



AI IN THE COURTS: How worried should we be?

Scholars and technologists see both benefits and dangers for AI in the courts. One thing they agree on: AI is here to stay.

As we enter 2024, it's tough not to think of 2023 as "the year of artificial intelligence." After all, last year saw the wide dissemination of ChatGPT (launched at the end of November 2022 by OpenAI), a free-to-use, large language model chatbot built to generate dialogue in response to human inquiry.¹

Unlike our old friend Google, a construct of 1998 that seems quaint by comparison, ChatGPT does not provide a list of results based on a web search. Instead, as a form of generative AI, it provides answers to prompts by drawing from knowledge through machine learning, or the process by which computers learn from examples.² The result is textual, human-like answers that are often detailed and context-specific.³ ChatGPT can produce essays, poems, computer code,⁴ and — yes — contracts, legal briefs, and a host of other documents relevant to the legal community.⁵

The legal industry, like many others, spent 2023 in a flurry of reactive

activity: Law schools amended honor codes to address AI-assisted learning,⁶ judges issued standing orders on AI-assisted briefing,⁷ and lawyers wondered how to harness the new power to research legal issues and even brainstorm strategy.⁸

So, where are we in 2024? We asked MAURA R. GROSSMAN, a professor in the School of Computer Science at the University of Waterloo; PAUL W. GRIMM, a retired federal judge and the David F. Levi Professor of the Practice of Law and Director of the Bolch Judicial Institute at Duke Law

School (which publishes *Judicature*); and CARY COGLIANESE, a professor at the University of Pennsylvania Carey School of Law and director of the Penn Program on Regulation, to discuss the pros and cons of AI in the legal space as we enter this brave new world.

Views about AI tend toward extremes: Either it will save the world from many of its current challenges, or it will destroy humanity as we know it. Where do you stand? Are you generally more positive or more negative about AI's potential impact, especially on the legal system?

Right now, the applications are many and the guardrails few. We have serious concerns about the use of untested, invalid, or unreliable AI systems, "function creep," discriminatory and inequitable outcomes, and the general hacker's philosophy of "move fast and break things."

[GRIMM/GROSSMAN]



GRIMM/GROSSMAN: We are enthusiastic about AI’s many potential positive benefits. But we also believe in the old Russian adage “trust but verify.” AI is a tool that can be wielded for good or for evil depending on how it is used and the safeguards that are placed around it. Right now, the applications are many and the guardrails few. We have serious concerns about the use of untested, invalid, or unreliable AI systems, “function creep,” discriminatory and inequitable outcomes, and the general hacker’s philosophy of “move fast and break things.”

We are most concerned about the ills we can already see, such as biased data and algorithms leading to discriminatory outcomes and greater inequality; the proliferation of misinformation and disinformation, which will threaten our judicial system, if not our entire democracy; increased crime and fraud as a result of easily created and hard-to-detect deepfakes; and increased threats to personal privacy through the accumulation of massive amounts of personal information in the hands of a few, unregulated big-tech companies with unabashedly selfish commercial interests. We are less troubled right now by existential risk, which seems possible only if AI is connected to a permanent energy source — so it cannot be unplugged or can learn to replicate itself; an unlimited financial source — so it can fund its activities without human oversight; or the introduction of lethal autonomous weapons — so it could make “kill” decisions independent of human involvement.

COGLIANESE: AI won’t be perfect, but the aim should be to have it do more good than bad — and to make the world better, on balance, than it is today. Moreover, any question about how

The challenge facing society is to ensure that the design, development, and deployment of AI will do more good than bad — and that it improves the status quo. . . . Along the way, it’s worth remembering that a world dependent solely on humans is imperfect, too.

[COGLIANESE]

“good” or “bad” AI will be cannot be answered across the board. AI is not a singular technology. It’s a proliferation of many varied technologies put to many varied uses. The types of AI algorithms vary, as do the datasets on which they train. Most importantly, the ways that AI algorithms are used vary widely. Some of these uses can be very good, such as detecting cancers or curing diseases through precision medicine. Other uses are good even if seemingly banal, such as helping the U.S. Postal Service sort mail by reading addresses on letters and packages.

AI can also be put to bad uses, such as fomenting political strife through misinformation campaigns or creating fraudulent images or documents. Even then, good AI tools may help spot the frauds and filter out the misinformation.

The highly varied uses for AI tools make it impossible to paint with a broad brush and declare that “AI is good (or bad).” Furthermore, the reality is that AI is here to stay. The challenge facing society is to ensure that the design, development, and deployment of AI will do more good than bad — and that it improves the status quo. This is where regulation comes in. Society needs ways to govern AI that can equitably reap its benefits while reducing its harms. If we can do that, then we can use AI to make the world better. Along the way, it’s worth remembering that a world depen-

dent solely on humans is imperfect, too. The key is to do better.

AI’s uses in the justice system are many. We have seen AI used in predictive policing, electronic discovery, evidentiary matters, sentencing, and actually helping to decide cases. What uses ought we be most optimistic about? What uses ought we be most concerned about? Are there any legal instances where AI should never be used?

GRIMM/GROSSMAN: We have long been proponents of using AI for electronic discovery and, frankly, we are baffled as to why — after a decade of sound empirical evidence — there is still hesitation to use technology-assisted review (TAR) to substantially reduce the time, cost, and burden of document review. We are excited about the prospect of using AI to increase access to justice, to help self-represented litigants determine whether they have a viable legal claim or defense, and to draft pleadings that properly address the jurisdictional, venue, and substantive elements required to state a proper claim. Our enthusiasm also extends to its use by attorneys to increase productivity and efficiency — provided that all AI-created pleadings are verified for accuracy before filing (both as to facts and as to citations). We can also see the benefits of online adjudication systems for small claims, housing, and traffic cases, where justice delayed is often justice denied.

Use of a faulty AI software application that recommends a bad movie to watch simply raises different concerns than using a faulty AI application for the purpose of determining the length of a criminal sentence to be imposed.

[GRIMM/GROSSMAN]

We worry a bit about self-represented litigants who may use AI to abusively flood the courts with meritless lawsuits, clogging the system and overwhelming court personnel and judges. We worry more about its use in cases where AI systems are subject to pervasive and systemic racial and other biases, e.g., predictive policing, facial recognition, and criminal risk/recidivism assessment. We lose the most sleep over a court relying on AI applications or evidence as the *primary* or *sole* evidence supporting a consequential outcome (such as the imposition of a criminal sentence) without sufficient assurances that the AI system has been trained on data and/or programmed so its validity and reliability can be demonstrated empirically (i.e., it accurately does what it was designed to do and it consistently produces accurate results when applied to substantially similar facts, respectively); and that it is equitable, unbiased, and/or fair (assuming that we can reach consensus on what it means for an algorithm to be “equitable,” “unbiased,” and/or “fair”). When one or more of those conditions cannot be met, we believe that AI evidence should be excluded.

While in some situations relying on AI is at best ill advised, we are loath to ban its use in toto for most applications. AI systems need to be considered

in light of their benefits and risks, the quality of the available alternative processes for decision-making, and the consequences of making a wrong decision because one or more of the conditions we specify have not been met. Use of a faulty AI software application that recommends a bad movie to watch simply raises different concerns than using a faulty AI application for the purpose of determining the length of a criminal sentence to be imposed.

COGLIANESE: It’s hard to say in the abstract that AI ought *never* be put to some uses. To be sure, any uses that are abhorrent if conducted without AI will still be abhorrent when conducted with AI. And there are other uses where, due to current limitations in data or algorithmic designs, AI is not ready for prime time. But judgments about whether specific AI tools are too biased, unjust, or unsafe will need to be made on a case-by-case basis and they’ll never be permanent or absolute. The technology is changing rapidly. But the appropriate test in all cases should be how well AI performs compared with the status quo.

With that comparative perspective in mind, we probably ought to be careful before concluding that some uses are just “too risky” ever to allow AI to handle them. If we think some

uses are too risky for AI, then presumably they’re risky without AI. Human decision-making is prone to bias and error, too. If AI tools can be shown to perform risky tasks better than humans, we ought to be open to considering AI.

Already, we see AI doing some amazing things. In November 2022, ChatGPT’s version 3.5 took the world by storm, but still only scored at the 10th percentile on the uniform bar exam. By March 2023, though, when OpenAI released ChatGPT version 4.0, this AI tool not only passed the uniform bar exam — but did so at the 90th percentile!

Still, as any lawyer can surely testify, the practice of law is not the same as the bar exam. Humans still outperform AI on tasks that call for creative problem-solving and out-of-the-box thinking. AI depends on very large sets of data to perform pattern recognition and forecasting. Although it can perform many of these kinds of tasks very well — even beating humans at detecting fraud, predicting recidivism, and finding errors in documents⁹ — many tasks will remain that humans do best. Truly *sui generis* judgments cannot be decided by AI tools, for example, even though they come before legal institutions with considerable frequency. And AI tools cannot make ultimate value judgments — however good they get at performing other sophisticated tasks. I agree with the thrust of Chief Justice John Roberts’s views about AI — expressed recently in his year-end report on the judiciary — that there will long remain important work for lawyers and judges to do.¹⁰

The challenge ahead will be to find the best ways for humans and computers to collaborate. We may also need to reimagine the work of lawyers, judges, and other personnel in our courts ►

and bureaucracies. If we can lighten humans' share of paperwork processing, routine order drafting, and other regular tasks, maybe we can unleash humans to do more of what they distinctively excel at. I'd like to see a world in which courts and bureaucracies provide greater human empathy and compassion. I'd also like to see them provide empathic support that is more accessible, consistent, and unbiased. Maybe the path to this more humane future will, ironically, depend on a technology that disrupts how we conceive of human effort in the legal profession.¹¹

AI offers exciting access-to-justice possibilities. AI lawyers¹² and AI judges,¹³ for instance, have already been employed, primarily abroad, to make the courts more accessible for individual litigants. What do you think about this new horizon?

GRIMM/GROSSMAN: When AI is used to assist self-represented litigants in drafting factually accurate pleadings that address the required elements for stating a claim, or when AI is used to create greater efficiencies for paying clients, this obviously has the potential to increase access to and reduce the staggering costs of justice. Similarly, a judge or judicial clerk who uses AI as an *initial* means of finding controlling authority to resolve a dispute — following that up with independent research and personal consideration of the facts — may be able to more quickly issue opinions and reduce court backlogs. Online legal and judicial resources are generally a positive development. They make it easier to access and navigate a complex and often painfully slow legal system. The key to the proper use of AI in the law is as a tool to *assist* litigants, counsel, and judges in performing legal tasks — not to *replace* them, including the independent professional judgment

they offer. We believe much remains to be said for having one's day in court, especially when the stakes are high.

COGLIANESE: People are already getting acclimated to AI in other spheres of their lives, which presumably will make them more comfortable with AI in their interactions with the legal system. Indeed, citizens may even come to demand the use of AI tools — especially if they are shown to lead to swifter, more accurate and consistent resolution of claims or disputes. Some survey research already reveals this public receptivity to consequential uses of AI. And private sector experience seems encouraging, too. eBay, for example, has a totally automated system for resolving disputes that reportedly leaves customers so satisfied that they return to eBay more frequently than customers who never had any disputes to resolve. Admittedly, the work private firms' online customer service complaint systems do is not exactly what courts do. But the point is that as AI-based automation gets woven into individuals' daily lives, it is likely that people will accept AI for uses traditionally performed by lawyers and courts. And if AI provides a comparable or superior vehicle for making traditional legal services and support more accessible to more people, we should all applaud the result.

ChatGPT — a recently popularized AI tool that can provide detailed responses to specific textual prompts — has caused disruption in many industries. Recently, for instance, a judge in Colombia used ChatGPT to assist in making a decision by feeding the AI tool a series of questions.¹⁴ What should the courts make of the tool's powers and possibilities? Should courts be using ChatGPT and, if so, how?

GRIMM/GROSSMAN: The courts must pay close attention to Generative AI (GenAI) tools like ChatGPT. GenAI — and, in particular, deepfakes — will unquestionably make future evidentiary issues far more challenging for both the bench and bar. Parties can be expected to present AI-generated evidence as genuine and accurate, and to challenge authentic evidence as deepfake. Cases involving GenAI are likely to require expensive forensic experts for the foreseeable future, and there is a real risk that juries will become increasingly skeptical of all evidence. ChatGPT already raised judicial eyebrows in 2023, when counsel in a federal case filed in the Southern District of New York used the tool in legal research and cited numerous nonexistent cases in a brief, leading multiple courts to promptly issue standing orders requiring disclosures and certifications when attorneys use GenAI tools to prepare pleadings.

Cases involving GenAI are likely to require expensive forensic experts for the foreseeable future, and there is a real risk that juries will become increasingly skeptical of all evidence.

[GRIMM/GROSSMAN]

While judicial officers might be tempted to use GenAI tools like ChatGPT in decision-making and drafting opinions, we have advocated for caution and restraint in this regard, until the tools are more trustworthy. The U.S. Constitution vests decision-making authority in human judges, not AI, and tools using GenAI are prone to error. There is simply no room for algorithmic hallucinations in judicial opinions. Judicial use of GenAI may also raise due process concerns if courts consider evidence or arguments presented by ChatGPT that were not presented by the litigants themselves.

COGLIANESE: In its canonical decision on procedural due process — *Mathews v. Eldridge* — the Supreme Court articulated a three-part balancing test for determining the fundamental fairness of a governmental process. One of the three parts — the interests of the individual — is entirely independent of any type of process, including one based on AI. But the other two parts — the accuracy of a process and the costs to the government — are ones that would presumably weigh in favor of using AI. These digital tools can make more accurate decisions than humans, and their use in automated systems

promises to lower the costs of adjudicatory services. On this basis, it's hard to argue that constitutional due process categorically precludes their use. Quite the contrary, someday adherence to constitutional values might very well demand the use of AI.

This is not to say that judges today should hand over decision-making to ChatGPT and the like. AI can still hallucinate. And uploading information on web platforms that are not safeguarded can raise privacy and ethical concerns. Caution is still in order. And just as some judges in the United States have issued guidance for the use of AI by lawyers appearing in their courtrooms, the judiciary in England and Wales has issued guidance for the use of AI by judges themselves.¹⁵ That guidance would be worthwhile for judges and their clerks in the United States to consider. Yet as technology gets more powerful and accurate, and as secure systems are developed to protect confidential information, I predict that judges around the world will make AI tools a regular part of their work. Just as judges have long relied extensively on law clerks to help them draft orders or opinions, and just as both judges and their clerks have relied for decades on electronic

databases in their research, they will eventually come to rely on large language models.

In *Wisconsin v. Loomis*, the Wisconsin Supreme Court held that the trial court's consideration, as part of the sentencing process, of an algorithmic risk assessment tool that estimates the risk of recidivism did not violate due process — even though the tool's internal methodology was not disclosed to the defendant who sought to challenge it. Much has been made of *Loomis* since it was decided in 2016. Do you agree with the decision? Should judges be able to rely on black-box algorithmic risk-assessment tools for decision-making? Relatedly, to what extent does AI have the potential to eliminate bias in our legal system versus to perpetuate existing biases? How important is explainability?

GRIMM/GROSSMAN: Reliance on AI for legal or adjudicatory functions must be conditioned on a satisfactory showing that the AI system is valid, reliable, and equitable, unbiased, and/or fair. Those three conditions can only be determined if the proponent of the use of AI (or of AI evidence) can demonstrate that it meets these criteria. Generally, that cannot be done unless the users, affected parties, and the court understand how the AI was developed, trained, operated, and achieved its results. This requires an appropriate level of transparency and explanation by the proponent of the AI system (or evidence), and a fair opportunity for the party opposing the AI system (or evidence) to understand enough to fashion a challenge to its use or admissibility.

Some courts have deprived parties seeking to challenge AI evidence of the opportunity to do so, or to undertake reasonable efforts to test it, because the AI developer resisted disclosure of information based on claims of propri-

If AI provides a comparable or superior vehicle for making traditional legal services and support more accessible to more people, we should all applaud the result.

[COGLIANESE]

etary trade secrets. *Wisconsin v. Loomis* is a good example of this. The court in *Loomis* reasoned that the defendant did not have a due process right to test the COMPAS AI program that calculated his likelihood of recidivism, in part, because the software used was only one factor considered by the judge, who had an independent obligation to determine the proper sentence given all of the evidence. But this fails to account for the fact that, without the defendant's ability to challenge the validity and reliability of the COMPAS AI program's prediction, the judge was entirely unequipped to assess the weight to be given to the algorithmic recidivism prediction. And balanced against the risk of imposing an unjustifiably long criminal sentence based, in part, on an erroneous AI-generated prediction, the alternative of allowing the defendant access to the program to test and challenge it — subject to a reasonable order protecting the developer's trade secrets — seems the better choice (and certainly the fairer one).

The better reasoned judicial decisions, of which there are a growing number, permit some level of discovery about the data, algorithm, functioning, and output of the AI system, subject to a protective order. The phrase "black box" typically refers to a system that lacks such transparency, but it can be a misleading concept. A judge need not see or even understand every nuance of how the AI system operates, so long as the proponent can explain the process by and circumstances under which it has been developed, trained, and — most importantly — tested. Without such testing, it is impossible to determine, for example, whether the AI system is more or less biased than any available alternative approach (including a fully human process). AI that has been

Machine-learning algorithms work in ways that are strikingly different than — and often counterintuitive to — the conventional statistics that judges may have learned in college. To use the new digital tools responsibly, and to pass judgment on their results as evidence put forward in litigated disputes, judges need to be armed with knowledge.

[COGLIANESE]

properly designed, developed, and deployed can be relied on for many legal purposes if exacting and independent auditing and validation of the system have occurred. Thus, the notion of "explainability" may more usefully be thought of as the ability to demonstrate that the AI system unequivocally meets the requirements of validity, reliability, and equity, lack of bias, and/or fairness, rather than as a description of its technical inner workings.

COGLIANESE: Lots of great questions! They deserve longer answers than we have space here, so maybe this is as good a time as any to refer to some relevant articles of mine (all available on SSRN): *Transparency and Algorithmic Governance; AI in Adjudication and Administration; From Negative to Positive Algorithm Rights; Moving Toward Personalized Law; and Procurement and Artificial Intelligence*.¹⁶ With that note, I happily offer brief answers here, with the understanding that anyone interested in more extended discussion will know where to find it elsewhere. Here goes:

Yes, *Loomis* was correctly decided on the basis of prevailing due process law. Procedural due process is a balancing act. It allows for innovations in processes as long as fundamental information is provided about how

decisions are made. Also, in *Loomis* the human judge was still kept "in the loop."

Yes, judges should be able to rely on algorithmic risk-assessment tools, provided they are well-calibrated, unbiased, validated, and sufficiently transparent. The tool in *Loomis*, it should be noted, was not a machine-learning algorithm — it was not black box in any intrinsic sense. Instead, it was black box because the outside firm that developed the risk-assessment tool claimed proprietary protection. That kind of black-box situation should and can be easily prevented. Courts that rely upon outside data analytic firms should insist, during procurement, on contractual assurances of robust transparency and adequate testing.

Yes, AI has the potential to eliminate bias — but yes, if used unthinkingly it also has the potential to perpetuate existing biases. I'm optimistic that AI ultimately will do more to eliminate bias. Why? Right now, too many biases are hidden. But because AI will only work with large datasets, using it will necessarily mean we have to collect a lot of data. When we have the data, we can begin to see unjust biases better and take steps to reduce them. Furthermore, the steps needed to de-bias AI will be mathematical. That's likely to be easier than rooting out the implicit biases in humans.

Until it is confirmed that the AI is the right tool for the job, that it can accomplish that job with sufficient accuracy and consistency, and that its input, functioning, or output is not subject to systemic bias, its use or acceptance as evidence should not be permitted in the justice system. Period.

[GRIMM/GROSSMAN]

And yes, explainability is important. But what counts as a sufficient explanation for an automated decision from a validated AI-based tool might not equate to a traditional judicial opinion. Moreover, let's not kid ourselves. As legal realists showed long ago, how *humans* explain *their* decisions doesn't necessarily equate to how they really reach those decisions. Perhaps human decision-making is as much of a black box as any.

Many universities have expressed serious concerns about students using large language models to complete their coursework. GPT-4 was able to pass the bar exam at the 90th percentile. Should law students be prohibited from using such tools to assist with their written work? Should faculty be prohibited from using them in preparing scholarly works for publication? What about junior associates using them for preparing pleadings?

GRIMM/GROSSMAN: As we move into an increasingly technology-enabled society and legal industry, we need to ensure that law students receive sufficient technical training and obtain sufficient digital literacy to practice effectively and efficiently in this new world. The rules of professional conduct in most states already require technical competence. Therefore,

students must be exposed to GenAI tools in law school so they understand their capabilities and limitations. That said, we cannot afford to teach budding lawyers to outsource to AI essential legal skills such as issue-spotting, critical thinking, and problem-solving. Those skills will remain necessary for all attorneys, even those who function with the assistance of AI adjuncts.

Law schools should provide opportunities to integrate AI into legal studies but also develop and assess students' reasoning and writing skills separate and apart from that. We see no problem with junior associates using GenAI as a starting point for drafting exercises — if they verify the accuracy and veracity of all AI output, including not only any facts but especially all case citations and other references.

COGLIANESE: When I was in law school, we were taught not to rely on headnotes, digests, and annotations as anything more than potentially helpful finding aids. I was also taught to be careful about using commercial outlines as study aids. We were given these warnings not merely because these materials were not the law but also because they could and did — and still do — contain inaccuracies. Undue reliance on them would also short-

change our opportunity to develop our skills and knowledge as lawyers. This same caution applies today with respect to AI. Law students and new attorneys still need to do the work themselves that's needed to learn to think analytically and write well. Indeed, only if they can do those tasks proficiently are they likely to be able to know how to use AI tools responsibly. There are no complete shortcuts.

What is the most important thing for judges to keep in mind about AI moving forward? What trends, promises, or pitfalls should they focus on?

GRIMM/GROSSMAN: It is important for judges to bear in mind that AI is not something to be feared. AI systems are merely computing tools designed to perform certain functions to augment or replace human effort. Those functions can be performed validly, reliably, and equitably — or not. The performance of such AI tools may exceed or lag behind fully human alternatives. Obtaining sufficient empirical information about the development and performance of AI systems is the critical gatekeeping role that judges must play.

Until it is confirmed that the AI is the right tool for the job, that it can accomplish that job with sufficient accuracy and consistency, and that its input, functioning, or output is not subject to systemic bias, its use or acceptance as evidence should not be permitted in the justice system. Period, full stop. The burden of proof rests with the proponent to demonstrate this, and the opposing party must be given a reasonable opportunity to challenge these assertions. Once the court has this information, it can weigh the benefit of using the AI system or evidence versus the prejudice or negative outcome that could occur if the AI is not ►

sufficiently valid, reliable, or fair, and make a just decision. When the risk or prejudice outweighs the benefits, the use of the AI system and the admission of AI evidence it produces should not be permitted. It is as simple and straightforward as that.

COGLIANESE: The most important thing judges can do is to increase their understanding of the mathematics behind AI tools. A variety of online courses are available, and there are and will continue to be accessible books and articles that judges can read. Machine-learning algorithms work in ways that are strikingly different than — and often counterintuitive to — the conventional statistics that judges may have learned in college. To use the new digital tools responsibly, and to pass judgment on their results as evidence put forward in litigated disputes, judges need to be armed with knowledge.

¹ Will Douglas Heaven, *The Inside Story of How ChatGPT Was Built From the People Who Made It*, MIT TECH. REV. (Mar. 3, 2023), <https://www.technologyreview.com/2023/03/03/1069311/inside-story-oral-history-how-chatgpt-built-openai/>.

² AIContentfy Team, *AI Writing vs. Traditional Writing: Pros and Cons*, AICONTENTFY (July 10, 2023), <https://aicontentfy.com/en/blog/ai-writing-vs-traditional-writing-pros-and-cons/> (“AI writing software is continuously improving by learning from vast amounts of existing written content, allowing it to generate increasingly accurate and contextually appropriate text.”).

³ Eric Griffith, *ChatGPT vs. Google Search: In Head-to-Head Battle, Which One Is Smarter?*, PC MAG. (Feb. 14, 2023), <https://www.pcmag.com/news/chatgpt-vs-google-search-in-head-to-head-battle-which-one-is-smarter>.

⁴ Beth McMurtrie, *ChatGPT Is Already Upending Campus Practices. Colleges Are Rushing to Respond*, THE CHRONICLE (March 6, 2023), <https://www.chronicle.com/article/chatgpt-is-already-upending-campus-practices-colleges-are-rushing-to-respondhow>.

⁵ Andrew Perlman, *The Implications of ChatGPT for Legal Services and Society*, HARV. L. SCH. CTR. ON LEGAL PRO.: PRAC. MAG. (Mar. 2023), <https://clp.law.harvard.edu/knowledge-hub/magazine/issues/generative-ai-in-the-legal-profession/the-implications-of-chatgpt-for-legal-services-and-society/>.

⁶ See, e.g., Karen Sloan, *University of California*

Berkley Law School Rolls Out AI Policy Ahead of Final Exams, REUTERS (Apr. 20, 2023, 4:28 PM), <https://www.reuters.com/legal/transactional/u-california-berkeley-law-school-rolls-out-ai-policy-ahead-final-exams-2023-04-20/>; Deborah Thompson Eisenberg, *ChatGPT Goes to Law School: Now What?*, MD. STATE BAR ASS’N (July 27, 2023), <https://www.msba.org/chatgpt-goes-to-law-school-now-what/>.

⁷ See, e.g., Standing Order Re: Artificial Intelligence (“AI”) Cases Assigned to Judge Baylson (E.D. Pa. 2023), <https://www.paed.uscourts.gov/documents/standord/standing%20Order%20Re%20Artificial%20Intelligence%206.6.pdf>; Standing Order for Civil Cases Before Magistrate Judge Fuentes (N.D. Ill. 2023), [https://www.ilnd.uscourts.gov/_assets/_documents/_forms/_judges/Fuentes/Standing%20Order%20For%20Civil%20Cases%20Before%20Judge%20Fuentes%20rev%27d%205-31-23%20\(002\).pdf](https://www.ilnd.uscourts.gov/_assets/_documents/_forms/_judges/Fuentes/Standing%20Order%20For%20Civil%20Cases%20Before%20Judge%20Fuentes%20rev%27d%205-31-23%20(002).pdf).

⁸ See, e.g., Bob Ambrogì, *New GPT-Based Chat App From LawDroid Is A Lawyer’s ‘Copilot’ for Research, Drafting, Brainstorming and More*, LAW SITES (Jan. 25, 2023), <https://www.lawnext.com/2023/01/new-gpt-based-chat-app-from-lawdroid-is-a-lawyers-copilot-for-research-drafting-brainstorming-and-more.html>.

⁹ See generally Cary Coglianese & Alicia Lai, *Algorithm vs. Algorithm*, 72 DUKE L. J. 1281, 1309–14 (2022) (discussing ways that machine-learning algorithms can outperform humans on a variety of tasks).

¹⁰ Chief Justice John G. Roberts, Jr., 2023 Year-End Report on the Federal Judiciary (Dec. 31, 2023), <https://www.supremecourt.gov/publicinfo/year-end/2023year-endreport.pdf>.

¹¹ Cary Coglianese, *Administrative Law in the Automated State*, 150 DÆDALUS 104 (2021).

¹² See Bobby Allyn, *A Robot Was Scheduled To Argue in Court, Then Came the Jail Threats*, NPR (Jan. 25, 2023, 6:05 PM), <https://www.npr.org/2023/01/25/1151435033/a-robot-was-scheduled-to-argue-in-court-then-came-the-jail-threats>.

¹³ See Tara Vasdani, *Robot Justice: China’s Use of Internet Courts*, LAW360 CAN. (Feb. 5, 2020, 11:07 AM), <https://www.law360.ca/ca/articles/1750396/robot-justice-china-s-use-of-internet-courts>.

¹⁴ Joanna York, *ChatGPT: Use of AI Chatbot in Congress and Court Rooms Raises Ethical Questions*, FRANCE 24 (Mar. 2, 2023, 8:49 PM), <https://www.france24.com/en/technology/20230203-chatgpt-use-of-ai-chatbot-in-congress-and-court-rooms-raises-ethical-questions>.

¹⁵ Courts and Tribunals Judiciary, *Artificial Intelligence (AI): Guidance for Judicial Office Holders* (Dec. 12, 2023), <https://www.judiciary.uk/wp-content/uploads/2023/12/AI-Judicial-Guidance.pdf>.

¹⁶ Cary Coglianese & David Lehr, *Transparency and Algorithmic Governance*, 71 ADMIN. L. REV. 1 (2019); Cary Coglianese & Lavi Ben Dor, *AI in Adjudication and Administration*, 86 BROOKLYN L. REV. 791 (2021); Cary Coglianese & Kat Hefter, *From Negative to Positive Algorithm Rights*, 30 WM. & MARY BILL RTS. J. 883 (2022); Cary Coglianese, *Moving Toward Personalized Law*, U. CHI. L. REV. ONLINE (2022) (available at https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4051776); Cary Coglianese, *Procurement and Artificial Intelligence*, HANDBOOK ON PUB. POL’Y & A.I. (forthcoming) (manuscript available at https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4591724).



PAUL W. GRIMM is a retired federal judge and the David F. Levi Professor of the Practice of Law and Director of the Bolch Judicial Institute at Duke

Law School (which publishes *Judicature*). Previously, he served as a district judge (and before that as magistrate judge) in the U.S. District Court for the District of Maryland.



MAURA R. GROSSMAN, JD, PhD, is a research professor in the David R. Cheriton School of Computer Science at the University of Waterloo, an adjunct professor at Osgoode Hall Law School of York University, and an affiliate faculty member of the Vector Institute of Artificial Intelligence.



CARY COGLIANESE is the Edward B. Shils Professor of Law, Professor of Political Science, and founding Director of the Penn Program on Regulation

at the University of Pennsylvania. The creator of *The Regulatory Review*, he specializes in administrative law and regulatory policy issues, with an emphasis on technology and regulation.