

FEATURES

AI and Legal Scholarship

June 1, 2026




by Kenny Berkowitz



that's already been completed.

In one study, a law professor and two computer scientists explored the difficulties that legislators face in defining AI-related copyright infringement. In another, a team from Cornell Bowers, MIT, and Stanford showed how delays in reporting patients' race and ethnicity hide disparities in health-care delivery. In a third, researchers from Cornell Law and Universiteit Leiden described how vulnerable communities are already being harmed by the commercialization of artificial intelligence, with “worse problems highly likely in the future.”

“When lawyers, engineers, and social scientists work together, they see different parts of the same problem,” says Jed Stiglitz, who directs the Center alongside his work as associate dean for academic affairs and the Richard and Lois Cole Professor of Law. “AI’s challenges don’t fit neatly into existing disciplines. They implicate the intersection of law, computer science, ethics, social science, and policy.”




The Center, which opened in September 2025, is home to two main modes of research: faculty like Stiglitz and Yun-chien Chang, who use AI as a tool to answer questions about the law, and faculty like Jessica Eaglin, James Grimmelman, Sarah Kreps, and Frank Pasquale, who investigate AI itself, probing ways the technology is transforming law and society.

In his work on AI, Stiglitz explores legal reasoning and machine learning, with recent articles on large language models that compile texts written over extended periods of time. His “Modeling Legal Reasoning” (2023) used AI to uncover qualitative changes in the rationale for Supreme Court opinions, and was followed by “Historical Trends in Macro-Jurisprudence” (2024), which analyzed randomly chosen paragraphs from the Court to chart changes in legal philosophies between Reconstruction and the present. A more recent paper, “Understanding Change in Jurisprudence” (2026), analyzed how the justices’ legal philosophies changed over their careers.

he drafted the Constitution. “We know the books that Madison owned, and we’re creating a program that can answer questions as though you’re sitting right there in his library, reading the same books he was reading,” says Stiglitz. “Essentially, we’re swimming in the same waters as Madison, thinking about legal questions within the intellectual currents of the day.”

Telescoping beyond Madison, Yun-chien Chang has been using artificial intelligence to compare constitutions and legal systems all over the world. In “Drawing the Family Tree” (2020), he and his co-authors—including Martin Wells, the Charles A. Alexander Professor of Statistical Sciences at Cornell Bowers—analyzed 170 dimensions of property law in 129 jurisdictions, creating new methodologies to pioneer the field of empirical comparative law. Five years later, in “The Genesis of Constitutions” (2025), along with a team of scholars, they built a large language model of 588 constitutions written between 1789 and 2020 to track the influence of a group of historic “core” constitutions and reached some unexpected results.



“With unsupervised machine learning, you do not tell the machine anything,” says Chang, the Jack G. Clarke Professor in East Asian Law and Director of the Clarke Program in East Asian Law & Cul-ture. “You only feed it zeros and ones and ask it to reveal the pat-terns. If you are a human scholar, you are immersed in a human tradition, and you’re very unlikely to cut those ties. But AI has no fear of contesting the tradition, it only knows the data presented to it. Because we use an algorithm to sort out these legal families, and because we are able to do that without the human bias of a hun-dred years of scholarship, the result is not just our opinion. It is the data, and the data speaks for itself.”

sentencing in the era of mass incarceration, where she's become a leading expert. In "On 'Color-blind' and the Algorithm" (2024), Eaglin described the use of algorithms "expanding like wildfire" among police, pretrial administrators, judges, and parole administrators, and cataloged legal interventions to address concerns about bias across the country.



There are a lot of examples to demonstrate that if you don't use AI, you will become slower than your competitors, no matter whether you are in practice or in academia. If you don't use AI, you're losing the opportunity to look at old topics in a completely new light. I know AI feels like a new world, but we need to embrace it whole-heartedly, albeit with caution.

— Yun-chien Chang, Jack G. Clarke Professor
in East Asian Law and Director of the
Clarke Program in East Asian Law & Culture



The next year, in "Opening the Black Box" (2025), Eaglin wrote about rising concerns over AI in the decade after the death of Michael Brown in Ferguson, Missouri, and the nationwide protests that followed. Suggesting that the period may be "foundational" in legal discourse about both policing and technology, she urges scholars and policymakers to embrace an

“While many jurisdictions have yet to implement AI-driven decision-making systems, that’s clearly the direction they’re heading,” says Eaglin. “Without question, AI is coming to criminal law. Legal actors are scrambling to find new ways to use AI, and researchers are trying to grasp the impact it’s going to have in reshaping law and policy. As legal scholars, it’s our responsibility to ask about the implications of using AI. Yes, we understand it can do some tasks faster, but is that really what we want? Should judges be using technology to pass judgment on individual defendants? These are the kinds of questions my research raises.”



What makes me hopeful? The promise of a space like the Center to raise up diverse work around AI and the law. It will connect people who are thinking about these issues with real-world practitioners. AI is here already, and it’s bound to expand. The Center is committed to giving us the tools we need to think about it, to understand it, to learn how to grapple with it.

— Jessica Eaglin, Professor of Law



hypothetical trickster named Judge Loki to dramatize




supposed to serve.” Law clerks and judges shouldn’t be using generative AI to draft their opinions—but if they do, they need to at least (a) understand the limits of the technology and (b) base their opinion on multiple variations of their prompts, querying competing apps and comparing the results to refine their own human thinking.

“Generative AI systems may seem to be very compelling, but they have many problems,” says Pasquale, professor of law, who teaches at both Cornell Law and Cornell Tech. “There are problems of hallucinations, problems of bias, problems of instability. In just one example, there’s research that shows if you give the same materials and the same questions to the same large language model, it may rule for the plaintiff the first time and the defendant the second time. And even if you solve all the problems of hallucination, opacity, manipulability, and instability, you’d still have others. You can use the technology as a tool, but it should never replace the core judicial function of understanding arguments from both parties.”

Widely known as the author of *The Black Box Society: The Secret Algorithms That Control Money and Information* (2016), *New Laws of Robotics: Defending Human Expertise in the Age of AI* (2020), and *Data Access and AI Explainability* (2025), Pasquale continues to expand his research. In just the past two years, he’s authored or co-authored articles about consumer use of AI agents to evaluate contracts, AI evaluations of workers, compensation for content creators whose work is appropriated by AI, alternatives to credit scores, the risks and benefits of automated legal systems, fallacies of the robot rights movement, the use of artificial intelligence in benefit claims, and the risks of AI-generated misinformation on democracy.

and technology; her upcoming *Harnessing Disruption: Building the Tech Future Without Breaking Society* (2026) offers a road map for navigating the future of technology that draws on her research. In 2018, four years before the public launch of ChatGPT, Kreps became one of OpenAI’s earliest collaborators, researching the platform’s potential impact on democracy. In a preliminary test, published as “All the News That’s Fit to Fabricate” (2020), she devised three experiments to measure whether AI-generated texts could appear credible enough to influence opinions on foreign policy.

The results? Yes, readers were generally unable to distinguish between bot-written and human-written texts. Yes, readers responded most favorably to stories that mirrored their own partisan positions. Yes, small models could produce misinformation as effectively as large models, and yes, none of the study’s models needed human intervention to improve their credibility.



AI companies are very invested in convincing everybody that all the issues around AI are going to be resolved quickly. Maybe they're right, and maybe they're not. Nobody knows.

— James Grimmelmann,
Tessler Family Professor of
Digital and Information Law

”

“We definitely proved the concept, and we were early in showing that AI can generate huge amounts of misinformation at scale,” says Kreps, an associate member of the Cornell Law faculty, the John L. Wetherill Professor in the Department of Government, and the director of the Cornell Tech Policy Institute at Cornell Brooks School of Public Policy. “Along with those questions, I was looking at the effects on democratic representation: If people believed these AI-generated stories, would they stop believing anything else? Well, I think that’s where we are now. There’s so much AI slop that people assume everything is AI, and people have no trust unless something comes from one of their preferred partisan sources. That’s very bad for democracy and democratic institutions.”

In trying to gauge the relationship between democracy, law, and technology, James Grimmelmann returns regularly to an online map of AI copyright lawsuits that includes Adobe, Anthropic, Apple, Github, Google, Microsoft,

United States, along with three contrasting decisions on “fair use.”

For Grimmelmann, the Tessler Family Professor of Digital and Information Law at Cornell Law and Cornell Tech, that’s progress. “Sometimes we find that existing laws don’t work,” he says. “It’s been happening for hundreds of years. There was an explosion in torts and rail-road law in the 19th century. The courts often filled the gap between what the legal system knew how to handle and the problems thrown up by new technologies. Sometimes new technologies call for legislation and regulation, but I tend to favor case-by-case learning, which is an evolutionary process where courts repeatedly confront variations until they find their way.”



As an example, he points to the foundational *Zeran v. AOL* (1997), which established immunity for internet service providers in cases where wrongs were committed by their users. (“It’s a great way to start the conversation about how AI should be regulated,” he says, “even if there are strong arguments against the decision.”) He turns to his own research too, including “Generative Misinterpretation” (2026), which found a “superficial fluency” in AI-generated arguments, and “Talkin’ ’Bout AI Generation” (2025), which begins with a prompt to create an image of “a weasel on the first day of high school,” and ends with a question about the future of AI: “Which way from here?”



lives, but it will take a lot of planning, regulation, and legal intervention to make it happen. By no means should we convince ourselves AI is moving too fast for us to do anything. It's exactly the opposite. AI design decisions that are being made now are going to have ramifying, exponential impacts down the road.

FRANK PASQUALE, PROFESSOR OF LAW

At this point, with Grimmelmann preparing the sixteenth edition of *Internet Law: Cases and Problems*, the answers are complex and the research ongoing. “I don’t know how much better generative AI is going to get,” he says. “There are people who confidently predict it will keep continuing to improve. But the technology might already be plateauing. If AIs don’t get better, then people will grow more cautious in how they use them, and usage will drop.

“There are big questions we’re still trying to answer,” continues Grimmelmann. “How will computers and AI shape the work of law? How do we use technology in ways that are good for society and for the legal

exploring this new intellectual landscape, and that’s what gets me to my desk in the morning.” ■



TAGS: [spring 2026](#)



◀ PREVIOUS ARTICLE

Light Speed: Cornell Law Opens a Cross-Disciplinary Center for Law and AI

NEXT ARTICLE ▶

AI and Legal Education

In This Issue



FEATURES



FACULTY ESSAY

Law Schools Should Teach How to Integrate AI Tools Into Practice

[Facebook](#) [X](#) [Instagram](#) [LinkedIn](#)

[© 2026 Cornell University](#) [Privacy Policy](#) [Accessibility](#) [Equal Education & Employment](#)