

# Enhancing Access to Justice Through Generative AI: A Preliminary Agenda (Feb. 2024)

*Christopher L. Griffin, et al.*

Further distribution restricted.  
Copyright 2024. All rights reserved.  
Reprinted with permission.



## **Enhancing Access to Justice Through Generative AI: A Preliminary Agenda**

*Christopher L. Griffin, Jr.,\* Cas Laskowski\*\* & Samuel A. Thumma\*\*\**

**February 1, 2024**

### **Abstract**

The artificial intelligence (AI) explosion has reached the legal profession. Attorneys, judges, scholars, and others who contribute to the justice system recognize the extraordinary potential of generative AI and, at time, fear its potency. Not only could humans lose jobs in the profession to algorithms, but we also still cannot fully appreciate the extraordinary risks that AI poses. In this essay, the authors—a law professor, a law librarian, and a judge—highlight both the promise and the peril of generative AI in the context of promoting access to justice. We do so, first, by asking an advanced chatbot to describe how it might advance quicker, less costly, and more inclusive dispute resolution. Different answers over just three months reveal how rapidly generative AI technology is evolving as well as the limitations of our—and the technology’s—understanding of AI’s role in access to justice. The essay then considers the most prominent, concerning shortcomings that generative AI presents as we work toward a more comprehensible and accessible civil justice system. The essay concludes with a primer on how to evaluate generative AI’s successes and argues that rigorous testing should accompany use of any AI-based tool, if not precede its launch.

---

\* Professor and Director of Empirical & Policy Research, University of Arizona James E. Rogers College of Law.

\*\* Head of Research, Data, & Instruction, University of Arizona James E. Rogers College of Law.

\*\*\* Judge, Arizona Court of Appeals, Division One, Phoenix, Arizona; Chair, Arizona Commission on Access to Justice and Co-chair of the Arizona Supreme Court COVID-19 Continuity of Court Operations During Public Health Emergency Workgroup (aka the Plan B Workgroup). The views expressed do not represent those of the Arizona Courts or the Arizona Court of Appeals.

## I. Introduction

By one definition, “[g]enerative artificial intelligence (AI) describes algorithms (such as ChatGPT) that can be used to create new content, including audio, code, images, text, simulations, and videos.”<sup>1</sup> Any form of AI, whether generative or “general,” meaning capable of processing information just like a human being, was once limited to the realm of science fiction. Although we have not created machines as advanced as the HAL 9000 in *2001: A Space Odyssey* (1968), or the human-like child played by Haley Joel Osment in *AI: Artificial Intelligence* (2001), the speed at which AI continues to evolve is staggering.

As noted in February 2023, “[a]fter years of research,” generative AI “is reaching a sort of tipping point, capturing the imaginations of everyone from students saving time on their essay writing to leaders at the world’s largest tech companies. Excitement is building around the possibilities that AI tools unlock, but what exactly these tools are capable of and how they work is still not widely understood.”<sup>2</sup> Now, a year later, a Google search for “ChatGPT” generates about 1.5 billion results. And that is just one of many platforms in the generative AI space.

The ostensible purpose of these technologies is to enhance our collective efficiency. Just as the Industrial Revolution heralded the replacement of human labor with automation, an AI-led transformation using powerful algorithms could save millions of hours of cognitive processing time. These tools are poised to transform any number of vocations, including the legal profession.<sup>3</sup> Attorneys could spend more time on client relations than contract drafting. Courts could identify better ways to help individuals through the legal system and resolving disputes. Self-represented litigants could navigate some legal problems without having to pay for representation.

Along with the extraordinary potential of generative AI, we should not lose sight of the extraordinary risks it poses. In this essay, we highlight both dimensions in the context of promoting access to justice. By access to justice, we mean any practice that helps litigants—especially in the nation’s civil courts—resolve their legal matters with minimal or no formal attorney representation. We also include efforts that help potential litigants avoid having to invoke the legal system in the first place as well as ways in which courts and other stakeholders can improve the legal system to better serve the public.

We start in Part II by outlining generative AI’s most promising features, recognizing that generative AI is so new that it is hard to offer more than a tabletop exercise of how it might enhance access to justice. We then address in Part III concerns about using generative AI to advance access to justice and assist self-represented litigants. Finally, we discuss in Part IV how to measure success in using generative AI to bridge the justice gap. At the end of the day, great

---

<sup>1</sup> *What is generative AI?*, McKinsey & Company (Jan. 19, 2023), <https://www.mckinsey.com/featured-insights/mckinsey-explainers/what-is-generative-ai> (last visited Oct. 9, 2023).

<sup>2</sup> *Artificial Intelligence, What is Generative AI? An AI explains*, World Economic Forum (Feb. 6, 2023), <https://www.weforum.org/agenda/2023/02/generative-ai-explain-algorithms-work/> (last visited Oct. 9, 2023).

<sup>3</sup> See, e.g., John Villasenor, *How AI Will Revolutionize the Practice of Law*, Mar. 20, 2023, <https://www.brookings.edu/articles/how-ai-will-revolutionize-the-practice-of-law/>.

care is needed in using generative AI to enhance access to justice, to ensure its long-term success, and to address a host of valid concerns.

## **II. Generative AI’s Potential to Enhance Access to Justice**

Predicting how generative AI will affect access to justice is difficult, mostly because the underlying technology is comparatively new and rapidly evolving. A simple example proves the point.

What follows is a very rough test in which we asked an AI chatbot to self-evaluate. Given the terabytes of information at its disposal, generative AI has access to vastly more data points than any of us could process in a lifetime. So, over the span of nearly three months, we asked the publicly available Bing AI search engine<sup>4</sup> the following question: “How can artificial intelligence help advance access to justice”? A comparison of the results is telling.

On August 14, 2023, the response was not all that instructive or optimistic.<sup>5</sup> First noting that AI “can help improve access to justice in many ways,” the response provided generalities like “a more responsive justice system”; “augmenting and even replacing lawyers”; and “provid[ing] a more just legal outcome than a human.” These responses are fairly standard reflections of what access to justice is supposed to deliver, with or without generative AI. The initial response was quickly followed by a proviso that technical advances in the law had not made services cheaper and more accessible, largely because of “the law’s apparent impenetrability.” The response ended on a more hopeful tone, suggesting that AI could help provide legal services at a lower cost to a larger number of people for two reasons: “Firstly, it can support the provision of legal services; and secondly, it can replace the role of legal experts. Legal technology that supports justice includes natural language processing (NLP), machine learning and chatbots.” All these observations are, generally speaking, true; they also seem comically simple to anyone who devotes their scholarly or practice-related time to access to justice.

On November 4, 2023—almost three months later—we replicated the search with the identical prompt on the same platform. What followed was far more instructive and helpful, providing a response that maps broadly onto three categories.<sup>6</sup>

First, Bing AI responded: “AI can **increase efficiencies** by automating tasks such as document preparation, legal research, and case management. This can reduce the workload and costs for lawyers and courts, and speed up the resolution of legal disputes.”<sup>7</sup> To the extent that machine learning can aid in performing (or outright perform) these tasks at a small fraction of the time a human would expend, AI has extraordinary potential for saving lawyers and others

---

<sup>4</sup> *What is Bing AI Search?*, <https://searchwith.ai/blog/understanding-how-bing-ai-works/#:~:text=Bing%20AI%20search%20is%20a%20search%20engine%20that,technology%20called%20Prometheus%20to%20enhance%20the%20search%20experience> (last visited Oct. 9, 2023).

<sup>5</sup> The verbatim response appears in the Appendix.

<sup>6</sup> The verbatim response appears in the Appendix.

<sup>7</sup> All bolded text in Bing AI responses reflects emphasis in the original.

resources that could be directed to more complex needs. Some extremely capable minds have declared that some of these tasks will be “resolved/solved in the near term” (if not currently) by generative AI platforms.<sup>8</sup> But is Bing AI aiming too low in identifying routine lawyer tasks? Why stop there?

A more ambitious agenda might take an even more prophylactic approach, aiding with dispute resolution before the lawyer begins their work, before a complaint reaches the court clerk’s window, or even before parties arrive at the courthouse. For example, some courts are turning to court-adjacent online dispute resolution (ODR) for high-volume civil disputes (e.g., consumer debt).<sup>9</sup> To our knowledge, the technology currently adopted in ODR platforms all requires some human facilitator to help litigants reach pretrial settlement. They usually are available in chat spaces or by individual email messages to help the parties reach a mutually acceptable plan. What if generative AI could help better facilitate that process in real time? Not only would courts save on human labor costs, but a well-designed algorithm should also be able to narrow the settlement space more accurately and quickly than even the most seasoned mediator.

As we up the ante on ambition, we might also expect generative AI to guide parties on how best to *avoid* litigation altogether? Might a sophisticated algorithm sift through and diagnose difficult issues before suits are filed and positions begin to calcify? How about helping parties arrive at a genuinely understandable and objectively fair residential lease agreement that precludes the need for many summary eviction cases? Could AI help educate a self-represented party on how to solve a problem they face without having to identify it as a legal issue? In addition to increasing the rate of dispute resolution, AI could assist with procedural engagement along the way. For example, generative AI might have the capacity to accurately translate materials and proceedings for non-English speakers in ways that promote procedural fairness alongside happiness with substantive outcomes.

It is a good and helpful thing for lawyers and judges to have more accurate information, to be more efficient, to have better tools for assessing risk, and to deliver more actionable advice. But for many parties to life-altering litigation, the possibility of having a lawyer to perform all those functions is unlikely, if not impossible. Common examples include unemployment benefits claimants, tenants in eviction suits, people experiencing consumer or medical debt, and family members arguing over custody or support arrangements.<sup>10</sup> One can easily imagine the impact of

---

<sup>8</sup> Katherine B. Forrest & Catherine Nyarady, Paul, Weiss, Rifkind, Wharten & Garrison, LLP, *AI and Access to Justice*, PowerPoint presentation at slide 31 (April 19, 2023), <https://www.nycourts.gov/LegacyPDFS/accesstojusticecommission/tc/2023/3A-AI-and-Algorithmic-Bias.pdf> (last visited Oct. 9, 2023).

<sup>9</sup> For one prominent example, see Utah Courts, *Online Dispute Resolution (ODR) Pilot Project*, <https://legacy.utcourts.gov/odr/> (last visited Nov. 29, 2023).

<sup>10</sup> See Samuel A. Thumma & Jaqueline E. Marzocca, *The Self Represented Party The Most Unique Party of Them All*, 59 ARIZONA ATTORNEY 24, 26 (June 2023) (“Nationwide, estimates provide that ‘more than 70 percent of civil and family cases involve at least one self-represented party.’ . . . In Arizona, the percentages may be even higher. For Maricopa County Superior Court cases closed during the 12 months ending June 30, 2021 (FY 2021), more than 90 percent of family court cases had at least one self-represented party, and more than 70 percent of the cases involved both parties being self-represented.”) (citations omitted).

generative AI most clearly, not in saving lawyers more time on their case, but by providing comprehensible information to self-represented litigants that they otherwise would never receive. When representing clients, a lawyer usually provides strategic advice and counsel and suggests which issues should be litigated fiercely (and which should not). Perhaps generative AI could perform the same functions for people who are not represented by lawyers and have little or no chance of retaining counsel.

Relatedly, and quite powerfully, generative AI might help courts and academics understand why self-represented parties eschew technology (e.g., electronic document filing) that attorneys are required to use.<sup>11</sup> Can generative AI shed light on why self-represented parties are obtaining childcare, taking time off work, finding transportation, walking through courthouse halls, and filing hard-copy documents when, instead, they could handle pleadings over the internet from the comfort of their own home? Can generative AI help identify the best days of the week, and best times of the day, to help ensure parties appear for court hearings?<sup>12</sup> The answers to these questions will, at least in part, reveal whether generative AI can improve the legal system (in court-based litigation, court-adjacent efforts, or completely outside the court system) for those who otherwise get lost in the legal shuffle.

Second, Bing AI responded: “AI can **democratize access to legal information** by providing online platforms and tools that can answer legal questions, generate legal documents, and offer guidance and advice. This can help people who cannot afford or access lawyers to solve their own legal problems or connect them with licensed professionals who can.” These are formidable examples of how generative AI can promote access to justice by delivering the law on demand to people’s digital devices. Some wise commentators have advocated such advancements.<sup>13</sup> But, again, is Bing AI “thinking” too narrowly?

For all the good that would follow more people interfacing with legal rules, their text and structure are often unnecessarily complex. What if generative AI could identify the most problematic bottlenecks in legal processes and simplify them? For example, if self-represented plaintiffs routinely find their claims dismissed, we might look first to service of process rules. Generative AI could sort among the many reasons why these plaintiffs fail to serve: the limited methods available, the time limits, something else? AI might be able to detect macro-level

---

<sup>11</sup> Data on file with the authors show that, from July 1, 2022 to May 31, 2023, of the 446,154 filings by self-represented litigants in family court cases in Maricopa County Superior Court, 432,797 (or 97 percent) were paper filings, while only 13,175 (or 3 percent) were documents that were efiled (which is required by lawyers representing litigants).

<sup>12</sup> Data on file with the authors show that, from July 12, 2023 through November 30, 2023, of the 5,667 total initial eviction hearings in the Pima County Consolidated Justice Court, the average appearance rate was 62%, but that the appearance rate on Mondays was 60%, while the appearance rate on Tuesdays was 63%.

<sup>13</sup> Clare Fraser, *AI: Opening the Door to Justice: How We Can Enhance Access to Justice – and Prevent Inequality – by Developing a Customised Artificial Intelligence Model with the Citizen as the End User*, Law Society of Scotland (August 2023) (advocating for creating “a customised large language model (“LLM”) within an environment where data such as case law, codes of practice and guidance [for the laws of Scotland] have been uploaded and embedded. The LLM is developed with the citizen as the predominant user and not the legal professional.”), <https://www.lawsco.org.uk/members/journal/issues/vol-68-issue-08/ai-opening-the-door-to-justice/> (last visited Oct. 9, 2023).

patterns that elude even the most intelligent lawyers and spur reforms that make rules more user-friendly for self-represented litigants—and even for seasoned lawyers. Procedural rule sets never (or almost never) get shorter and simpler over time. Maybe AI could presage a reversal of that trend and streamline rules so that *everyone* can understand them. Instead of a legal system dominated by centuries-old, arcane, and at times foreign language—or at least a seemingly different dialect of American English—technology might generate rules that facilitate rather than frustrate their use.

Any of these functions inherently makes the law more inclusive and applicable. AI likely won't replace the three years of formal training that licensed attorneys need to practice, but it can help them express legal arguments more plainly and, hopefully, more effectively.

Third, Bing AI Responded: “AI can **improve the quality and consistency of legal decisions** by using data and algorithms to analyze cases, predict outcomes, and recommend actions. This can help judges and lawyers to make more informed and objective decisions, and reduce the risk of human errors and biases.” The platform added: “AI can **enhance the transparency and accountability of the justice system** by making legal data and processes more accessible and understandable to the public. This can increase the trust and confidence of the people in the rule of law, and encourage participation and feedback.”

If true, this prediction would be monumental. As a system administered by human beings, the justice system is not free from bias or discrimination. And generative AI provides the possibility of offering great advances in reducing those flaws. But we ask again: is Bing AI not ambitious enough?

Just as it might lend a hand in rationalizing an overly complex legal system, generative AI might help the system determine which justice indicators are valid and consistent and which are not. Generative AI could supply a macro-level vision for curing longstanding problems, using enormous data sets to identify and help set better standards. For example, many courts and social scientists suggest that lowering the recidivism rate indicates success in a specialized treatment court. What if the premises underlying that choice are wrong? For example, the recidivism rate might not be well-defined. Does it only encompass repeat offenses during the treatment period? After two or three years following treatment? Recidivism could be limited to drug offenses or expand to include violent crimes. Generative AI might help select the best indicator of success, conditional on a jurisdiction's resources, values, and objectives. Similarly, is affording judicial officers more, or less, discretion when deciding pretrial release or sentence conditions the better path? Or are there several avenues, not just one, that generative AI can help identify?

On this dimension, the future is decidedly uncertain. As a result, generative AI's influence will extend only as far as stakeholders accept its results as valid. Consider the criminal justice system again. Law enforcement officers, prosecutors, defenders, courts, prisons, and boards of parole and clemency all come to the system with their professional experience and conventional wisdom. For generative AI to break through the inertia, it has to prove its own efficacy by teaching human users how to look at the world correctly, or at least better.

Generative AI provides a tool, but not a panacea, for addressing time-worn, intractable issues with new and perhaps counterintuitive suggested resolutions. It's time to look hard and deeply at those potential solutions that generative AI makes possible. But in doing so, it is essential to address best practices and recognize concerns generative AI presents, and also have a careful eye focused on how to measure success.

### **III. Concerns over Using AI and Best Practices**

To leverage AI toward access to justice, we must understand its limitations and cultivate best practices toward empowering users instead of augmenting inequities.

Machine learning model outputs are no more than information predictions and collections. We describe these models as “learning” things because they undergo a process designed to mirror the way humans absorb information. AI algorithms are initially “educated” on a set of training data, mapping patterns in those data until they can receive new information and generate accurate connections or identify valid patterns. For example, if we are training an algorithm to perform facial recognition tasks, we might feed it a series of images of people's faces (as well as pictures of other items). The more faces it “sees,” the better it can identify what factors are most important to correctly picking out faces “in the crowd.”

When they work, these systems are truly impressive. Understanding their limitations in any particular context (e.g., aiding criminal investigations) is critical to mitigating the risks of incorrect prediction and ensuring due process in their implementation. In the access to justice realm, inaccurate predictions could be devastating. If self-represented litigants rely on generative AI to navigate civil legal issues, incorrect guidance on answering a lawsuit could lead to a default judgment. In generative text models, like OpenAI's GPT-4, the answer to a question or prompt is also a prediction: the most likely next word or phrase based on a large language model. As with facial recognition technology, the accuracy and usability of an AI response to a question about handling an eviction case will depend on the quality, scale, and variability of the data on which the algorithm was trained, as well as the structure of the prompt itself.

How instructive any response is to the prompt a person enters will, at first, depend on the data used to train the model. Although it might seem obvious, an inanimate algorithm cannot (at least not yet) learn from information to which it was not exposed. This truth leads to a shortcoming of generative AI known as *exposure bias*. Exposure bias emerges when a computer model trained on a specific set of data 1) does not perform well when introduced to different data and 2) fails to creatively and accurately interpret the new data. This problem literally feeds on itself. In the context of generative AI, exposure bias grows rapidly and perniciously. Any prediction of the next word in a generative response relies on *all the previous words* the model generated. As a result, any error early on compounds the “garbage in-garbage out” dilemma.

A recent cautionary example comes from the “Tessa” generative chatbot used by the National Eating Disorder Association (NEDA). The chatbot was designed to replace humans at a call center for people dealing with disordered eating. Because generative AI models must be trained on a wide cross-section of data to provide sufficient responses, those training data needed to include enough examples of helpful reactions to someone in distress. Unfortunately, *dieting*

*advice* is much more prevalent in everyday human conversation; the training data, therefore, “taught” the algorithm to use language more consistent with restricted eating. As a result, by early June 2023, NEDA had suspended its AI chatbot for giving harmful advice. NEDA’s chief executive, Elizabeth Thompson told the *New York Times* she was “waiting for an explanation about how that content was introduced into a closed program.”<sup>14</sup>

The lesson for access to justice advocates is that AI tools must be trained on data reflecting the legal problems facing people across socioeconomic, educational, and geographic distributions—not just the average or endpoints of the distribution. This is particularly true for racial and ethnic minorities; their experiences might not be recognized by the algorithm because of training data limitations, leading to serious errors in advice or decision-making.<sup>15</sup> At the very least, AI tools designed for self-represented litigants should oversample the cases and circumstances that those individuals most frequently encounter. Otherwise, they might be worse off than without the technology, as in the NEDA example above.

Generative AI is also subject to *hallucinations*, which are inaccurate sentences or phrases produced by the system. While there are methods to reduce such risk, no technique exists to completely eliminate it. Generative AI models create content by predicting the next object in some sequence that is probabilistically ideal, be it a word or a pixel. More advanced generative models produce much better prediction outputs but might be cost-prohibitive for adoption in access to justice spaces.

In addition to fictitious sentences, well-known hallucinations include generating false citations. False citations arise when the algorithmic model is designed to predict the right combination of words and numbers that mirrors the structure of citations from training data, *without regard to the truth*. A now-infamous example involved two plaintiff’s attorneys who used ChatGPT to write a legal brief. The AI platform hallucinated six case citations in the document, which defense counsel could not locate in actual reporters.<sup>16</sup> The court ended up dismissing the case and sanctioning the plaintiff’s attorneys. In the sanctions order, the court said “that there is nothing inherently improper in lawyers using AI ‘for assistance,’ but he said lawyer ethics rules ‘impose a gatekeeping role on attorneys to ensure the accuracy of their filings.’”<sup>17</sup>

---

<sup>14</sup> Lauren McCarthy, *A Wellness Chatbot Is Offline After Its ‘Harmful’ Focus on Weight Loss*, N.Y. TIMES, <https://www.nytimes.com/2023/06/08/us/ai-chatbot-tessa-eating-disorders-association.html> (June 8, 2023; updated June 9, 2023).

<sup>15</sup> See, e.g., Brianna Rauenzahn, Jamison Chung & Aaron Kaufman, *Facing Bias in Facial Recognition Technology*, THE REG. REV., Mar. 20, 2021, <https://www.theregreview.org/2021/03/20/saturday-seminar-facing-bias-in-facial-recognition-technology/>.

<sup>16</sup> See Larry Neumeister, *Lawyers Submitted Bogus Case Law Created by ChatGPT. A Judge Fined Them \$5,000*, AP (June 22, 2023), <https://apnews.com/article/artificial-intelligence-chatgpt-fake-case-lawyers-d6ae9fa79d0542db9e1455397aef381c>.

<sup>17</sup> Sara Merken, *New York Lawyers Sanctioned for Using Fake ChatGPT Cases in Legal Brief*, REUTERS (June 26, 2023), <https://www.reuters.com/legal/new-york-lawyers-sanctioned-using-fake-chatgpt-cases-legal-brief-2023-06-22/>.

Even worse, some models have been trained to produce real citations but still apply them incorrectly or look to true citations that are not the best choice for the proposition stated. For example, when we asked a prototype legal chatbot “Can a school prevent a student article from being printed in a school publication?”, it responded, in part: “[S]chool authorities can exercise prior restraint on publications distributed on school premises during school hours if they can reasonably forecast substantial disruption of or material interference with school activities due to the distribution of such printed material USCS Const. Amend. 1, Religious and political freedom.” Although the answer does follow from a necessary First Amendment analysis, the better source for citation purposes is the actual United States Supreme Court decision.<sup>18</sup> Now imagine a self-represented litigant using a chatbot to draft a pleading or other court document. Without the first clue about how to verify a citation’s accuracy, the litigant could wind up submitting subpar—perhaps completely fabricated information—and drawing the court’s ire.

Another concern for AI-informed access to justice is the transparency of algorithmic processes. Due process is founded on notice, the opportunity to be heard at a meaningful time in a meaningful way, and the chance to challenge evidence offered against a party.<sup>19</sup> Many AI systems are not capable of providing the reasoning behind their outputs. Without a clear understanding of the factors involved in systems, the bases for a decision, or the ability to challenge it after the fact, due process is imperiled.<sup>20</sup>

Developing best practices for legal AI systems are essential and should embrace, among other guidelines, the following:

- **Use Diverse, Representative Data.** Bias in AI outputs often stems from biased training data.<sup>21</sup> Ensuring that training datasets reflect diversity across the many dimensions that matter for access to justice (e.g., race, ethnicity, income, education) is crucial. Without a wide range of demographics, perspectives, and scenarios in the data, any AI tool will surely underserve its intended user base. When representative data are not available, data scientists can apply technical strategies for reducing bias or improving data collection methods for future analysis. Similarly, testing prototypes with the populations who often represent themselves in court can inform the development process and help identify potential areas of bias. Too often, innovation occurs without the input of the intended

---

<sup>18</sup> See *Hazelwood Sch. Dist. v. Kuhlmeier*, 484 U.S. 260 (1988); see also Mehul Bhattacharyya et al., *High Rates of Fabricated and Inaccurate References in ChatGPT-Generated Medical Content*, 15 CUREUS, May 19, 2023, at 2, 4. (showing that, when ChatGPT-3.5 was asked “to generate 30 unique short papers on various biomedical topics,” the platform listed 115 total references, of which “47% were fabricated, 46% were authentic but inaccurate, and only 7% were authentic and accurate”).

<sup>19</sup> See, e.g., *Armstrong v. Manzo*, 380 U.S. 545, 552 (1965); *Grannis v. Ordean*, 234 U.S. 385, 394 (1914).

<sup>20</sup> *But see State v. Loomis*, 881 N.W.2d 749, 761 (Wisc. 2016) (finding that “[the defendant] is correct that the risk scores do not explain how the COMPAS program uses information to calculate the risk scores” but declining to find a due process violation because “to the extent that [the defendant’s] risk assessment is based upon his answers to questions and publicly available data about his criminal history, [the defendant] had the opportunity to verify that the questions and answers listed on the COMPAS report were accurate”).

<sup>21</sup> See, e.g., Sandra G. Mayson, *Bias In, Bias Out*, 128 YALE L.J. 2218 (2019).

user community. In the access to justice context, that means testing academic and practitioner assumptions against the lived experience and needs of the target audience.<sup>22</sup>

- **Create “Human-in-the-Loop” Systems.** Human oversight in AI decision-making processes must be included in any algorithmic platform, especially in a domain like high-volume, high-stakes civil litigation. Keeping people “in the loop” will not guarantee success, but timely human intervention can override decisions that the AI system does not “understand” will be detrimental to users. The level of human oversight needed as well as the timing of oversight depends on the level of risk involved and the potential implications of delay.
- **Develop Impact Assessments.** AI models that courts and lawyers deploy should be reviewed regularly to ensure that the outcomes they expect align with the outcomes they observe, to the extent possible. When they do not, developers should refine the model to account for those unexpected outputs and to incorporate new data and changing societal norms, both of which reduce bias over time. For example, an impact assessment might have flagged the issues with NEDA’s Tessa chatbot before it caused any harm, by identifying how poorly the system responded to people with eating disorders.
- **Be as Transparent as Possible.** Stakeholders should strive for transparent explanations of how any AI model was developed, and how it works, so that users can see what factors informed the decision-making process and how they were weighted. Being open about algorithmic inputs and calculations builds trust and understanding among everyone involved in the civil justice system. Educating users about the capabilities and limitations of AI models, as well as providing clear guidelines on how to use these models effectively and responsibly can help mitigate risks.

#### IV. How to Measure Success

As the justice system grapples with the questions and caveats in the previous two sections, it should simultaneously deploy a suite of evaluation tools for measuring generative AI’s benefits. Legal academics and social scientists now have at their disposal a variety of methodologies for program evaluation.<sup>23</sup> A complete review of those methods is beyond the scope of this essay. For now, we highlight some key criteria in choosing how to evaluate AI systems in the access to justice context.

The first consideration is *what* to evaluate. This question is one of outcomes. If we want to know whether an AI tool promotes inclusivity and transparency, we might focus on user

---

<sup>22</sup> A recent example is STACY BUTLER, SARAH MAUET, CHRISTOPHER L. GRIFFIN, JR. & MACKENZIE PISH, THE UTAH ONLINE DISPUTE RESOLUTION PLATFORM: A USABILITY EVALUATION AND REPORT (2020), available at [https://papers.ssrn.com/sol3/papers.cfm?abstract\\_id=3696105](https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3696105).

<sup>23</sup> A very accessible overview appears at Stanford Legal Design Lab, *Evaluation Methods of Justice Innovations*, <https://justiceinnovation.law.stanford.edu/resources/evaluation/> (last visited Nov. 30, 2023).

comprehension of how the tool works. If we want to know more about whether self-represented litigants successfully resolve their legal matters, we will choose “win rates” as the relevant outcome variable. And if we want to understand better how AI promotes efficient dispute resolution, we might use time to disposition as the main indicator.

In some sense, there is no “right” choice when it comes to outcome variables. What matters to the empirical analysis is what matters to the community deploying the AI tool. Thus, measuring success is somewhat in the eye of the beholder. The outcome variables included in any evaluation should reflect the values and needs of those administering algorithmic systems. For example, a jurisdiction that wants its online dispute resolution tool to be useful without relying too much on human technical support would care a lot about whether users can find answers in the frequently asked questions (FAQ) section.<sup>24</sup> But constantly turning to the FAQ can also signal that the platform is non-intuitive or too cumbersome to follow.

The second consideration that courts and administrators should confront is *how* to evaluate. Again, there are many more evaluative methods from applied statistics than space to review here. Suffice it to say that there are three primary approaches: (1) subjective surveys; (2) observational data; and (3) experimental methods. Surveys, by construction, can only reveal (if anything) how and why users interact (or don’t) with a legal innovation. They can be informative about efficacy—insofar as user satisfaction measures how well something works—but they are usually unable to say much objectively about outcomes. Observational studies rely on large datasets including measures of the chosen outcome (the dependent variable) and all the discernable factors that could plausibly impact it (the independent variables). The social scientist using observational methods often wants to generate causal inference. They can’t, however, because the processes that created the data are subject to selection effects and other “confounding” influences. For example, consider a court that deploys an AI dispute resolution platform that, when used, resolves cases more quickly than the status quo ante. That result *could* reflect the utility of the AI tool. It could also pick up the unobserved impact of inherent diligence if the people who choose to use the tool get things done more quickly (on average) than those who choose not to use the tool. At the extreme, the tool could be useless, and the results only emerged because users are faster in general than non-users in taking care of business.

The gold standard methodology for assessing whether any legal innovation, AI-based or otherwise, is the randomized control trial (RCT). In short, RCTs follow the procedure of a clinical trial: participants are divided into a control group and one or more treated groups using some randomizing device (e.g., coin flip, wheel spin). The treated groups are exposed to the innovation, and the control group is shielded from the innovation as much as possible. Allocating the new tool or resource ensures (on average) that any selection effects or confounds will wash out in the analysis. Experimentation along these lines is anathema to many jurists and lawyers. As such, the legal profession lags behind others in the evidence basis for its practices.<sup>25</sup> But if we really want to learn what works and what doesn’t—if we want to uncover causality in legal process—we should embrace experimental methods more readily.

---

<sup>24</sup> See BUTLER, ET AL., *supra* note 2022, at 35-39.

<sup>25</sup> See, e.g., D. James Greiner & Andrea Matthews, *Randomized Control Trials in the United States Legal Profession*, 12 ANN. REV. L. & SOC. SCI. 295 (2016).

Finally, stakeholders must ask *when* to evaluate AI tools for access to justice. This question might be the most practically important for courts and litigants. If an AI-related innovation enters public spaces before being subjected to rigorous testing, any of the adverse consequences outlined in the previous section could accrue. If so, the stakeholders involved would have to admit that they deployed a new procedure without fully understanding (or understanding at all) its likely effects. Even with relatively benign interventions like self-help materials in courthouses, failure to evaluate beforehand risks all sorts of unintended consequences. Thus, justice system stakeholders should at all costs avoid launching AI platforms at scale without findings from a proper evaluation in hand.

The double gold standard path forward, as it were, would be to pilot an AI-backed tool with a small, but statistically powerful, number of users in an RCT. Doing so both provides preliminary evidence of whether the tool works and helps developers weed out bugs. One court in a state, or one courtroom in a county, could be the pilot jurisdiction. Armed with solid evidence of effectiveness, administrators could refine the effort and scale up the pilot to more locations and repeat the evaluation. Repeated findings that the AI platform works offers a proper evidentiary basis for full deployment. Aside from this admittedly time-consuming process, courts and lawyers should at least pursue rigorous evaluation no later than the tool's launch date. Along with the downsides mentioned above, officers of the law are reluctant to abandon practices they *believe* are useful. The more entrenched an innovative practice becomes over time, the harder it may be to discard—even if later evaluation shows that it is (and perhaps never was) useful.

## **V. Conclusion**

Generative AI is opening doors to rooms that, until very recently, we didn't even know existed and could not imagine. Its capacity for processing all the information in the country's law libraries and more has enormous potential for enhancing access to justice. The most commonly used chatbot today provides decent answers to the question that we set out to answer in this essay. But those answers are incomplete. The AI platform is less ambitious, does not appreciate the significant risks for self-represented litigants, and cannot self-assess its responses. These truths reinforce the great care needed when using generative AI to enhance access to justice, to ensure its long-term success, and to address a host of valid concerns. At the risk of hyperbole, in the future, the sky is the limit—provided we understand generative AI's promise and pitfalls now.

## Appendix

Bing AI Search Engine (<https://www.bing.com>) responses to the following query on different dates: “How can artificial intelligence help advance access to justice.”

**August 14, 2023 at 3:47 pm (Arizona MST) search = 220 words and two sources:**

Artificial intelligence (AI) can help improve access to justice in many ways. According to the Centre for Public Impact (CPI), AI can provide a more responsive justice system by augmenting and even replacing lawyers. Intelligent machines can help provide a more just legal outcome than a human<sup>1</sup>. In many areas of human activity, technological advances have made goods and services cheaper and more accessible. Thus far, however, in law (as in medicine) technological advances have tended to have the opposite effect. A major contributor to this fact is the law’s apparent impenetrability. The law is perceived as a complicated organism that requires the costly intervention of legal experts to be accurately deciphered. Accordingly, even as our technology gets smarter and faster, the trend in the law has been for legal services to get bigger and more costly. But will this always be the case? Intelligent machines now increasingly shape our society. We routinely rely on algorithms to perform tasks, often with greater precision and reliability than humans<sup>1</sup>.

AI can provide a legal service at a significantly reduced unit cost to a much larger portion of the population. Firstly, it can support the provision of legal services; and secondly, it can replace the role of legal experts. Legal technology that supports justice includes natural language processing (NLP), machine learning and chatbots<sup>2</sup>.

1. <https://www.centreforpublicimpact.org/insights/joel-tito-ai-justice>
2. <https://www.abajournal.com/magazine/article/artificial-intelligence-has-made-great-inroadsbut-not-as-far-as-increasing-access-to-civil-justice>

**November 4, 2023 at 7:44 am (Arizona MST) search = 587 words and eight sources:**

Artificial intelligence (AI) is a powerful technology that can help advance access to justice in many ways. Some of the possible benefits of AI for the justice system are:

- AI can **increase efficiencies** by automating tasks such as document preparation, legal research, and case management. This can reduce the workload and costs for lawyers and courts, and speed up the resolution of legal disputes.
- AI can **democratize access to legal information** by providing online platforms and tools that can answer legal questions, generate legal documents, and offer guidance and advice. This can help people who cannot afford or access lawyers to solve their own legal problems or connect them with licensed professionals who can.

- AI can **improve the quality and consistency of legal decisions** by using data and algorithms to analyze cases, predict outcomes, and recommend actions. This can help judges and lawyers to make more informed and objective decisions, and reduce the risk of human errors and biases.
- AI can **enhance the transparency and accountability of the justice system** by making legal data and processes more accessible and understandable to the public. This can increase the trust and confidence of the people in the rule of law, and encourage participation and feedback.

However, AI also poses some challenges and risks for the justice system, such as:

- AI can **create ethical and legal dilemmas** by raising questions about the responsibility, liability, and accountability of the developers and users of AI systems. For example, who should be held accountable if an AI system makes a wrong or harmful decision, and how can the affected parties seek redress?
- AI can **exacerbate existing inequalities and injustices** by reflecting and reinforcing the biases and discrimination that exist in the data and algorithms used by AI systems. For example, an AI system that predicts the risk of recidivism for criminal defendants may be influenced by factors such as race, gender, and socioeconomic status, and result in unfair and discriminatory outcomes.
- AI can **undermine the human dignity and autonomy of the people** by reducing their involvement and control over their own legal affairs. For example, an AI system that provides legal advice may not be able to capture the nuances and emotions of the human situation, and may not respect the preferences and values of the people.

Therefore, it is important to ensure that AI is used in a way that respects the principles and values of the justice system, such as fairness, equality, impartiality, and human rights. Some of the possible measures to achieve this are:

- Developing and implementing **ethical and legal standards and guidelines** for the design, development, and use of AI systems in the justice system. These should ensure that AI systems are transparent, explainable, accountable, and respectful of human dignity and autonomy.
- Establishing and enforcing **oversight and regulation mechanisms** for the monitoring, evaluation, and auditing of AI systems in the justice system. These should ensure that AI systems are reliable, accurate, consistent, and compliant with the ethical and legal standards and guidelines.
- Providing and promoting **education and training** for the developers, users, and beneficiaries of AI systems in the justice system. These should ensure that they have the necessary skills, knowledge, and awareness to use AI systems effectively, responsibly, and ethically.

*Please do not cite or distribute without permission*

AI has the potential to transform the justice system and improve access to justice for many people. However, it also requires careful and thoughtful consideration of its implications and impacts, and the adoption of appropriate measures to ensure that it is used for good and not for evil.

Learn more:

1. [papers.ssrn.com](https://papers.ssrn.com)
2. [bing.com](https://bing.com)
3. [jolt.law.harvard.edu](https://jolt.law.harvard.edu)
4. [centreforpublicimpact.org](https://centreforpublicimpact.org)
5. [lawtimesnews.com](https://lawtimesnews.com)
6. [yjolt.org](https://yjolt.org)
7. [capgemini.com](https://capgemini.com)
8. [papers.ssrn.com](https://papers.ssrn.com)