). About the Leadership Initiative. https://leadership-initiative.law.columbia.edu/content/about-leadership-initiative

Djanegara, N. D. T., Zhang, D., Badi Uz Zaman, H., Meinhardt, C., Watkins, G., Nwankwo, E., Wald, R., Kosoglu, R., Koyejo, S., & Elam, M. (2024). Exploring the impact of AI on Black Americans: Considerations for the Congressional Black Caucus's policy initiatives (p. 12). Stanford Institute for Human-Centered Artificial Intelligence.

Niederhoffer, K., Kellerman, G. R., Lee, A., Liebscher, A., Rapuano, K., & Hancock, J. T. (2025, September 22). AI-generated "workslop" is destroying productivity. Harvard Business Review.

Kaltman Law. (n.d.). AI detectors in academia. https://kaltmanlaw.com/ai-detectors-academia/

Poushter, J., Fagan, M., & Corichi, M. (2025). How people around the world view AI (p. 5). Pew Research Center. https://www.pewresearch.org

Scribbr. (2024–2025). AI detector accuracy study. https://www.scribbr.com/ai-detector/best-ai-detector/

Stanford Institute for Human-Centered Artificial Intelligence. (2025). AI on trial: Legal models hallucinate in 1 out of 6 (or more) benchmarking queries.

University of Chicago Law School. (2023–2025). Generative AI policy. https://www.law.uchicago.edu/students/handbook/generative-ai-policy

University of Kansas School of Law. (2025). Integrating generative AI in legal pedagogy: AI for Lawyers course. https://www.cambridge.org/core/journals/international-journal-of-legal-information/article/integrating-generative-ai-in-legal-pedagogy