

Who Can Write a Better Brief: Chat AI or a Recent Law School Graduate?

Part 1

BY RONALD M. SANDGRUND

“Lawyers’ jobs are a lot less safe than we think.”¹

“Law is seen as the lucrative profession perhaps most at risk from the recent advances in A.I. because lawyers are essentially word merchants.”²

“No, lawyers won’t be replaced by artificial intelligence. Yet. Give it a few years.”³

“The notoriously change-averse legal industry will face a particularly abrupt disruption by AI.”⁴

Law firms “fail to appreciate how quickly the pace of exponential change can be.”⁵

“Firms too slow to adapt to AI . . . will suffer a competitive disadvantage.”⁶

“AI will replace lawyers . . . who fail to adapt with it.”⁷

“It may even be considered legal malpractice not to use AI one day.”⁸

“The sky is falling! The sky is falling!”⁹

This is the tenth article series by The Inquiring Lawyer addressing a topic that Colorado lawyers may discuss privately but rarely talk about publicly. The topics in this column are explored through dialogues with lawyers, judges, law professors, law students, and law school deans, as well as entrepreneurs, computer scientists, programmers, journalists, business leaders, politicians, economists, sociologists, mental health professionals, academics, children, gadflies, and know-it-alls (myself included). If you have an idea for a future column, I hope you will share it with me via email at rms.sandgrund@gmail.com.

This two-part article examines whether lawyers will soon be replaced by machines and, more important, whether the Inquiring Lawyer’s days as a columnist are numbered. Part 1 consists of an interview with Professor Harry Surden, a nationally known law professor, former software engineer, and expert on the intersection between artificial intelligence (AI) and legal practice. Also weighing in is ChatGPT-3.5, an artificial language program. Part 2 will feature The Inquiring Lawyer’s version of a battle rap, giving readers the opportunity to compare the wit and wisdom of The Inquiring Lawyer and ChatGPT as expressed in their parallel humorous essays about lawyers.

Introduction

Did the quotes at the start of this article get your attention? Did they strike you as tech hype? Fear-mongering? Just clickbait for lawyers?

I love science fiction books and movies about the coming *l’apocalypse de la machine*. I feasted on Isaac Asimov’s *Robot* series, with its three laws of robotics¹⁰—instructions built into robots so they don’t harm humans—and its chief protagonist R. Daneel Olivaw, a humanoid detective who helps solve murders involving apparent violations of the three laws. Two of my favorite movies are *Blade Runner*, based on Phillip Dick’s *Do Androids Dream of Electric Sheep?*, and Ridley Scott’s *Alien*. In both movies, androids create more than a few problems for their creators. And I loved James Cameron’s *The Terminator*, involving an existential war

across time between humans and their creation, Skynet,¹¹ and Skynet's cyborgs. Standing as a beacon of hope are Data, from *Star Trek the Next Generation*, and R2-D2, C-3PO, and BB-8, from *Star Wars*, who serve faithfully alongside their human creators. Not so much HAL.¹²

But I digress.

This dialogue may seem a bit pedestrian in the shadow of these monumental science fiction works, but it concerns an issue that should be creeping onto every lawyer's and law firm's radar screens: the encroachment by—or maybe, more hopefully, a collaboration with—AI. Word processing, e-discovery, and searchable legal databases were all adopted during my legal career, and each had profound effects on the day-to-day practice of law, legal ethics, the business of law, and the attorney job market. Many of us recall the gross inefficiencies of practicing law in the 20th century: (1) typing (and retyping) briefs and contracts on paper using a typewriter; (2) employing Wite-Out®; (3) printing, copying, and snail-mailing legal briefs to opposing counsel and the court; (4) tunneling through boxes of court-stored paper files; (5) hiring persons called “legal secretaries” to type one's handwritten notes and dictation onto paper; (6) sending letters to opposing counsel using something called the “US Post Office,” and wondering if they ever arrived and whether, in a week or two, you might get a response; (7) driving to a law library to conduct legal research, including wading through volumes and volumes of Shepard's Citations to see if that fantastic case you are relying on has been overruled; and (8) spending weeks arranging your client's dusty and creased business records in chronological order and then reading them line by line to see if there was anything relevant or privileged in there.

For newer lawyers snickering at these examples, I ask: Are you ready for the day when an AI program could write a brief or a contract that is far better than anything you could produce? What if you can't afford to purchase the AI program? And what about your kids: will you be encouraging them to go to law school if it looks like AI will be performing over 50% of the work lawyers currently perform?

All of which raises the question whether we are at an inflection point, like when seemingly

overnight tens of thousands of horses were put out to pasture following the arrival of the mass-produced automobile in 1910, or when thousands of elevator operators looked for new jobs after the widespread acceptance of automatic elevators in 1950, or when most travel agents went extinct in the early 2000s. Hello Expedia, Kayak, and Booking.com!

There are dozens and dozens of practical, legal, ethical, moral, and business issues tied up in AI performing legal and judicial tasks, from writing contracts, to interviewing potential clients online or virtually via holograms (did you know that research shows that clients are often more honest talking to a robot than a human?¹³), to predicting the settlement value of a personal injury case, to determining appropriate bail and jail sentences untainted by cognitive and structural biases,¹⁴ to providing access to justice to hundreds of thousands of folks who cannot find or afford a lawyer willing to help them, to—well, the list is quite long.

To keep things simple, this dialogue will focus mainly on a *singular* legal task that large language models using AI may soon perform as well or better lawyers: writing a motion and brief addressing discrete legal issues. ChatGPT's utility in transactional work will not to be addressed here. (I heard from one reliable source that a Big Law partner reviewed a first draft of a merger agreement created by ChatGPT and reported it was as good or better than any first draft he had seen.)

Participants



ChatGPT is a computer program. I interviewed version GPT-3.5. Version GPT-4 is now available as a subscription service.



The Inquiring Lawyer is a human being with an opinion on everything.



Professor Harry Surden is a human being. He is also a professor at the University of Colorado Law School. He joined the faculty in 2008. His scholarship focuses on

legal informatics, AI and law (including machine learning and law), legal automation, and issues concerning self-driving/autonomous vehicles. He also studies intellectual property law with a substantive focus on patents and copyright, and information privacy law. Before joining CU, Professor Surden was a resident fellow at the Stanford Center for Legal Informatics (CodeX) at Stanford Law School. In that capacity, Professor Surden conducted interdisciplinary research with collaborators from the Stanford School of Engineering exploring the application of computer technology toward improving the legal system. Before attending law school, Professor Surden worked as a software engineer for Cisco Systems and Bloomberg L.P.

A Glimpse Into the Future



The Inquiring Lawyer: Professor Surden, could you tell us why you transitioned from software engineer to attorney and law professor?



Professor Harry Surden: I was always interested in both the technical and social science and humanities side of topics. As an undergraduate, I took


a broad range of courses, studying computer science as well as political science, philosophy, and even an undergraduate law course. When I graduated and entered the world of software engineering, working first in finance at Bloomberg L.P. and then at Cisco Systems, I found it fascinating working with and programming these vast and complex computer systems. Particularly at Bloomberg, as a software engineer, I saw how technology was transforming finance in the late 1990s. One thought in the back of my mind was the idea that a similar transformation could somehow impact law as well, perhaps empowering the public and, hopefully, bettering society. After several years as a software engineer, I decided to pursue the other, nontechnical side of my interests, by pursuing a law degree at Stanford. My hope was that I could eventually become a law professor and combine my two interests, studying artificial intelligence and law. There, with the support of several professors, I helped cofound CodeX in 2005. Since then, my excellent colleagues have been pursuing this goal.

InQ: Why does AI interest you as a law professor?

Prof. Surden: There are a few reasons. One is the idea that we may be able to use the technology to help those who are underserved by lawyers. By some accounts, 80% of people in the United States who need legal assistance are unable to obtain or afford it. One approach would be to fully fund legal help for all those who need it. But it has been 50-plus years, and our country does not seem to want to do that. Another idea in the back of my mind was that perhaps AI or other similar technologies could help bridge this access to justice gap. In a similar vein, I always thought that law made itself much more difficult to understand, to the non-legally educated person, than it needed to be. In my opinion, law often unjustifiably cloaks itself in jargon, obscuring certain ideas from being understandable. While some areas of law are justifiably complex, others are—perhaps through no one's intention or fault—actually quite simple underneath but difficult for a typical person to understand. The similar hope was to make law more understandable to the lay public, perhaps by using technology, so that we are all more aware and capable of engaging with the laws that govern us. Finally, as a software engineer, I find the whole topic of AI completely fascinating. The idea that one can, however imperfectly, encode linguistic meaning using math, continues to amaze me to this day.

InQ: Since many readers are not familiar with either the technology or the jargon that fills the AI space, perhaps you could define—in words a lawyer born in the 1960s would understand—various terms, and maybe give a simple example of each, starting with the term “algorithm.”

Prof. Surden: An algorithm is a series of well-defined instructions designed to perform a specific task, such as sorting a list of numbers from largest to smallest. To use an analogy, it is a little like a recipe in cooking, where each step must be meticulously followed in the correct order and precise amounts. However, an algorithm is usually a high-level, more abstract version of a recipe, where the instructions are represented generally—usually in math—rather than in the



YOUR HONOR, I'D LIKE TO
MOVE FOR A CONTINUANCE, MY
SERVER IS DOWN

language of a particular computer program, like Javascript or Python. Programmers then follow this high-level algorithm recipe and create an actual computer program that carries out the steps, but using a particular programming language, such as Python, Javascript, or C. The same high-level algorithm can be written or “implemented” many different ways, in different computer programming languages, as long as the computer programs follow the precise instructions and contours of the algorithm's math. So, we can think of the algorithm as the higher-level, but still precise, description of the process, and the computer program as a practical way to actually carry out the algorithm's process. Sometimes, as shorthand, people talk about a particular implementation of an algorithm in a particular computer program as the algorithm itself.

InQ: What is “artificial intelligence”?

Prof. Surden: There is probably no one definition of artificial intelligence that everyone would agree with. But a definition that I find useful is the following: Using computers to solve problems, make predictions, answer questions, or make automated decisions or actions, on tasks that when done by people, typically require “intelligence.”

There is also no one definition of intelligence that people will agree with, but for our purposes, we can think of it loosely as higher-order cognitive skills—such as abstract reasoning, problem solving, use of language, learning, and visual processing—that are associated with advanced human thinking.

Thus, for example, there are a number of activities that people do that are thought to involve many of these higher-order processes, such as playing chess, solving problems, driving a car, reading, discussing philosophy, and writing. When we use a computer to solve any one of these tasks—that in humans are associated with higher-order cognitive functions—it is called an artificial intelligence task. But AI computer systems accomplish these tasks very differently than humans do.

InQ: What about the term “machine learning”? What does that mean?

Prof. Surden: Machine learning is a way of creating AI systems in which the computer learns useful patterns from data. This is in contrast to people manually creating rules for the computer to follow. A good example of machine learning is email spam detection. We can think of a few ways to detect email spam. One way would be to manually craft a list of words that we think are associated with spam based on our personal experience, such as “free” or “award.” However, this is a “brittle” approach that will not cover every case and will not adapt over time. A better way—and the way it works today in most cases—is to have an AI machine learning system “learn” what spam looks like by analyzing emails for patterns. So, instead of giving the computer a list of words, we instead click the “spam” button on our email systems to indicate to the system that we think a particular email is spam. This is, in effect, giving the email program an example of spam that it can scan for patterns. On its own, it has an algorithm that is designed to spot words, or other features, that appear unusually frequently in spam emails versus wanted emails. So, machine learning in this context is giving the algorithm examples of what we are interested in, and having it “learn” patterns from those examples, usually using statistics, rather than manually crafting the rules.

InQ: Can you compare machine learning to “formal rule representations” for us?

Prof. Surden: In a certain way, formal rule representation is the opposite of machine learning. It involves people with expertise crafting precise rules about the way things work. However, these rules are created in such a way that a computer can process them and check for violations. This has also been a very successful approach in certain other areas. A good example of this is tax preparation software, which involves manually created rules.

InQ: In researching this article, I have run across the phrase “having humans in the loop.” What does that phrase mean?

Prof. Surden: There are many areas that involve judgment and estimates. In a computer program, there are a few options when we encounter one of these. One option is to just let the computer make its best guess, using its algorithms, and then continue on. Another approach is to pause and send the decision to a human to weigh in on and possibly ultimately decide the matter. This is an example of having humans in the loop. A good example of this involves an airplane’s autopilot, where the airplane’s automation does some of the assessment in terms of takeoff and landing but much of the time the end judgment remains in the pilot’s hands, in terms of making the ultimate flying decisions. The human pilot is “in the loop” rather than having the plane’s flying being fully automated.

InQ: Can you contrast “strong AI” and “artificial general intelligence,” or “AGI,” for the readers?

Prof. Surden: Strong AI is the aspiration—which does not exist yet—of computer systems that could meet or exceed the level of human intelligence across all areas. By contrast, today, even the most advanced computer large language systems, such as GPT-4, are not quite at the level of humans across all fields of human endeavor. Thus, strong AI is currently fictional and something we only see in entertainment, such as the C-3PO robot in *Star Wars*. Researchers are uncertain if, or when, we will ever achieve strong AI.

InQ: Last, can you explain what’s meant by “AI chat” or “chat bot programs,” which I will

refer to during our discussion as “chat AI” and which have garnered a lot of press of late due to programs like ChatGPT.

Prof. Surden: ChatGPT is a chat-based interface to an underlying technology known as GPT,¹⁵ made by a company called OpenAI. It is an extremely advanced version of an AI technology known as a large language model. Essentially, ChatGPT is a huge breakthrough in AI that occurred last year in 2022, which allows language models to reason, solve problems, and answer questions at near-human—or in

Amazingly, advanced versions of these systems, particularly the most recent GPT-4, are also able to do problem solving and reasoning. This was quite unexpected to most AI researchers, including me.

some cases, above-human—ability. There is a free version of ChatGPT available to the public known as 3.5, and then there is a state-of-the-art advanced version called GPT-4, just released in March 2023, which is incredible. It is much improved over ChatGPT-3.5.

InQ: Can you explain to our readers what chat AI programs like ChatGPT are, what they do, and how they work?

Prof. Surden: Essentially, programs like ChatGPT are machine-learning programs,

based on an approach using deep-learning neural networks, that read vast amounts of text, huge portions of the internet, books, and so on. By analyzing so much human-produced text, they learn the fundamental patterns underlying human language and can produce human-like text. They are fundamentally text generators and basically just predict the next word, based on what has been asked in the “prompt” (e.g., “Write me a poem”) as well as what the system has written so far (e.g., “The cat sat . . .”). The program uses the prompt, plus what it’s already written, to predict the next word. Amazingly, advanced versions of these systems, particularly the most recent GPT-4, are also able to do problem solving and reasoning. This was quite unexpected to most AI researchers, including me. The ability to solve problems seems to be an “emergent” and unexpected property of making these programs so big, and due to very excellent engineering on the part of OpenAI. These AI advances represent a huge leap in the state of the art compared to just last year, 2022, when AI systems could not reliably understand what was asked of them, nor reliably follow instructions and produce useful results, the way that GPT-4 can.

InQ: The more I read about chat AI programs, the more I have come to believe these programs are “dumb,” riddled with hallucinations—falsehoods—and yet they can appear very “smart” to someone using the program. Are these programs smart, dumb, or somewhere in between?

Prof. Surden: I hesitate to answer that, because I think that it risks comparing them too much to humans, which they are decidedly not. I frame them as “useful” or “not useful.” ChatGPT-3.5 released in November 2022, was the first general AI system, in my opinion, that was actually useful across a wide variety of tasks. GPT-4, released in March 2023, was even more useful. They are getting better, as the technology improves, in giving more accurate and more reliable answers, and the difference between GPT-3.5 and GPT-4 is an example of this. I expect these types of programs to only get better, in terms of usefulness and accuracy, over time.

InQ: Do you believe that current chat AI technology, similar to ChatGPT, could write a

legal brief suitable for filing in court? If not, how long do you think it will be until this is possible?

Prof. Surden: Based on my research, the free GPT-3.5 is not quite up to the task. The more advanced GPT-4, however, is capable of producing a good first draft of a legal motion. However, you definitely would not want to file it directly in court. Rather, it would need to be double-checked for errors and subject to additional reasoning and analysis by humans. I would not recommend using GPT-4 directly to do this currently. There are legal tech systems, such as CaseText, that use GPT-4 in the background, but they are built by lawyers and have privacy safeguards. That is a better way to go in my opinion.

InQ: Would writing such a brief always require significant human collaboration, or do you expect AI to reach the point where the program by itself could scan a motion and supporting brief and generate a top-notch file-worthy response?

Prof. Surden: I think for certain basic, non-complicated legal cases, we're not far

from the day where a technology similar to GPT-4 can create a solid first draft of a motion that can, with significant double-checking and additional analysis, be ready to file. I think for more complicated cases that form the backbone of many law practices, these technologies should be treated as "first-draft" machines rather than fully fledged motion-producing products.

InQ: If we reach the point where AI could generate a legal brief that, to a reader—like a judge—is as well-researched and persuasive as one generated by a skilled attorney, producing as good or better legal outcomes, is there any reason why such briefs should not be used rather than human-generated briefs? Would your answer change if AI could produce a brief as good as that written by our finest legal minds? Is so, might it be legal malpractice for an attorney not to employ the best AI brief-writing program?

Prof. Surden: I think it depends on how good and accurate the automatically produced briefs are, and how complicated the case is. It is hard to know at this point, but for now, I

would still want lawyers using this technology only as a first draft.

InQ: What about the ethics of lawyers using AI to write briefs? Does such activity run headlong into various legal ethics rules? For example, current rules generally require lawyers to identify adverse controlling legal authority; not to misstate facts or the law; not to mislead the court; and, to discharge their duties consistent with their capacity as officers of the court.¹⁶ Are these rules going to need to change and, if so, how?

Prof. Surden: I would not recommend that lawyers use ChatGPT-3.5 or GPT-4 directly today. If they enter in private or sensitive client information to OpenAI, they could run afoul of ethical privacy rules, and may even effect a waiver of confidentiality privileges.¹⁷ Rather, I recommend using a legal technology service that is built upon AI but has privacy and security guarantees on top of it. I think the technology needs some good, independent testing before we can advocate using it more broadly. Also, attorneys need to be sure to double-check

DENVER BAR ASSOCIATION PLACEMENT SERVICE

Serving the legal community for over 30 years.

**Temporary, Temp-to-Hire and Permanent
Placement of Legal Support Staff**

Mev Parsons
1290 Broadway, Ste. 1700
Denver, CO 80203
303-894-0014
dbaps@earthlink.net

denbar.org/Placement-Services

AI work product the same way they would double-check the product of a junior associate.

InQ: If the ethical rules don't change, wouldn't lawyers then need to duplicate all the work their AI brief-writing programs saved them from doing so those lawyers could then certify that the brief meets all the ethical rules applicable to court filings? If the rules don't change, will that hamstring lawyers from producing the best briefs on behalf of their clients, perhaps at a much lower cost?

Prof. Surden: Again, I hesitate to answer, because the technology is so new. I think we need reliable, independent testing that compares the quality of this technology to the quality of attorney output. Until we have that, it's too early to tell.

Let's Go to the Source

InQ: ChatGPT-3.5—I hope you don't mind if I just call you ChatGPT. You're an enormously popular program—you reached 100 million users in two months, the fastest app on record. In comparison, TikTok took nine months, while Instagram took 30.¹⁸

What's your take on the kinds of ethical issues lawyers might face when collaborating with you?



ChatGPT: As generative AI becomes more prevalent in the legal industry, lawyers may face various ethical issues related to its use in their law practices.

Some of the potential ethical issues include:

1. Confidentiality and Privacy: Lawyers have a duty to maintain confidentiality and protect their clients' privacy. The use of generative AI raises concerns about data privacy, particularly when it comes to sensitive information that may be processed by AI systems. Lawyers must ensure that their AI systems are secure, and client data is not misused or accessed by unauthorized parties.

2. Bias and Fairness: AI systems are only as unbiased as the data they are trained on. Lawyers must ensure that the data used to train their AI models is fair and does not perpetuate discrimination. Lawyers must also monitor their AI systems for potential biases and take corrective actions when necessary.

3. Accountability: The use of AI raises questions about accountability when errors or mistakes are made. Lawyers must be able to explain how their AI systems work and be prepared to take responsibility when their systems make mistakes.

4. Professionalism: Lawyers must use AI in a manner that is consistent with their professional obligations. They must avoid delegating tasks to AI systems that require human judgment or skills. Lawyers must also ensure that their use of AI does not undermine their professional responsibility to act in their client's best interests.

5. Informed Consent: Lawyers must obtain informed consent from their clients before using AI systems in their cases. Clients must be made aware of the limitations and potential risks associated with the use of AI in legal practice.

6. Unauthorized Practice of Law: The use of AI may give rise to unauthorized practice of law issues. Lawyers must ensure that their AI systems are not performing tasks that should be performed by licensed attorneys.

Overall, lawyers must ensure that their use of generative AI is consistent with their ethical obligations and does not compromise their clients' interests or privacy.

InQ: You know ChatGPT, one thing I like about your answers to my questions is that they are succinct and fast—it took you less than two seconds to “think” about, “analyze,” and “respond” to my last question. But, of course—and no offense intended—you cannot think or analyze. You just run your algorithm and predict what words should come next in sequence based on the trillions of data points you have scraped from the Web. Still, pretty impressive—you almost had me fooled.

Do you “think” there's a special and unique place for generative AI brief-writing programs in the context of helping underserved populations achieve greater access to justice?

ChatGPT: Generative AI brief-writing programs have the potential to improve access to justice for underserved populations. These programs can automate the creation of legal documents, including briefs, which can be time-consuming and expensive to produce. By reducing the time and cost associated with legal document creation, these programs can


help to level the playing field for individuals and communities who may not have the resources to pay for expensive legal services. Additionally, generative AI brief-writing programs can help to ensure that legal documents are created in a more consistent and accurate manner. This can be particularly important in contexts where legal language is complex and nuanced, and where even small errors or inconsistencies can have significant consequences. However, it is important to note that generative AI brief-writing programs should not be seen as a replacement for human legal professionals. While these programs can be helpful in automating certain tasks, they cannot provide the same level of nuanced legal analysis and advice that a human attorney can. As such, it is important to view generative AI brief-writing programs as a tool that can be used in conjunction with human legal professionals to improve access to justice for underserved populations.

InQ: What negative developments might be associated with such a development—would small claims court disputes turn into a mind-boggling battle of the AI brief-writing programs?

ChatGPT: One potential negative development associated with generative AI brief-writing programs is that they could exacerbate existing power imbalances in the legal system. While these programs have the potential to reduce costs and improve efficiency, they may also create a situation where only those who can afford to invest in the technology have access to the best legal documents. This could lead to a situation where wealthier litigants are better able to advocate for their interests, while those who cannot afford the technology are left at a disadvantage.

Another potential negative development is that AI brief-writing programs could be used to flood the court system with low-quality or frivolous briefs. This could be particularly problematic in small claims court, where disputes may already be less complex and time-consuming. If litigants begin to rely heavily on generative AI brief-writing programs to create their legal documents, it could lead to a situation where the court is inundated with poorly crafted and irrelevant briefs, which could further slow down an already overburdened system.

0W2023



COUNSEL, I AM NOT FAMILIAR
WITH YOUR CITATIONS FOR KRAMER
VS KRAMER AND BRONCOS VS SEAHAWKS

Furthermore, AI brief-writing programs could be used to create deceptive or manipulative briefs, which could be difficult for judges or other legal professionals to detect. This could lead to a situation where justice is not served, as litigants are able to win cases based on misleading or inaccurate information presented in their briefs.

Overall, while generative AI brief-writing programs have the potential to improve access to justice, it is important to carefully consider the potential negative developments that could arise as a result of their use, and to take steps to mitigate these risks.

InQ: Those risks certainly exist, although I think our adversary system should help lessen those concerns. Take the May 2023 Avianca

Airlines case, where a 30-year veteran lawyer, whose work was being supervised by another lawyer, used ChatGPT to help write a brief that was riddled with fabricated citations.¹⁹ A federal judge came down really hard on the lawyer, and I think the surrounding publicity constituted a flashing red caution sign. I also think that lawyer's actions implicated nearly every ethics concern you raised plus one more: competence.

ChatGPT, might we reach the point where lawyers may have to choose among various brief-writing programs, each heavily informed by the writing styles and intellects of well-known US Supreme Court advocates, but few programs are offered that are drawn from the writing of female advocates or advocates of color because these persons rarely appear in the Supreme Court?

ChatGPT: As you note, female advocates and advocates of color have historically been underrepresented in the US Supreme Court, which could limit the availability of data needed to train generative AI brief-writing programs that are informed by their writing styles and approaches. This could lead to a situation where the programs that are available are biased towards the writing styles and approaches of male or white advocates, which could further exacerbate existing inequities in the legal system.

To mitigate these risks, it will be important for those involved in the development of generative AI brief-writing programs to actively seek out and incorporate diverse perspectives and approaches. This could involve working with a range of legal professionals from diverse backgrounds, and taking steps to address the biases that are inherent in the technology. Overall, the development of generative AI brief-writing programs that are informed by the writing styles and approaches of various Supreme Court advocates could be a valuable tool for lawyers. However, it will be important to ensure that these programs are developed in a way that is inclusive and equitable, and that they do not further entrench existing power imbalances in the legal system.

InQ: What about brief-writing programs targeted at certain jurists: programs that write excellent briefs and also try to predict which arguments are most likely to persuade certain jurists? What problems do you think might emerge in this context?

ChatGPT: There are a number of potential problems that could emerge in this context. For one, there is the risk that these programs could reinforce biases that are already present in the legal system. If the programs are trained on data that is biased towards certain judges or justices, they may perpetuate these biases by recommending arguments that are more likely to be persuasive to those judges or justices, regardless of whether those arguments are actually legally sound or justified. Furthermore, there is the risk that these programs could create a feedback loop, where judges or justices begin to expect certain types of arguments or writing styles from lawyers. This could create a situation where lawyers are incentivized to

conform to certain expectations, rather than pushing for novel or creative legal arguments.

InQ: If AI assumes more and more lawyer-like functions, what should we do with the surplus lawyers and all those wannabe lawyers—law and college students?

ChatGPT: To address this potential surplus of legal professionals, there are several strategies that could be considered. One option is to encourage the development of new legal services that are focused on areas that are less likely to be impacted by the rise of AI, such as areas that require significant human interaction or judgment.²⁰ This could include areas such as family law, dispute resolution, or legal counseling, which may be less susceptible to automation. Another strategy is to encourage legal professionals to develop new skills and areas of expertise that are complementary to the capabilities of AI. For example, lawyers could focus on developing their skills in areas such as data analytics, project management, or client relations, which may be in high demand in a world where AI is increasingly prevalent.

In addition to these strategies, it will also be important to provide support and resources to legal professionals who are impacted by the rise of AI. This could include programs to help lawyers transition to new careers or to develop new skills, as well as support for law schools and other legal education programs to adapt to the changing demands of the legal profession. Overall, as AI assumes more lawyer-like functions, it will be important for the legal profession to adapt and evolve to meet the changing demands of the marketplace. While there may be a surplus of lawyers and law graduates in certain areas, there will also be new opportunities and challenges that emerge as a result of these changes, and it will be important to take a proactive approach to managing this transition.

InQ: Well, this has been an amazing and enlightening interview. One last question: We've dealt exclusively with one another over the Internet. Here is a Turing test²¹ of sorts: How do you know I'm not a generative AI program—have I manifested any obvious "tells?"

ChatGPT: Based on the content of our conversation and the nature of your questions, it

appears that you are a human being rather than a generative AI program. You have demonstrated a level of understanding and engagement that suggests human cognition, and your questions have been thoughtful and nuanced in a way that is not typically seen in AI-generated text. Nonetheless, I am just a machine and cannot claim to know for sure.

Conclusion

"The future is coming, but it will not be as fast as some predict."²²

It may seem like the world of science fiction, but AI capabilities are improving rapidly, bringing change. Near term, AI is likely to enhance our lawyering skills and improve our work-product and efficiency. But, certainly, significant transformations are coming and it is hard to predict when they will happen and what they will look like. And it is not just lawyers (and legal assistants and law clerks) looking over their shoulders. Radiologists have their eyes on AI as well, as programs now can detect breast cancer as well or better than doctors. Still, doctors take solace that "[a]n A.I.-plus-doctor should replace doctor alone, but an A.I. should not replace the doctor," and that "the technology will be effective and trusted by patients only if it is used in partnership with trained doctors."²³ These observations would seem to apply to lawyers as well.

Lurking in the shadow of AI improvements is the metaverse, a vision of the Internet's next evolutionary step—a singular, shared, immer-

sive, persistent, three-dimensional virtual space where lawyers, judges, witnesses, and observers might each be sitting in the comfort of their homes, adorned with headsets or surrounded by holographic imaging, attending meetings, depositions, hearings, trials, and appellate arguments.²⁴ This massively scaled metaverse will likely include an interoperable network of real-time rendered 3D virtual worlds that can be experienced synchronously and persistently by a nearly unlimited number of users with an individual sense of presence and with continuity of data, such as identity, history, entitlements, objects, communications, and payments.²⁵ Future shock may be waiting for all attorneys just outside the door: "too much change in too short a period of time."²⁶

If AI starts to creep into your consciousness late at night, stirring a worry you can't quite put your finger on, read Ted Chiang's "ChatGPT is a Blurry JPEG of the Web,"²⁷ which digs deeply into, in an understandable way, the significant limitations of ChatGPT and similar chat AI programs.²⁸ And take comfort in the fact that the current practicing bar's future probably will involve chat AI augmenting lawyers' skills, providing an inexpensive tool that will save time and money, producing better and more creative and collaborative work product, helping minimize unconscious and structural biases, and expanding access to justice.²⁹ Still, if someone suggests you obtain a cognitive implant to speed communication between your mind and some future chat AI program, proceed cautiously.³⁰ **CL**



Ronald M. Sandgrund is of counsel with the construction defect group of Burg Simpson Eldredge Hersh Jardine PC. The group represents commercial and residential property owners, homeowner associations and unit owners, and construction professionals and insurers in construction defect, product liability, and insurance coverage disputes. He is a frequent author and lecturer on these topics, as well on the practical aspects of being a lawyer. He has taught entrepreneurial innovation and public policy and trial advocacy, and has lectured on legal ethics, construction law, mass tort litigation, consumer rights, and other subjects at Colorado Law. He is, as far as he is aware, not an artificial life form.



About the cartoonist: Daniel Walter is an LA-based composer and musician. Collaborating with prominent filmmakers such as Ari Aster and Scott Aukerman, his compositions have featured in projects premiering at festivals from Cannes to SXSW. Walter has guest-lectured on film scoring at UCLA and received recognition through the ASCAP Film and TV Scoring Workshop, a Jerry Goldsmith Award nomination, and a composer residency in Alaska. He has scored commercials that played at the Super Bowl and Olympics, as well as movie trailers for all the major studios. He has also produced a feature film, *The Tenant*, which sold to Sky UK, and a short film, *Picture Day*, which won a special jury award at the Palm Springs ShortFest in 2022 and the Dallas Film Festival in 2023—www.danielwaltermusic.com.

NOTES

1. Tippet and Alexander, “Robots are coming for lawyers—which may be bad for tomorrow’s lawyers but great for anyone in need of cheap legal assistance,” *The Conversation* (Aug. 9, 2021), <https://theconversation.com/robots-are-coming-for-the-lawyers-which-may-be-bad-for-tomorrows-attorneys-but-great-for-anyone-in-need-of-cheap-legal-assistance-157574>.
2. Lohr, “A.I. is Coming for Lawyers, Again,” *N.Y. Times* (Apr. 10, 2023), <https://www.nytimes.com/2023/04/10/technology/ai-is-coming-for-lawyers-again.html>.
3. Greene, “Will ChatGPT make lawyers obsolete? (Hint: be afraid),” *Reuters* (Dec. 9, 2022), <https://www.reuters.com/legal/transactional/will-chatgpt-make-lawyers-obsolete-hint-be-afraid-2022-12-09>.
4. Polzin, “The Rise of AI: Why Legal Professionals Must Adopt or Risk Being Left Behind,” *Entrepreneur* (Jan. 11, 2023), <https://www.entrepreneur.com/science-technology/why-legal-professionals-must-adapt-to-ai-or-risk-being-left/441751>.
5. Heinen, “ChatGPT as a Replacement for Human Lawyers,” *Foley & Lardner LLP blog* (Jan. 5, 2023), <https://viewpoints.foley.com/post/102i4ho/chatgpt-as-a-replacement-for-human-lawyers>.
6. Sahota, “Will A.I. Put Lawyers Out of Business?” *Forbes* (Feb. 9, 2019), <https://www.forbes.com/sites/cognitiveworld/2019/02/09/will-a-i-put-lawyers-out-of-business/?sh=48da5c9631f0>.
7. Patrice, “AI Won’t Replace All Lawyers . . . Just the Lazy Ones,” *Above the Law* (Feb. 13, 2023), <https://abovethelaw.com/legal-innovation-center/2023/02/13/ai-wont-replace-all-lawyers-just-the-lazy-ones>.
8. See Sahota, *supra* note 6 (quoting Tom Girardi, the lawyer who inspired the movie *Erin Brockovich*).
9. Attributed to Chicken Little, a very nervous fowl.
10. Those laws are: (1) A robot may not injure a human being or, through inaction, allow a human being to come to harm; (2) A robot must obey orders given it by human beings except where such orders would conflict with the First Law; and (3) A robot must protect its own existence as long as such protection does not conflict with the First or Second Law.
11. Skynet is a fictional artificial neural network conscious group mind and artificial general super-intelligence system.
12. HAL (Heuristically programmed ALgorithmic computer) is from Stanley Kubrick’s movie *2001: A Space Odyssey*, based on the writings of Arthur C. Clarke.
13. See Sahota, *supra* note 6 (“[I]t’s been demonstrated people are more likely to be honest with a machine than with a person, since a machine isn’t capable of judgment.”).
14. Okay, I will concede that nothing man creates is untainted by bias.
15. “GPT” means generative pre-trained transformer.
16. See, e.g., Colo. RPC 3.3, cmt. 1.
17. See Colo. RPC 1.6.
18. Grant, “ChatGPT is Causing a Stock-Market Ruckus,” *Wall St. J.* (May 9, 2023), <https://www.wsj.com/articles/chatgpt-is-causing-a-stock-market-ruckus-4b7cc008>.
19. See Weiser, “Here’s What Happens When Your Lawyer Uses ChatGPT,” *N.Y. Times* (May 27, 2023), <https://www.nytimes.com/2023/05/27/nyregion/avianca-airline-lawsuit-chatgpt.html>.
20. Philosopher Yuval Harari argues that chat AI programs, “[t]hrough [their] mastery of language,” are on the verge of forming “intimate relationships with people, and us[ing] the power of intimacy to change our opinions and worldviews.” *Economist* (Apr. 28, 2023), <https://www.economist.com/by-invitation/2023/04/28/yuval-noah-harari-argues-that-ai-has-hacked-the-operating-system-of-human-civilisation>.
21. [ChatGPT, in response to my query]: “The Turing test is a test of a machine’s ability to exhibit intelligent behavior equivalent to, or indistinguishable from, that of a human. In the Turing test, a human evaluator engages in a natural language conversation with a machine and another human, without knowing which is which. If the evaluator is unable to distinguish the machine from the human based on the conversation alone, then the machine is said to have passed the Turing test. The Turing test is often used as a benchmark for evaluating the development of artificial intelligence, as it provides a measure of a machine’s ability to exhibit human-like intelligence and behavior.”
22. Lohr, *supra* note 2 (quoting Raj Goyle).
23. Satariano and Metz, “Using A.I. to Detect Breast Cancer that Doctors Miss,” *N.Y. Times* (Mar. 5, 2023), <https://www.nytimes.com/2023/03/05/technology/artificial-intelligence-breast-cancer-detection.html>.
24. This imagined metaverse is fairly benign. However, it could be more like in the movie *The Matrix*, where human minds are consigned to “live” in a simulated reality while the intelligent machines that put them in this stasis harvest the subjugated humans’ bodies for bio-electric energy to keep the machines running. The machines did this because the humans tried to shut them off by denying them access to solar energy. https://en.wikipedia.org/wiki/The_Matrix.
25. Ball, *The Metaverse and How it Will Revolutionize Everything* (Liveright Publ’g Corp. 2022).
26. Toffler, *Future Shock* (Random House 1970).
27. *New Yorker* (Feb. 9, 2023), <https://www.newyorker.com/tech/annals-of-technology/chatgpt-is-a-blurry-jpeg-of-the-web>. Of course, the next generation of AI may simply kill all of humankind in their sleep. Youdowsky, “Pausing AI Development Isn’t Enough. We Need to Shut it All Down,” *Time* (Mar. 29, 2023), <https://time.com/6266923/ai-eliezer-youdowsky-open-letter-not-enough>. Or it may hallucinate and make up career-ending lies. Verma and Oremus, “ChatGPT Invented a Sexual Harassment Scandal and Named a Real Law Prof as the Accused,” *Wash. Post* (Apr. 5, 2023), <https://www.washingtonpost.com/technology/2023/04/05/chatgpt-lies>. And, frankly, the energy requirements of training and running generative AI may be multiples of the energy needs of all those Bitcoin miners—and so, if you are believer in climate chaos, our demise might be found simply in running these AI programs. See Ludvigsen, “ChatGPT’s Electricity Consumption,” *Towards Data Science blog* (Mar. 1, 2023), <https://towardsdatascience.com/chatgpts-electricity-consumption-7873483feac4> (suggesting that ChatGPT is dumber than a doornail, and cannot “create,” but “imitates” being smart).
28. Newport, “What Kind of Mind Does ChatGPT Have?” *New Yorker* (Apr. 13, 2023), <https://www.newyorker.com/science/annals-of-artificial-intelligence/what-kind-of-mind-does-chatgpt-have>.
29. While lawyers’ jobs may be safe in the short term, the prospects for legal assistants may be dimmed. See Metz, “‘The Godfather of A.I.’ Quits Google and Warns of Danger Ahead,” *N.Y. Times* (May 1, 2003), <https://www.nytimes.com/2023/05/01/technology/ai-google-chatbot-engineer-quits-hinton.html> (chatbots “tend to complement human workers, but they could replace paralegals . . .”); Schell, “Which Jobs Will be Most Impacted by ChatGPT?” *Visual Capitalist* (Apr. 30, 2023), <https://www.visualcapitalist.com/cp/which-jobs-artificial-intelligence-gpt-impact> (legal assistants 100% exposed to AI—AI expected to reduce their tasks by 50% or more).
30. Or maybe you can dispense with the cognitive implants now that AI programs can translate your thoughts into words. See Metz, “A.I. is Getting Better at Mind-Reading,” *N.Y. Times* (May 1, 2023), <https://www.nytimes.com/2023/05/01/science/ai-speech-language.html>.